CRITICAL THINKING

FOR INTELLIGENCE ANALYSIS:

Connecting the Dots

Gheorghe Tecuci and David Schum

To Susan Durham

Table of Contents

PREFACE

Book Purpose

This textbook has been written for those studying the process of drawing conclusions from masses of evidence resulting from extensive investigations in a variety of contexts, including intelligence analysis, law, cybersecurity, criminal investigations, and military and business inferences and decisions. Many universities now offer undergraduate and graduate courses concerning these activities. These courses are offered in order to provide introductory preparation for persons contemplating future work in these contexts. These courses have also been of interest to persons having various levels of past experience in these activities, but who are seeking additional knowledge concerning matters their current work requires.

Our major objective in this book is to provide accurate, useful, and extensive information about the evidential and inferential issues encountered by persons whose tasks require them to *draw defensible and persuasive conclusions from masses of evidence of* all kinds that come from a variety of different sources.

The “*connecting the dots*” metaphor that illustrates the extraordinary complexity of these evidential and inferential reasoning issues, may have gained its current popularity following the terrorist attacks in New York City and Washington, DC, on September 11, 2001. It was frequently said that the intelligence services did not connect the dots appropriately in order to have possibly prevented the catastrophes that occurred. Since then, we have seen and heard this metaphor applied in the news media to inferences in a very wide array of contexts in addition to the aforementioned intelligence, legal, military, and business contexts. For example, we have seen it applied to allegedly faulty medical diagnoses; to allegedly faulty conclusions in historical studies; to allegedly faulty or unpopular governmental decisions; and in discussions involving the conclusions reached by competing politicians. What is also true is that the commentators or “talking heads” on television, radio, Internet, or the sources of written accounts of inferential failures, never tell us what they mean by the phrase “connecting the dots.” A natural explanation is that they have never even considered what this phrase means and what it might involve.

The truth is that analysts are routinely asked to perform tasks for which they have received little, if any, adequate tutoring. Many courses are indeed offered in various agencies. Frequently, however, such courses dwell mainly, or only, on the <u>production</u> of intelligence analyses and <u>not</u> on <u>the analytic process</u> itself. By "production" we simply mean the manner in which the results of an intelligence analysis are compiled and then displayed or "packaged." Imagine that someone is given a gift that comes in very attractive packaging. This person opens the box and there's either nothing in it, or what is in it is the wrong size, color, or style. In any case, what is in the box is not pleasing to the recipient. Many intelligence analyses come in very attractive packaging in the form of carefully constructed Powerpoint or other visual presentations. But what they display are arguments and conclusions that cannot always be defended when they are subjected to criticism from other, often very well informed, persons who must make decisions based on intelligence analyses. In addition, many intelligence analysis courses mainly involve accounts, in the form of "war stories," provided by analysts whose experience qualifies them to describe their own analytic methods and results.

Conventional courses in logic, probability, and statistics do not prepare a person for the task of drawing conclusions based on masses of evidence whose items suggest many, often complex and interrelated, lines of arguments on hypotheses of interest. The evidence of interest to intelligence analysts usually concerns events that are unique, singular, or one-of-a-kind, and are thus not replicable or repeatable. This means that there are rarely any useful or relevant statistical records available to draw upon in making inferences about the capabilities and intentions of potential or real adversaries. We had no existing statistical records regarding the intentions of foreigners who showed up in our civilian flying schools wishing only to learn how to steer multiengine aircraft and not how to perform takeoffs or landings.

Lacking existing statistical records to draw upon, the analysts must generate new information by inquiry, the asking of questions. What is therefore needed is formal training in the basic analytic tasks of *generating novel and productive hypotheses* and in the *construction of defensible and persuasive arguments based on masses of evidence*. But these two major requirements are exactly what this book is all about.

How to Use the Book

This book may be regarded as a new, expanded, and improved edition of our previous book, "I*ntelligence Analysis as Dicovery of Hypoheses, Evidence and Arguments: Connecting the Dots",* pulished by Cambride University Press in 2016. In reflects the progress made in our research over the past 10 years, in our long-term effort to improve intelligence analysis, an objective that was so much advocated by William Nolte (2004) and Stephen Marrin (2011). New material has ben added, the Disciple-CD analytical tool was replaced with Cogent, a much mor powerful and easier to use tool, many new exercises have been added, and solutions to all of them have been provided.

As the previous book, we have structured it to be used by a wide variety of users with different prior backgrounds, training, and interests. This allows the book to be used either as the main textbook for entire courses on intelligence analysis, or as a textbook for a part of a course, based on the desired coverage of intelligence analysis topics (introduction, basic topics, advanced topics), and the desired level of use of Cogent (no use, demonstration of use, actual use).

Introduction to Critical Thinking: Chapter 1-2
Critical Thinking: Chapters 1-3
Critical Thinking with Cogent: Chapter 1-3, 6
Advanced Critical Thinking: Chapter 1-6

Naturally, we hope that your learning venture with or without the assistance of Cogent will be a most valuable experience in which you discover many very important elements of intelligence analysis, some of which you might not have heard anything about before. We also hope that this venture will be directly relevant to tasks you face, or will face, every day in your analytic careers. Finally, we hope that it will be as enjoyable as it is informative. So, as you begin this learning venture, we wish you **Bon Voyage!**

Acknowledgments

We are most endebted to Susan Durham, to whom this book is dedicated. It was her support that allowed us to start this line of research, and than revive it. We are also grateful to William Nolte, Joan McIntyre, John Donelan, and Cindy Ayers for their support and encouragement. Phillip Hwang, Joan Vallancewhitacre, John Greer, Michelle Quirk, and Steven Rieber also provided support to our research.

Our colleagues and students from the Learning Agents Center contributed to the research on which this book is based, particularly Dorin Marcu, who was the main developer of Cogent, and Louis Kaiser who contributed to the assessment of relevance and to many of the review questions.

About the Authors

Gheorghe Tecuci (Ph.D., University of Paris-Sud and Polytechnic University of Bucharest, 1988) is Professor of Computer Science and Director of the Learning Agents Center in the School of Computing of George Mason University, Member of the Romanian Academy, and former Chair of Artificial Intelligence in the Center for Strategic Leadership of the U.S. Army War College. He has followed a career-long interest in the development of a computational theory and technology allowing non–computer scientists to develop cognitive agents that incorporate their problem-solving expertise and can act as cognitive assistants to experts, as expert consultants to nonexperts, and as intelligent tutors to students. He has published over 200 papers, including 11 books, with contributions to artificial intelligence, knowledge engineering, cognitive assistants, machine learning, evidence-based reasoning, and intelligence analysis. He has received the U.S. Army Outstanding Civilian Service Medal (for “groundbreaking contributions to the application of artificial intelligence to center of gravity determination”) and the Innovative Application Award from the American Association for Artificial Intelligence.

David A. Schum (Ph.D., Ohio State University, 1964) was Professor of Systems Engineering, Operations Research, and Law, as well as Chief Scientist of the Learning Agents Center at George Mason University, and Honorary Professor of Evidence Science at University College London. He followed a career-long interest in the study of the properties, uses, discovery, and marshaling of evidence in probabilistic reasoning. His major lines of research have involved the tracking of evidential and inferential subtleties in complex inference; the design of various strategies for assisting persons in the performance of complex inference tasks; the study of the task of assessing the inferential force of various forms and combinations of evidence; and study of ways to enhance the process of discovery of new ideas and their relevant evidential tests. Dr. Schum has published over 100 papers in a variety of journals, and eight books, including *The Evidential Foundations of Probabilistic Reasoning, Analysis of Evidence, Evidence and Inference for the Intelligence Analyst,* and *Probabilistic Analysis of the Sacco and Vanzetti Evidence*, being recognized as one of the founding fathers of the Science of Evidence.

1 INTRODUCTION

1.1 Critical Thinking

Critical thinking is a complex concept that was developed over the past 2,500 years through the work of some of the greatest minds, including Aristotle (384BC–322BC), Galileo Galilei (1564–1642), Isaac Newton (1642–1727), John Locke (1632–1704), William Whewell (1794–1866), Charles Peirce (1839–1914), and John Wigmore (1863–1943), who have tried to understand and reason about the world through a process of discovery and testing of hypotheses based on evidence.

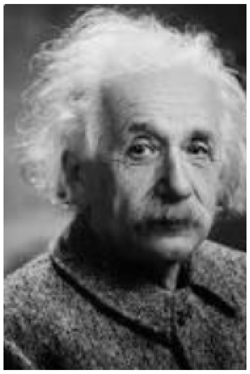

Albert Einstein
Thinking is hard work; that's why so few do

Critical thinking is now defined as the intellectually disciplined process of actively and skillfully conceptualizing, applying, analyzing, synthesizing, and/or evaluating information gathered from, or generated by, observation, experience, reflection, reasoning, or communication, as a guide to belief and action. Critical thinking is incorporated into scientific thinking, mathematical thinking, historical thinking, anthropological thinking, economic thinking, moral thinking, and philosophical thinking (Scriven and Paul, 1987). It is at the core of the problem-solving and decision-making tasks in a wide variety of disciplines, including military science and intelligence, computing, natural and social sciences, education, agriculture, and medicine. Critical thinking requires a skillful combination of imaginative and critical reasoning.

Developmets in artificial intelligence have made it posssible to build intelligent computer systems that may aid wih critical thinking, based on the complementariness between humans and computers reasoning process. Humans are slow, sloppy, forgetful, implicit, and subjective, but have common sense and intuition, and may find creative solutions in new situations. In contrast, computer systems are fast, rigorous, precise, explicit, and objective, but they lack common sense and the ability to deal with new situations (Turoff, 2007; Tecuci et al., 2007). Moreover, in contrast to a computer system, a human has a very limited attention span and can analyze only a small number of alternatives at a time (Pomerol and Adam, 2006). Such a system is *Cogent*, or cognitive agent for cogent analysis (Tecuci et al., 2015; 2018).

1.2 Imaginative Reasoning

We often hear it claimed that some people have imaginative reasoning capabilities and others don't. If you don't have it, you are out of luck. The truth of the matter is that nature has endowed <u>all of us</u> with imaginative reasoning capabilities (Howe, 1999). The trouble is that we are not always given the opportunity or encouragement to be imaginative or creative in our thinking. Our work on Cogent is based on the idea that you are naturally required to exercise your imagination in the act of trying to make sense out of the masses of evidence you will encounter. Our role in this process is to assist you in various ways. What *you* think about the evidence you will encounter is all-important. You may be able to assign possible meanings to evidence that others do not perceive. Another very important matter concerns how productive the exercises of our imaginations are. We all encounter persons who seem to be imaginative in the new ideas they generate. However, many of these same persons do not always generate new ideas that are helpful in the analytic tasks at hand. So, what needs to be encouraged in intelligence analysis is productively imaginative thought. But there are other ingredients necessary in efforts to help you become more like Sherlock or Mycroft Holmes than Inspector Lestrade, the famous characters of Conan Doyle.

The Cogent system we have developed can only assist you in several ways, and the rest depends on you and your analytic capabilities. You will be able to exercise your imaginative reasoning capabilities to the fullest only when you are *driven by curiosity or wonder* to find solutions to the analytic problems you encounter. If you don’t care whether anyone finds a solution to these problems, you stand very little chance of generating a productively imaginative solution. Experience in many areas has shown that the most productively imaginative persons are also those who have the greatest degree of commitment to find solutions to problems that confront them.

The final ingredient we mention here concerns the diligence with which you approach each new analytic problem you face. There is an old saying that fortune favors the prepared mind. Unless you have done your homework in the particular substantive areas your analytic problems involve, you also stand little chance of generating productively imaginative new ideas. Your brain requires something to work with; as we all learn, this requires burning the midnight oil. But being well acquainted only with the specifics of the substance of your analytic problems is often not quite enough. Productively imaginative persons usually also have a breadth of knowledge and experience to draw upon. Productive new ideas so often spring from the analogies we perceive; these analogies are often stated in the form of metaphors. But the forming of useful metaphors requires knowledge that goes beyond the boundaries of the believed substance of an analytic problem.

One of the most difficult problems we have faced in our work on Cogent is assisting you to construct defensible and persuasive arguments from a mass of evidence to hypotheses being considered. How well we are able to marshal our thoughts and evidence is vitally important in constructing defensible and persuasive arguments. The task of constructing arguments from a mass of different kinds of evidence is inherently difficult; perhaps it is the most difficult element of intelligence analysis. Though methods for performing complex argument construction have been around for a long time, such as the Wigmorean methods (Wigmore 19…), few people have made particular use of them until quite recently. In this volume we have combined concerns about these argument methods with concerns about thought and evidence marshaling.

Argument construction involves the interplay of imaginative and critical reasoning processes. As a result, different persons will imagine different reasoning routes from the same evidence to the same hypotheses. In addition, different persons may believe that the same body of evidence favors entirely different hypotheses. In short, *there is no such thing as a uniquely correct argument from some collection of evidence to hypotheses being entertained*. Add to this the fact that our evidence is always incomplete and any conclusion drawn today may have to be revised tomorrow in light of recently discovered evidence.

A final point concerns the argument construction methods themselves. The methods we discuss in connection with Cogent may appear overly compulsive and may seem to require "too much thought." One response here is to remind persons reading our works that careful intelligence analyses always require careful thought, regardless of what methods are being used. Using the methods we describe, we construct "pictures" of a complex argument in the form of what today are called inference networks. You may have had some exposure to the use of various software systems that now exist for the probabilistic analysis of inference networks. The trouble is that no such system tells the user how to construct an inference network appropriate in the analysis of some existing mass of evidence. These systems all assume that the imaginative and critical reasoning steps necessary in inference network construction have already been performed by the user. Having experience with the methods we discuss will offer analysts great assistance in seeing what is involved in the construction of defensible and persuasive arguments, regardless of whether you try to apply these methods in every analysis you undertake. Far too many persons are looking for a book entitled: *Intelligence Analysis Made Simple*. We do not see any hope for any *serious* book or course having this title. Intelligence analysis is an inherently difficult task; the methods we describe form one way of coping with the complexity of such tasks. Our Cogent system assists in performing these complex tasks.

1.3 Intelligence Analysis as Connecting the Dots

"Connecting the dots" is an appropriate metaphor in characterizing the evidential and inferential processes involved in intelligence analysis. This metaphor has gained its current popularity following the terrorist attacks in New York City and Washington, D.C., on September 11, 2001. It was frequently said that the intelligence services did not connect the dots appropriately in order to have possibly prevented the catastrophes that occurred. Since then, we have seen and heard this metaphor applied in the news media to inferences in a very wide array of contexts, in addition to intelligence analysis, including legal, military, and business contexts. For example, we have seen it applied to allegedly faulty medical diagnoses; to allegedly faulty conclusions in

historical studies; to allegedly faulty or unpopular governmental decisions; and in discussions involving the conclusions reached by competing politicians. What is also true is that the commentators or “talking heads” on television, radio, the Internet, or the sources of written accounts of inferential failures, never tell us what they mean by the phrase “connecting the dots.” A natural explanation is that they have never even considered what this phrase means and what it might involve. But we have made a detailed study of what connecting the dots entails and have found this metaphor very useful, and quite intuitive, in illustrating the extraordinary complexity of the evidential and inferential reasoning required in the contexts we have mentioned. Listening or seeing some media accounts of this process may lead one to believe that it resembles the simple tasks we performed as children when, if we connected some collection of <u>numbered</u> dots correctly, a figure of Santa Claus, or some other familiar figure, would emerge. Our belief is that critics employing this metaphor in criticizing intelligence analysts have very little awareness of how astonishingly difficult the process of connecting the dots can be in so many contexts, especially in intelligence analysis.

A natural place to begin our examination is by trying to define what is meant by the metaphor *"connecting the dots"* when it is applied to evidence-based reasoning tasks performed by intelligence analysts and others:

“Connecting the dots” refers to the task of marshaling thoughts and evidence in the generation or discovery of productive hypotheses and new evidence, and in the construction of defensible and persuasive arguments on hypotheses we believe to be most favored by the evidence we have gathered and evaluated.

Many noted persons were concerned about connecting dots in imaginative and productive ways:

- *Charles Sanders Peirce* (1839-1914) said: "Abduction is an act of insight … it is true that different elements of the hypothesis were in our minds before; but it is the idea of putting together what we had never before dreamed of putting together which flashes the new suggestion before our contemplation."
- *Henri Poincaré* (1854-1912), noting how many different rule combinations there are in mathematics and sciences, said: "The true work of the inventor consists in choosing among these combinations so as to eliminate the useless ones or rather to avoid the trouble of making them".
- *John Henry Wigmore* (1863-1943) said that we need a system "to provide the logical (or psychological) process of consciously juxtaposing the detailed related ideas, for the purpose of producing rationally a single final idea."
- *Jacques Hadamard* (1865-1963): "Indeed, it is obvious that invention or discovery, be it in mathematics or anywhere else, takes place by combining ideas."
- *Albert Einstein* (1879-1955) noted that "combinatorial play" seems to be an essential

feature of productive thought.

- *Arthur Koestler* (1905-1983) describing his works on the "bisociation" of ideas, said: "the bisociative act connects previously unconnected matrices of experience."

The following represents an account of nine complexities in the process of connecting the dots.

1.3.1 How Many Kinds of Dots Are There?

It is so easy to assume that the only kind of dot to be connected concerns details in the observable information or data we collect that may eventually be considered as evidence in some analysis. We might refer to these dots as being evidential dots. Sherlock Holmes had another term for the details in observations he made, calling them trifles. As he told Dr. Watson, “You know my method, it is founded on the observance of trifles” (Baring-Gould W.S., 1967, Vol. II, page 148). A related problem here is that most items of intelligence information may contain many details, dots, or trifles, some of which are interesting and others not. What this means is that incoming intelligence information must be carefully parsed in order to observe its significant evidential dots. In Section3.1 we give special attention to the problem of what qualifies as an evidential dot. *Not all data or items of information we have will become evidence in an analysis task*.

Consider the bombing during the Boston marathon that took place on April 15, 2013. Many images have been taken during this event. One is a widely televised videotape of two young men, one walking closely behind the other, both carrying black backpacks. This is the evidential dot shown in the bottom left of Figure 1. Why should we be interested in this evidence dot? Because it suggests to us ideas or hypotheses of what might have actually happened. Consider our ideas or thoughts concerning the relevance of the backpack dot just described. We have other evidence that the two bombs that were set off were small enough to be carried in backpacks.

This allows the inference that the backpacks carried by the two young men might have contained explosive devices and that they should be considered as suspects in the bombing. A further inference is that these two men were the ones who actually detonated the two bombs.Thus, the second type of dot concerns ideas we have about how some evidential dot, or a collection of evidential dots, is connected to matters we are trying to prove or disprove. We commonly refer to the matters to be proved or disproved as hypotheses. Hypotheses commonly refer to possible alternative conclusions we could entertain about matters of interest in an analysis. These other dots, that we call idea dots, come in the form of links in chains of reasoning or arguments we construct to link evidential dots to hypotheses. Of course, hypotheses are also ideas. Each of these idea dots refer to sources of uncertainty or doubt we believe to be interposed between our evidence and our hypotheses. This is precisely where imaginative reasoning is involved. The essential task for the analyst is to imagine what evidential dots mean as far as hypotheses or possible conclusions are concerned. Careful critical reasoning is then required to check on the logical coherence of sequences of idea dots in our arguments or chains of reasoning. In other words, does the meaning we have attached to sequences of idea dots make logical sense?

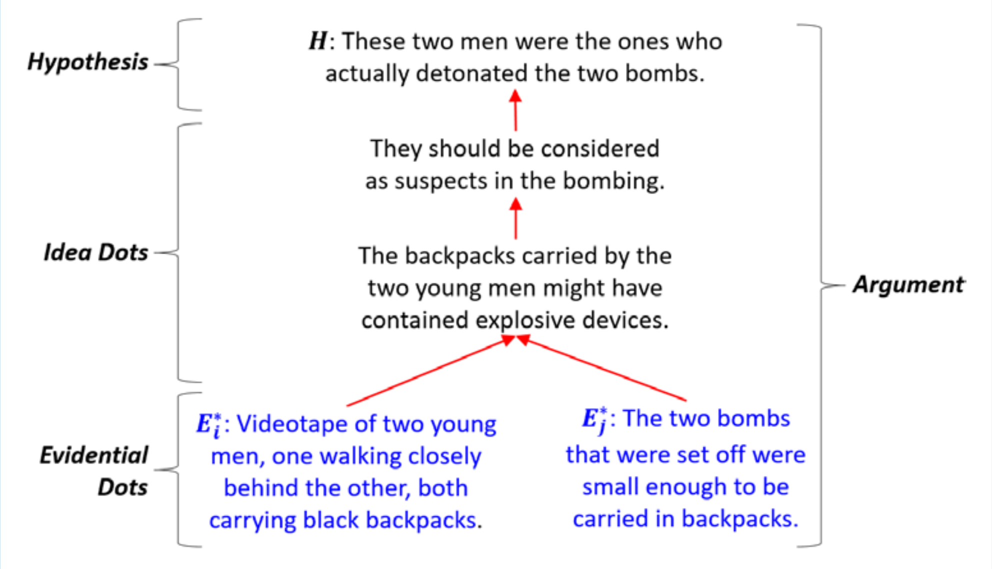
Hypothesis
Idea Dots
Evidential Dots
Argument
$\boldsymbol{H}$: These two men were the ones who actually detonated the two bombs.
They should be considered as suspects in the bombing.
The backpacks carried by the two young men might have contained explosive devices.
E_i^*: Videotape of two young men, one walking closely behind the other, both carrying black backpacks.
E_j^*: The two bombs that were set off were small enough to be carried in backpacks.

Figure 1. Types of dots to be connected: evidence, ideas, and hypotheses.

1.3.2 Which Evidential Dots Can Be Believed?

The next problem we discuss is one of the most important, challenging, and interesting problems raised in any area of intelligence analysis. From some source, a sensor of some sort, or from a person, we obtain an evidential dot saying that a certain event has occurred. But just because this source says that this event occurred does not entail that it did occur. *What is vitally necessary is to distinguish between evidence of an event and the event itself.* We adopt the following notational device to make this distinction:

- E_i^* represents the reported occurrence of event E from source I.
- E represents the actual occurrence of this event.

So, a basic inference we encounter is whether or not *E* did occur based on our evidence E_i^*. Clearly, this inference rests upon what we know about the credibility of source I. There are some real challenges here in discussing the credibility of source I. Chapter 4 discusses in detail the task of assessing the credibility of our sources of intelligence evidence.

Example 1. Consider again the evidential dot concerning the two men carrying backpacks. This is an example of tangible evidence. We can all examine this videotape to our heart's content to see what events it might reveal. The most important attribute of tangible evidence is its authenticity: *Is this evidential dot what it is claimed to be?* The FBI claims that this videotape was recorded on April 15, 2013 on Boylston St. in Boston, MA, where the bombings occurred, and recorded before the bombings occurred. Our imaginations are excited by this claim and lead to questions such as those that might have arisen in the minds of defense attorneys when this case came to trial. Was this videotape actually recorded on April 15, 2013? Maybe it was recorded on a different date. If it was recorded on April 15, 2013, was it recorded before the bombings occurred? Perhaps it was recorded after the bombings occurred. And, was this videotape actually recorded on Boylston St. in Boston, MA? It may have been recorded on a different street in Boston, MA, or perhaps on a street in a different city.

There is no better way of illustrating the importance of evidence crediility assessments than to show how such assessments form the very foundation for all arguments we make from evidence to possible conclusions. In many situations people will mistakenly base inferences on the assumption that an event E has occurred just because we have evidence E_i^* from source I. This amounts to the suppression of any uncertainty we have about the credibility of source I (whatever this source might be).

Figure 2 is a simple example illustrating this credibility foundation; it will also allow us to introduce the next problem in connecting the dots. What this figure shows is an argument from evidence E_i^* to hypothesis H. As shown, the very first stage in this argument concerns an inference about whether the event E actually occurred. This is precisely where we consider whatever evidence we may have about the credibility of source I. We may have considerable uncertainty about whether the event E occurred. All subsequent links in this argument concern the relevance of event E on hypothesis H. These relevance links connect the idea dots we discussed. As Figure 2 shows, each idea dot is a source of uncertainty associated with the logical connection between whether event E did occur and whether H is true.

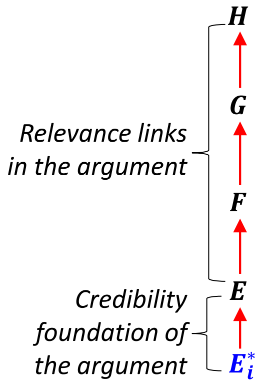
H
G
Relevance links in the argument
F
Credibility foundation of the argument
E
E_i^*

Figure 2. Credibility foundation of argument.

1.3.3 Which Evidential Dots Should Be Considered?

In all of the contexts we have considered, there is usually no shortage of potential evidential dots. In fact, in many of these contexts, persons drawing conclusions about matters of importance are swamped with information or data. This situation is currently being called the "big data problem." Here we begin to consider vital matters concerning the discovery or investigative tasks, and the imaginative or creative reasoning these tasks involve. Unfortunately, in many situations people or organizations try to collect everything in the hope of finding something useful in an inference task. This wasteful practice is one reason why the big data problem exists, since only a minute fraction of the information collected will be relevant in any inference of concern. In our work we have paid great attention to the process of discovery that necessarily takes place in a world that keeps changing all the while we are trying to understand parts of it of interest to us in our inference tasks. As will be discussed in Section 1.5 this is an ongoing seamless activity in which we have evidence in search of hypotheses, hypotheses in search of evidence, and the testing of hypotheses all going on at the same time. Hypotheses you entertain, questions you ask, particular evidence items, and your accumulated experience, all allow you to examine which evidential dots to consider. Part of our objectives here is to make the process of discovery more efficient. As we will also discuss, these discovery tasks involve mixtures of three different forms of reasoning: abduction (imaginative, creative, or insightful reasoning that shows that something is possibly true), deduction, (that shows that something is necessarily true) and induction (that shows that something is probably true). These forms of reasoning provide the bases for our idea dots.

1.3.4 Which Evidential Dots Should We Try to Connect?

Here comes a matter of great complexity. It usually happens that hypotheses we entertain are generated from observations we have made involving potential evidential dots. On limited occasions, we can generate a hypothesis from a single evidential dot. For example, in a criminal investigation finding a fingerprint will suggest a possible suspect in the case. But in most cases, it takes consideration of combinations of evidential dots in order to generate plausible and useful hypotheses, as illustrated in the following example based on accounts given in Time magazine and the Washington Post.

Example 2. From European sources came word that terrorists of Middle Eastern origin would make new attempts to destroy the World Trade Center, this time using airliners. Many threats are received every day, most of which come to nothing. However, from several civilian flying schools in the USA came word (to the FBI) that persons from the Middle East were taking flying lessons, paying for them in cash, and wanting only to learn how to steer and navigate heavy aircraft but not how to make takeoffs and landings in these aircraft. By itself, each of these two items of information, though admittedly strange, may not have seemed very important. But, taken together, and recalling that, during World War II kamikaze Japanese pilots used their airplanes as bombs, they might have caused even an Inspector Lestrade to generate the hypothesis that there would be attacks on the World Trade Center using hijacked airliners. The hijackers would not need to learn how to make takeoffs; the aircrafts' regular pilots would do this. There would be no need for the hijackers to know how to land aircraft, since no landings were intended, only crashes into the World Trade Center and the Pentagon. Why were these two crucial items of information not considered together? The answer seems to be that they were not *shared* among relevant agencies. Information not shared cannot be considered jointly, with the result that their joint inferential impact could never have been assessed. For all time, this may become the best (worst) example of failure to consider evidence items together. This is just one reason why we will so strongly emphasize the importance of evidence marshaling strategies in this book. Even Sherlock Holmes would perhaps not have inferred what happened on September 11, 2001 if he had not been given these two items of information together.

The problem, however, is that here we encounter a combinatorial explosion, since the number of possible combinations of two or more evidential dots is exponentially related to the number of evidential dots we are considering. Suppose we have some number N of evidential dots. We ask the question: How many combinations C of two or more evidential dots are there? The answer is given by the following expression: $C = 2^{\mathbf{N}} – (N + 1)$. This expression by itself does not reveal how quickly this combinatorial explosion takes place. Here are a few examples showing how quickly C mounts up with increases in N:

- for N = 10 C = 1013
- for N = 25 C = 33,554,406
- for N = 50 $C = 1.13 \times 10^{15}$
- for N = 100 $C = 1.27 \times 10^{30}$

There are several important messages in this combinatorial analysis for intelligence analysis. The first concerns the size of N, the number of potential evidential dots that might be connected. Given the array of sensing devices and human observers available to our intelligence services, the number N of potential evidential dots is as large as you wish to make it. In most analyses N would certainly be greater than 100 and would increase as time passes. Remember that we live in a non-stationary world in which things change and we find out about new things all the time. So, in most cases, even if we had access to the world's fastest computer, *we could not examine all possible evidential dot combinations*.

Second, trying to examine all possible evidential dot combinations would be the act of looking through everything with the hope of finding something. This would be a silly thing to do, even if it were possible. The reason of course is that most of the dot combinations would tell us nothing at all. What we are looking for are combinations of evidential dots that interact or are dependent in ways that suggest new hypotheses or possible conclusions. If we examined these dots separately or independently, we would not perceive these new possibilities. The bottom left-hand side of Figure 3 is an abstract example; a tragic real life illustration is in the right-hand side of the figure.

In the top left-hand side of Figure 3 we show an instance where three evidential dots have been examined separately or independently in which case they tell us nothing interesting. Then someone notices that, taken together, these three dots combine to suggest a new hypothesis H that no one has thought about before, as shown in bottom left-hand side of the figure. What we have here is a case of evidential synergism in which two or more evidence items mean something quite different when they are examined jointly than they would mean if examined separately or

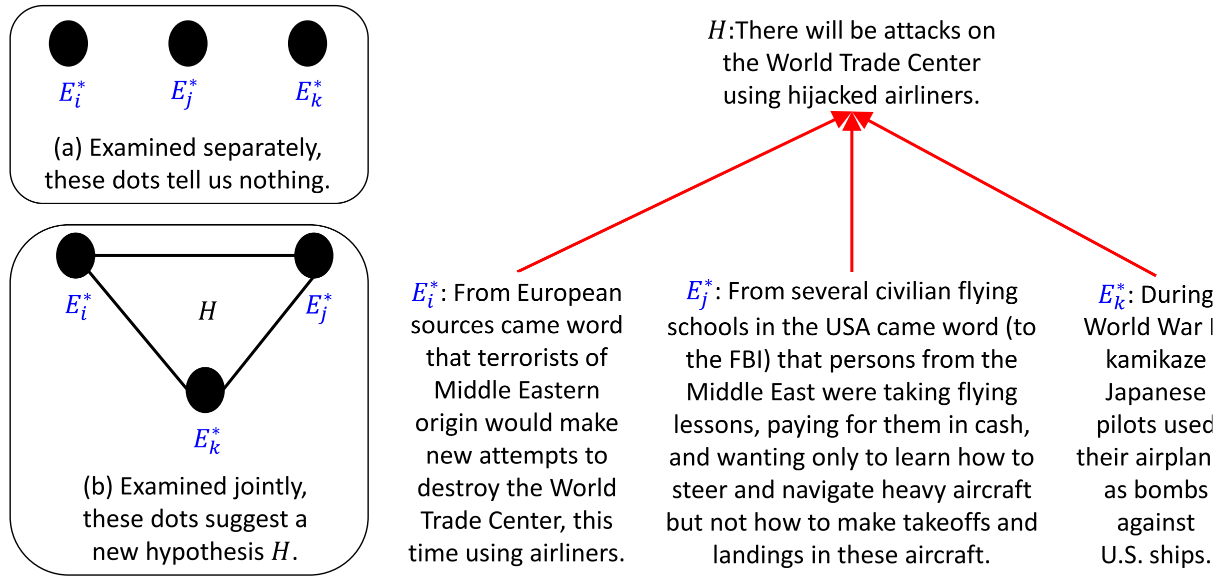
E_i^*
E_j^*
E_k^*
(a) Examined separately, these dots tell us nothing.
E_i^*
H
E_j^*
E_k^*
(b) Examined jointly, these dots suggest a new hypothesis H.
H:There will be attacks on the World Trade Center using hijacked airliners.
E_i^*: From European sources came word that terrorists of Middle Eastern origin would make new attempts to destroy the World Trade Center, this time using airliners.
E_j^*: From several civilian flying schools in the USA came word (to the FBI) that persons from the Middle East were taking flying lessons, paying for them in cash, and wanting only to learn how to steer and navigate heavy aircraft but not how to make takeoffs and landings in these aircraft.
E_k^*: During World War I kamikaze Japanese pilots used their airplan as bombs against U.S. ships.

Figure 3. Illustration of evidential synergism.

independently. Here we come to one of the most interesting and crucial evidence subtleties or complexities that have, quite frankly, led to intelligence failures in the past: failure to identify and exploit evidential synergisms. We will address this matter in other problems we mention concerning connecting the dots.

What is absolutely crucial in selecting dot combinations to examine is an analyst's experience and imaginative reasoning capabilities. What we should like to have is a conceptual "magnet" that we could direct at a base of evidential dots that would "attract" interesting and important dot combinations, as discussed in Section 2.4.

1.3.5 How to Connect Evidential Dots to Hypotheses?

As we will discuss in Section 3.2, all evidence has three major credentials or properties: relevance, credibility, and inferential force or weight. No evidence ever comes to us with these three credentials already attached; they must be established by defensible and persuasive arguments linking the evidence to the hypotheses we are considering. As we will see, relevance answers the question: *"So what? How is this datum or information item linked to something we are trying to prove or disprove?"* If such relevance linkage cannot be established, this datum is irrelevant or useless. As discussed above, credibility answers the question: *"Can we believe what this evidence is telling us?"* The inferential force or weight credential asks: *"How strong is this evidence in favoring or disfavoring the hypothesis?"* This is where probability enters our picture since, for very good reasons, the force or weight of evidence is always graded in probabilistic terms.

A relevance argument is precisely where the idea dots become so important. Considering an item of information, an analyst must imagine how this item could be linked to some hypothesis being considered before it could become an item of evidence. These idea dots forming this linkage come in the form of propositions or statements indicating possible sources of doubt or uncertainty in the imagined linkage between the item of information and hypotheses being considered. For a simple example, look again at Figure 2 where we show a connection between evidence E_i^* and hypothesis H. An analyst has an item of information from source I concerning the occurrence of event E that sounds very interesting. This analyst attempts to show how event E, if it did occur, would be relevant in an inference about whether hypothesis H is true. So the analyst forms the following chain of reasoning involving idea dots. The analyst says, "If event E were true, this would allow us to infer that event F might be true, and if F were true, this would allow us to infer that event G might be true. Finally, if event G were true, this would make hypothesis H more probable." If this chain of reasoning is defensible, the analyst has established the relevance of evidence E_i^* on hypothesis H.

In forming this argument the analyst wisely begins with the credibility foundation for this whole argument: *Did event E really occur just because source I says it did?*

There are several important things to note about relevance arguments; the first concerns their defense. Suppose the argument in Figure 2 was constructed by analyst $\mathcal{A}$. $\mathcal{A}$ shows this argument to analyst $\mathcal{B}$ who can have an assortment of quibbles about this argument. Suppose $\mathcal{B}$ says, "You cannot infer F directly from E; you need another step here involving event K. From E you can infer that K occurred, and then if K occurred, then you can infer F." Now comes analyst $\mathcal{C}$ who also listens to $\mathcal{A}$'s argument. $\mathcal{C}$ says, "I think your whole argument is wrong. I see a different reasoning route from E to hypothesis H. From E we can infer event R, and from R, we can infer event S, and from S we can infer T, which will show that hypothesis H is less probable." Whether or not there is any final agreement about the relevance of evidence E_i^*, analyst $\mathcal{A}$ has performed a real service by making the argument openly and available for discourse and criticism by colleagues. There are several important messages here.

First, *there is no such thing as a uniquely correct argument from evidence to hypotheses*. What we all try to avoid are disconnects or non sequiturs in the arguments we construct. But even when we have an argument that has no disconnects, someone may be able to come up with a better argument. Second, we have only considered the simplest possible situation in which we used just a single item of potential evidence. But intelligence analyses are based on masses of evidence of many different kinds and that come from an array of different sources. In this case we are obliged to consider multiple lines of argument that can be connected in different ways. It is customary to call these complex arguments inference networks.

But analysts gain much assistance in developing such analyses by learning about argument construction methods devised nearly a hundred years ago by a world-class evidence scholar named John H. Wigmore (1863-1943). Wigmore (1913; 1937) was the very first person to carefully study what today we call inference networks. In Chapter 6 we present in detail the modern version of a Wigmorean inference netwok, called augmented Wigmorean argumentation.

There is also a message here for critics such as news writers and the taking heads on television. These critics always have an advantage never available to practicing intelligence analysts. Namely, they know how things turned out or what actually happened in some previously investigated matter affecting the nation's security. In the absence of clairvoyance, analysts studying a problem will never know for sure, or be able to predict with absolute certainty, what will happen in the future. A natural question to ask these critics is: *What arguments would you have constructed if all you knew was what the analysts had when they made their assessments?* This would be a very difficult question for them to answer fairly, even if they were given access to the classified evidence the analysts may have known at the time.

1.3.6 What Do Our Dot Connections Mean?

The previous item concerns efforts designed to establish the defensibility of complex arguments. But what do these arguments mean to persons for whom these arguments are being constructed? This question raises matters concerning how persuasive our arguments are when they are taken all together. Our view is that the persuasiveness of an argument structure depends, in large part, upon the nature of the probabilities we assess and combine in our arguments and in stating our major conclusions.

Here we consider the direction and force (or weight) of our arguments based on the combined evidence we have considered. Direction refers to the hypothesis we believe our evidence favors most. Force or weight means how strongly we believe the evidence favors this hypothesis over alternative hypotheses we have considered. There are two uncontroversial statements we can make about the force or weight of evidence. The first is that the force or weight of evidence has vector-like properties. What this means is that evidence points us in the direction of certain hypotheses or possible conclusions with varying degrees of strength. The second is that the force or weight of evidence is always graded in probabilistic terms indicating our uncertainties or doubts about what the evidence means in terms of its inferential direction and force. But beyond these two statements, controversies begin to arise.

Before we consider assorted controversies, it is advisable to consider where our uncertainties or doubts come from in the conclusions we reach from evidence. Have a look once again at Figure 2 involving a simple example based on a single item of evidence. Our evidence here was E_i^*, from source I, saying that event E occurred. We ask the question: *How strongly does this evidence E_i^* favor hypothesis H*? As we discussed, this argument was indicated by what we termed idea dots, each one indicating what the analyst constructing this argument believed to be sources of doubt or uncertainty associated with the argument from the evidence to the hypothesis. As you see, there are two major origins of uncertainty: those associated with the credibility of source I, and those associated with links in the analyst's relevance argument. So, the force of evidence E_i^* on hypotheses H depends on how much uncertainty exists in this entire argument involving each one of its credibility and relevance links. The interesting message here is that the evidence inferential force or weight credential depends on its other two credentials, credibility and relevance.

In the simple example just discussed there are four major origins of uncertainty, one associated with credibility and three associated with relevance. But this is the easiest possible situation since it involves only one item of evidence. Think of how many sources of uncertainty there might be when we have a mass of evidence together with multiple complex and possibly interrelated arguments. The mind boggles at the enormity of the task of assessing the force or weight of a mass of evidence commonly encountered in intelligence analysis when we have some untold numbers of sources of credibility and relevance uncertainties to assess and combine. We are certain that critics of intelligence analysts have never considered how many evidential and idea dots there would be to connect.

So, the question remains: *How do we assess and combine the assorted uncertainties in complex arguments in intelligence analysis, and in any other context in which we have the task of trying to make sense out of masses of evidence?* Here is where controversies arise. The problem is that there are several quite different views among probabilists about what the force or weight of evidence means and how it should be assessed and combined across evidence in either simple or complex arguments. Each of these views has something interesting to say, but no one view says it all. As you will see in Chapter 5, we consider four systems of probability. We do consider the conventional or Bayesian system that involves numerical probability judgments, but there are some severe limitations to this approach. Therefore, we also consider the Belief Functions, Baconian, and Fuzzy probability systems. But we devote considerable attention to a combination of the Baconian and Fuzzy systems that require probabilities to be expressed in words rather than in numbers. The Baconian system, resting upon the view of Sir Francis Bacon, is especially relevant in the contexts we have mentioned. It is the only system of probability that concerns the completeness, as well as the strength, of the evidential coverage we can claim in the conclusions we reach from our evidential dots.

Later in this book we will discuss how the Cogent system allows you to assess and combine probabilistic judgments in situations in which many such judgments are required. There is further difficulty as far as judgments of the force or weight of evidence is concerned. Analysts, or teams of analysts, may agree about the construction of an argument but disagree, often vigorously, about the extent and direction of the force or weight this argument reveals. There may be strong disagreements about the credibility of sources of evidence or about the strength of relevance linkages. These disagreements can only be resolved when arguments are made carefully, and are openly revealed so that they can be tested by colleagues. A major mission of the Cogent system is to allow you to construct arguments carefully and critically, and encourage you to share them with colleagues so that they can be critically examined.

There is one final matter of interest in making sense out of masses of evidence and complex arguments. Careful and detailed argument construction might seem a very laborious task, no matter how necessary it is. Now consider the task of revealing the conclusions resulting from an analysis to some policy-making "customer" who has decisions to make that rest in no small part on the results of an intelligence analysis. What this "customer" will probably not wish to see is a detailed inference network analysis that displays all of the dots that have been connected and the uncertainties that have been assessed and combined in the process. A fair guess is that this "customer" will wish to have a narrative account or a story about what the analysis predicts or explains. In some cases, "customers" will require only short and not extensive narratives. This person may say: *Just tell me the conclusions you have reached and briefly explain why you have reached them.* So the question may be asked: *Why go to all the trouble to construct defensible and persuasive arguments when our 'customers' may not wish to see their details?*

There is a very good answer to the question just raised. *Your narrative account of an analysis must be appropriately anchored on the evidence you have*. What you wish to be able to tell is a story that you believe contains some truth; i.e., it is not just a good story. The virtue of careful and critical argument construction is that it will allow you to anchor your narrative not only on your imagination, but also on the care you have taken to subject your analysis to critical examination. There is no telling what questions you might be asked about your analysis. Rigor in constructing your arguments from your evidence is the best protection you have in dealing with "customers" and other critics who might have entirely different views regarding the conclusions you have reached. The Cogent system is designed to allow you and others to critically evaluate the arguments you have constructed.

1.3.7 Whose Evidential Dots Should Be Connected?

There are several very easy answers to this question. One obvious answer is that all the potential evidential dots collected by any intelligence service that bear upon a problem involving our nation's security should be shared or brought together. Since September 11, 2001, so many examples of potentially relevant evidence, gathered by different intelligence services, were never shared across agencies and offices. The basic problem this creates is that the extremely important evidential synergisms we discussed above can never be detected and exploited in reaching analytic conclusions. In some cases, this has resulted in our failure to reach any conclusion at all in some important matter. This forms the basis for one of the major criticisms of the intelligence services in their failure to "connect the dots." In some instances in the past, potential evidence may have been viewed as a "proprietary" commodity to be shared only at the discretion of the agency or person that collected it. In other cases, there have been various statutory rules preventing sharing of evidence across intelligence-related services. Whatever the causes for this lack of sharing of intelligence information, this problem has been of great concern in the past few years.

But there is one way that Cogent can assist in the detection and inferential exploitation of possible evidential synergisms and it is something that rests on analysts, and analyst teams, at work on an intelligence problem. Careful argument construction will help reveal the incompleteness of available evidence. The analysts might easily observe that not all questions that should be asked about the problem at hand have in fact been answered. So, this forms the basis for asking questions such as:

- Have any other agencies or offices attempted to answer these questions that we believe have gone unanswered?
- If these other agencies have gathered such evidence, how can we best justify or be able to have ready access to it?
- What collection efforts should be mounted to gather evidence necessary in order to provide more complete assessments of evidence necessary to form more productive conclusions?

In many cases such evidence may have never been collected. In these cases, analysts can play very important roles in directing effective and productive evidence collection efforts. In so many instances it seems that we try to collect everything with the hope of finding something. This is one reason why we often correctly believe that we are drowning in information. More imaginative efforts are required in order to collect potential evidential dots of actual relevance in inference problems faced by intelligence analysts. This is another area in which the imagination of analysts becomes so important.

1.3.8 How Persistent Are the Dot Connections?

Among the many complexities of intelligence analysis none seem more difficult than the fact that analysts must provide their explanations or predictions in a non-stationary world. In short, the world keeps changing as analysts are trying their best to understand it well enough to provide explanations or to make predictions. One consequence is that we have continuing streams of new information, some items of which we will assess as being relevant evidence regarding our explanations or predictions. An explanation for some pattern of past events analysts have previously regarded as correct may now seem incorrect in light of new evidence just discovered today. A prediction regarded as highly likely today may be overtaken by events we will learn about tomorrow. In fact, the very questions we have asked yesterday may need to be revised or may even seem unimportant in light of what we learn today. One consequence of all of this is that the process of discovery or investigation in intelligence analysis is a ceaseless activity. It would be a drastic mistake to view discovery in intelligence analysis as being a stationary activity in a non-stationary world.

1.3.9 How Much Time Do We Have to Connect the Dots?

A major objective of intelligence analysis is to help ensure that the policies and decisions reached by the governmental and military leaders, at all levels, are well informed. Analysts face different requirements in their efforts to serve these policy and decision-making "customers". In some cases, teams of analysts participate in lengthy analyses that combine evidence from every available source to make long-term assessments on matters of current and abiding interest. But in many other cases current analyses are required to answer questions that are of immediate interest and that do not allow analysts time for extensive research and deliberation on available evidence regarding the questions being asked. How Cogent can help alleviate this difficulty is discussed in Chapter 6.

Identifying the complexities of intelligence analysis is actually the easy part. What is not so easy are efforts to assist analysts in coping with the complexities of the evidential reasoning tasks they routinely face.

1.4 What Ingredients of Analysis are to be Generated by Imaginative Thought?

It would be a very rare occurrence if you encountered an analytic task in which all possible hypotheses, all available evidence, and all arguments connecting the evidence and the hypotheses were supplied for you. All these ingredients you will have to generate or discover for yourself. This is where your imaginative reasoning becomes necessary. Now it happens that imaginative reasoning, though necessary, is not sufficient. Suppose you have generated some alternative hypotheses from the evidence you have discovered, or selected from some larger collection of evidence, that seems relevant to these hypotheses. As we will discuss in this book, you must also establish the relevance, credibility, and inferential force “credentials” of the evidence you have. This involves critical as well as imaginative reasoning on your part. You must be able to construct arguments from evidence to hypotheses that are both defensible and persuasive; this is where critical reasoning also becomes vitally necessary. You may have generated entirely plausible hypotheses as well as evidence that you believe bears on these hypotheses, but, if your arguments linking your evidence and your hypotheses have non sequiturs, disconnects or “short circuits” that are recognized by others, your analysis will fail to be defensible or persuasive.

We understand that intelligence analysis is a very complex activity often involving many persons in many locations. It may certainly be the case that potential evidence in your current analytic task is actually generated by other persons. For example, you may have a steady stream of message traffic or regular reports of some kind that arrive at your desk every day. Though you did not yourself generate or discover these items of information yourself, you must decide which items from the mass of items you receive could indeed be evidence relevant in an inference task you presently have. But it is also true that your imaginative reasoning is involved when you request information, and potential evidence, that no one has at present.

1.4.1 Generating Main Hypotheses to be Defended by Evidence and Argument

The main hypotheses may be generated as possible answers of the intelligence questions asked. There are different types of such questions.

Types of Questions

Binary (yes/no) questions, such as:

- Does Hakka have chemical weapons?
- Did John steal the car?

Questions having a finite set of answers. The number of these questions is determined by a finite set of possibilities based on prior knowledge:

- Which missile was tested?
- Which terrorist group performed the terror attack?
- Who passed the documents?

Note that, because of information gaps, it is possible that all of the potential answers identified in the problem are inconsistent with the available information. For example, if two terrorist groups appear as the prime suspects in a terror attack, but there is information that tends to rule out or minimize the likelihood of their involvement, the correct answer might be an unidentified third terrorist group.

Questions having a potentially unlimited set of answers. The number of these questions is ultimately constrained by the available information. Examples of such questions are:

- Why did the foreign minister resign?
- Why did President Marcos reject the treaty?
- Why is country Redland developing nuclear weapons?

In other cases, the hypotheses will arise from observations we make. In this case we have evidence in search of hypotheses, or possible explanations for what we have observed. In some cases, when our evidence is scant, it may even be appropriate to refer to an initial hypothesis as a guess.

Types of Hypotheses

Generally, our major hypotheses refer to events or situations that we are presently unable to observe directly. These events may have happened in the past, are now possibly happening, or may possibly happen in the future. Here are three examples of hypotheses concerning past, present, or future events.

Hypotheses concerning a past event
A terrorist incident occurred two months ago in which several lives were lost. After an investigation, two suspects X and Y have been identified. Here are some hypotheses we could entertain about this past event:

- H_1: Person X was the one involved in this incident.
- H_2: Person Y was the one involved in this incident.
- H_3: Both X and Y were involved in this incident.
- H_4: Neither X nor Y were involved in this incident.

Hypotheses concerning an event that may be happening "now"
You might have reason to suspect that Country X is still holding prisoners of war taken years ago during a conflict we had with it. Your suspicion here forms one hypothesis, H_5: Country X is now holding some of our POWs. This example illustrates why it is true that we always have more than one hypothesis. Another possibility is $notH_5$, that is, H_6: Country X is <u>not</u> holding any of our POWs.

Hypotheses concerning a future event or situation
We have been closely monitoring the deteriorating relations between Countries A and B that share a common border. We now entertain the possibility that there will be armed conflict between these two countries "in the near future." Thus, we have as major hypothesis H_7: There will be armed conflict between A and B in the near future. Another hypothesis, of course, is $notH_7$, that is, H_8: There will be no armed conflict between A and B in the near future.

This example allows us to see that we will often need to make our hypotheses more specific. The hypothesis that there will be armed conflict between A and B is actually not very informative if it is our final stated conclusion. Decision makers will wish to know such things as: *Who will start it? How will the conflict proceed? How long will it last? and Who will win?*

All of these examples concern events/situations that <u>might have happened</u>, are <u>now possibly happening</u>, or <u>might happen in the future</u>. We have no certainty about any of these hypotheses. At the moment, they are all simply possibilities. If, at the moment, we reported any of these hypotheses in the form of a conclusion, we would not be taken seriously. We have given no one else any reasons why the hypothesis we have chosen to report as a conclusion should be favored over any of the other hypotheses that are possible. This is where our next two ingredients come in, evidence and arguments.

1.4.2 Generating the Evidential Grounds for Arguments

The second major ingredient of intelligence analyses is evidence that can be defended as relevant in showing why some hypothesis is true. Here is an example of its importance.

Take any of the three situations just mentioned in Section 1.4.1 concerning hypotheses about either past, current, or future events:

H_1: Person X was the one involved in this incident.
H_5: Country X is now holding some of our POWs.
H_7: There will be armed conflict between A and B in the near future.

All of these situations involve events that are not currently directly observable to us. We were not at the scene of the terrorist incident; we have no direct observations of the presence of the POWs; and we cannot see into the minds of the leaders in Countries A and B in order to read their intentions. But, we can observe other events or things that can serve as evidence, signs, or indicators of any of these hypotheses. So, we might define evidence in the following way:

Evidence is any observable sign, indicator, or datum we believe is relevant in deciding upon the extent to which we infer any hypotheses we have entertained as being correct or incorrect.

Here is some examples of evidence we might find concerning the above hypotheses.

For H_1: We might find evidence showing that X was in the near vicinity of the incident one hour before it happened.

For H_5: A recent visitor to Country X shows us a dog-tag he says was given to him by a resident of X. On this tag is the name of a soldier who has been missing since our conflict with Country X ended.

For H_7: We might obtain evidence bearing upon the state of military preparedness of either country.

1.4.3 Generating Arguments Linking Evidence and Hypotheses

The third major ingredient of intelligence analysis concerns the arguments we must construct in defense of the relevance, credibility, and force or weight of our evidence. Again, no item of evidence comes to us with these credentials already established; they must be established by arguments. The arguments we make form logical links between the evidence we have and the hypotheses we entertain. One way to look at an argument is to say that it forms a chain of reasoning from evidence to hypotheses. Often, there will be many links in a chain of reasoning.

Figure 4 shows an argument from the evidence E^* (X was in the near vicinity of the incident one hour before it happened) to the hypothesis H (Person X was the one involved in this incident). Our argument might run as follows:

- We have evidence that X was in the near vicinity of the incident one hour before it occurred. Therefore, it is possible that X was indeed in the near vicinity of the incident one hour before it occurred.
- Then he might have been at the scene of the incident when it occurred.
- Then, if he was at the scene of the incident at the time it occurred, he might have been a participant in the incident.

The argument just constructed is one made in defense of the relevance of evidence that X was in the near vicinity of the incident an hour before it occurred. Notice that, if you regard this argument as plausible, we have a direct link between the evidence and our hypothesis.

However, all the argument in Figure 4 shows is that X might have been a participant in the incident. Remember from Section 1.4.1 that we were considering three other hypotheses, in addition to this one. Therefore, what we would like to know is which of the four hypotheses is most likely. This would require the analysis of all the hypotheses, and not just based on single items of evidence. It would also require an assessment of how likely each hypothesis is, based on the relevance, the credibility, and the inferential force of evidence.

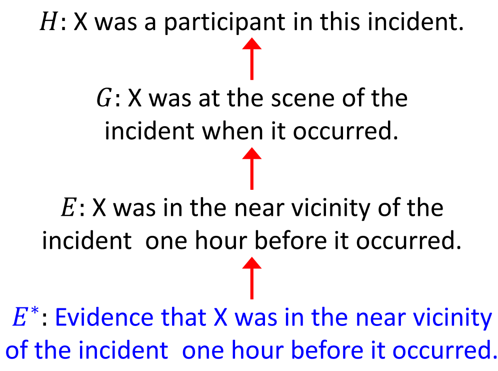
H: X was a participant in this incident.
G: X was at the scene of the incident when it occurred.
E: X was in the near vicinity of the incident one hour before it occurred.
E^*: Evidence that X was in the near vicinity of the incident one hour before it occurred.

Figure 4. Sample argument.

1.5 A Computational Approach to Intelligence Analysis

1.5.1 The Arch of Knowledge

We have found the metaphor of an *arch of knowledge* to be very useful in summarizing the many ideas expressed over the centuries concerning the generation of new thoughts and new evidence. This metaphor comes from the work of the philosopher David Oldroyd in a valuable work having this metaphor as its title (Oldroyd 1986). Figure 5 shows this metaphor applied in the context of science. Based upon existing records, it seems that Aristotle (384BC-322BC) was the first to puzzle about the generation or discovery of new ideas in science. From sensory observations we generate possible explanations, in the form of hypotheses, for these observations. It was never clear from Aristotle's work what label should be placed on the upward, or discovery-related arm of the arch in Figure 5. By some accounts, Aristotle's description of this act of generating hypotheses is called "intuitive induction" (Cohen and Nagle, 1934; Kneale, 1949). The question mark on the upward arm of the arch in Figure 5 simply indicates that there is still argument about what this discovery-related arm should be called. By most accounts, the downward arm of the arch concerns the deduction of new observable phenomena assuming the truth of a generated hypothesis (Schum, 2001b).

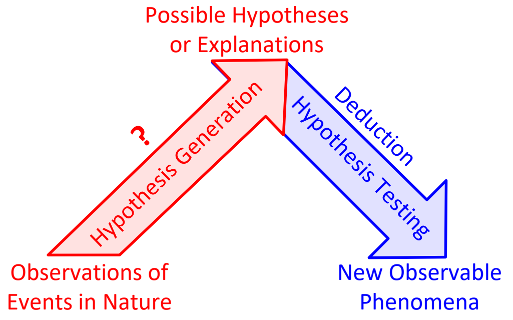

Figure 5. The "Arch of Knowledge" in Science.

Over the millennia since Aristotle, many people have tried to give an account of the process of discovering hypotheses and how this process differs from ones in which existing hypotheses are justified. Galileo Galilei (1564-1642) thought that we "reason backward" inductively to imagine causes (hypotheses) from observed events, and we reason deductively to test the hypotheses. A similar view was held by Isaac Newton (1642-1727), John Locke (1632-1704), and William Whewell (1794-1866). Charles S. Peirce (1839-1914) was the first to suggest that new ideas or hypotheses are generated through a different form of reasoning, which he called *abduction* and associated with imaginative reasoning (Peirce, 1898; 1901). His views are very similar to those of Sherlock Holmes, the famous fictional character of Conan Doyle (Schum, 1999). The next section introduces a systematic approach to intelligence analysis which is based on the scientific method and is supported by the Cogent system.

1.5.2 Intelligence Analysis in the Framework of the Scientific Method

Within the framework of the scientific method, intelligence analysis can be viewed as *ceaseless discovery of evidence, hypotheses, and arguments* in a non-stationary world, involving collaborative computational processes of *evidence in search of hypotheses*, *hypotheses in search of evidence*, and *evidentiary testing of hypotheses* (see Figure 6).

Intelligence analysis may start with an intelligence question: *What are the possible answers to this question?* The imagined answers are the hypotheses to be analyzed. The process may also start with an interesting observation that needs to be explained: *What hypotheses would explain this observation?* Answering such questions involves *abductive (imaginative) reasoning* that shows that something is

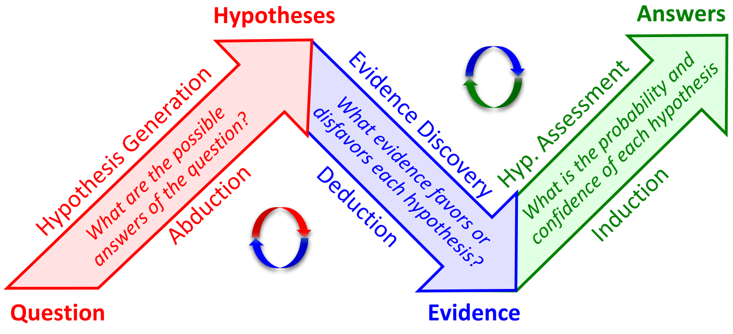
Hypotheses
Answers
Hypothesis Generation
What are the possible answers of the question?
Abduction
Evidence Discovery
What evidence favors or disfavors each hypothesis?
Deduction
Hyp. Assessment
What is the probability and confidence of each hypothesis
Induction
Question
Evidence

Figure 6. Intelligence analysis as discovery

possibly true. We call this process *evidence/question in search of hypotheses*.

To determine which of the hypotheses is true, one needs evidence. One approach to discover evidence, is to put each of the generated hypothesis to work by asking the question: *What evidence would be observable if this hypothesis were true?* One then decomposes the hypothesis into simpler and simpler hypotheses, and uses the simplest hypotheses to generate new lines of inquiry and discover new evidence. The reasoning might go as follows:

- If H were true then the sub-hypotheses H_1 and H_2 would also need to be true.
- But if H_1 were true then one would need to observe evidence E_1^*, and so on.

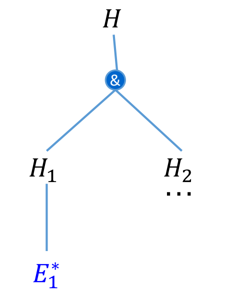

Figure 7. Hypothesis H in search of evidence.

This process leads to the discovery of new evidence by identifying the necessary conditions for hypothesis H to be true (See Figure 7).

A broader question to guide the discovery of evidence is: *What evidence would favor or disfavor this hypothesis?* In this case one would look for *sufficient conditions*, or even *indicators*, for a hypothesis to be true or false. The reasoning, illustrated in Figure 8, might go as follows:

- H would be true if H_1 and H_2 would be true.
- Then H_2 would be true if either H_{1a} or H_{1b}would be true.
- Searching for evidence relevant to H_{1a} we discover E_{1a}^*.
- Searching for evidence relevant to H_{1b} we discoverE_{1b}^* and E_{1b}^*.

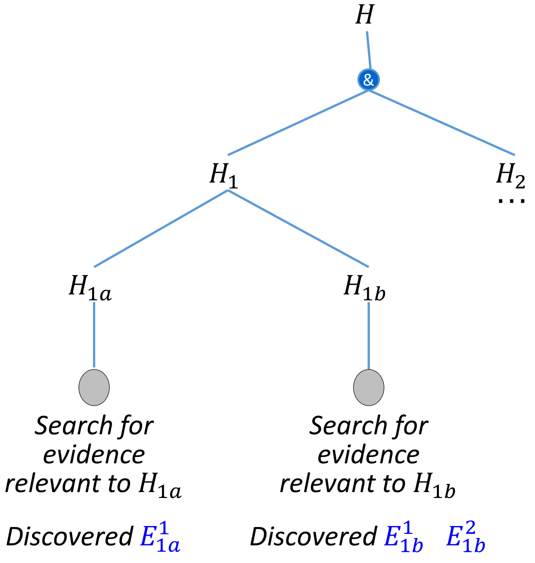
H
&
H_1
H_2
...
H_{1a}
H_{1b}
Search for evidence relevant to H_{1a}
Search for evidence relevant to H_{1b}
Discovered E_{1a}^1
Discovered E_{1b}^1 E_{1b}^2

Figure 8. Another strategy for hypothesis in search of evidence.

We call this process *hypothesis is search of evidence*. It involves *deductive reasoning* that shows that something is *necessarily* true.

Through *inductive reasoning,* that shows that something is *probably* true, one tests the hypotheses. First one assesses the probabilities of the bottom-level hypotheses based on the discovered evidence. Then the probabilities of the upper level hypotheses are computed based on the logical structure of the argumentation.

Evidence in search of hypotheses, *hypotheses in search of evidence*, and *evidentiary assessment of hypotheses* are collaborative processes that support each other in recursive calls, as shown by the circles in Figure 6. For example, the discovery of new evidence may lead to the modification of the existing hypotheses or the generation of new ones that, in turn, lead to the search and discovery of new evidence. Also, inconclusive assessments of the hypotheses lead to the need of discovering additional evidence. Since these processes are generally very complex and involve both imaginative and critical reasoning, they can be best approached through the synergistic integration of the analyst's imaginative reasoning and computer's knowledge-based critical reasoning. In the following we illustrate this systematic approach to intelligence analysis with an example of anticipatory analysis.

1.5.3 Anticipatory Intelligence

Anticipatory intelligence is the complex task of identifying and assessing new, emerging trends, changing conditions, and underappreciated developments to challenge long-standing assumptions, encourage new perspectives, identify new opportunities, and provide warning of threats to national interests, based on current information of all kinds that come from a variety of different sources (ODNI, 2019). *It is aimed at potential events, including low-probability high threat events and involves active attention management, focusing attention on likely sources of critical information* (Klein et al., 2007).

Analyst's Standpoint

Mavis, an intelligence analyst, follows incoming information that is relevant to her "account". When interesting information arrives, she attempts to generate possible explanations for this information and may also generate predictions based on her explanations for the events in this information. We will show how evidence about a missing cesium-137 canister from a company warehouse leads her to anticipating that a dirty bomb will be set off in the Baltimore-Washington D.C. area. But first we will present Mavis's standpoint that helps understand her reasoning (Anderson et al., 2005, pp. 115-117). There are five basic questions an analyst answers in a declaration of standpoint. We should note that the answers she provides will be relevant to the particular situation involved in this example. The reason is that her standpoint might change in different situations she encounters.

Question I: *Who am I?* [i.e., What role am I playing or what "hat" am I wearing in this present analysis?]. I am a career intelligence analyst having seven years of experience. My current account concerns counterterrorism. At present my account involves identifying possible terrorist activities that might occur here in the USA. These actions might be planned and implemented by persons holding US citizenship or by foreign nationals who are here illegally. I add here that there are several other analysts who share my account. In short, I am not the only person involved in the detection of possible terrorist actions.

Question 2*: At what stage of what process am I?* Counterterrorism activities, such as those involving my account, are ongoing and have been so for several years. I have been managing my Counterterrorism account for six years now and am reckoned to be one of the most experienced among the analysts who share this account. Others are less experienced and often rely on my expertise.

Question 3: *What are my objectives?* My major objectives involve the timely discovery of any possible form of terrorist actions that are based on my defensible and persuasive analysis of the evidence I have been able to gather. My overriding objective is to correctly predict the occurrence of any such terrorist actions so that they may be prevented from occurring,

Question 4: *What kinds and how much information will I have access to?* I believe it correct to say that my account involves what is termed "all-source intelligence analysis". We have access to any form of classified intelligence information. Our "need to know" for most of this information is rarely challenged. In addition, we have facilities for gathering information from a variety of "open sources," including all the media and internet sources.

Question 5: *How much time do I have to complete my analysis?* I would not say that my work involves "current intelligence" in which some policy or decision making "customer" requires answers in a very short time. However, the matters addressed in my account have their own sense of urgency since we may initially discover possible terrorist actions that are planned to occur in a very short time. This is why we have emphasized the timeliness of our discovery of possible terrorist actions. We cannot wait to discover them by witnessing their occurrence.

Evidence in Search of Hypotheses

Mavis reads an article in today's Washington Gazette that concerns how safely radioactive materials are stored in this general area. Willard, the investigative reporter and author of this piece begins by noting how the storage of nuclear and radioactive materials is so frequently haphazard in other countries and wonders how carefully these materials are guarded here in the USA, particularly in this general area. In the process of his investigations the reporter notes his discovery that a canister containing cesium-137 has gone missing from the STEMQ Company in Maryland, just three days ago. The STEMQ Company manufactures devices for sterilizing medical equipment and uses cesium-137 in these devices along with other radioactive materials. This piece arouses Mavis's curiosity because of her concern about terrorists planting dirty bombs in our cities. The question is: *What hypotheses would explain this*

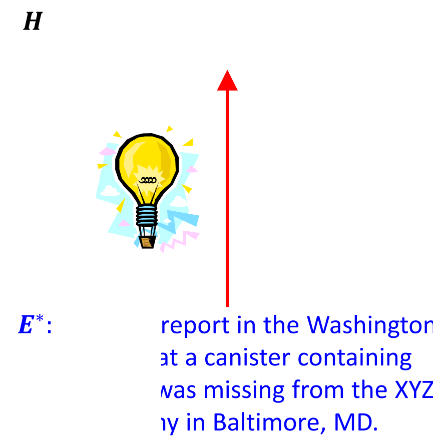

Figure 9. Hypothesis generation through imaginative reasoning.

observation? She experiences a flash of insight that a dirty bomb may be set off in the Baltimore-Washington D.C. area (see Figure 9). However, no matter how imaginative or important this hypothesis is, no one will take it seriously unless Mavis is able to justify it. She needs to build an argument linking the evidence to her hypothesis. In doing this, she is assisted by the six honest serving men from "The Elephant's Child” poem of Rudyard Kipling [1865 - 1936]:

I kept six honest serving men,
(They taught me all I knew),
Their names are What and Why and When
and How and Where and Who.

Starting from E* up, Mavis’s hypothesis generation proceeds step by step as shown inFigure 10.

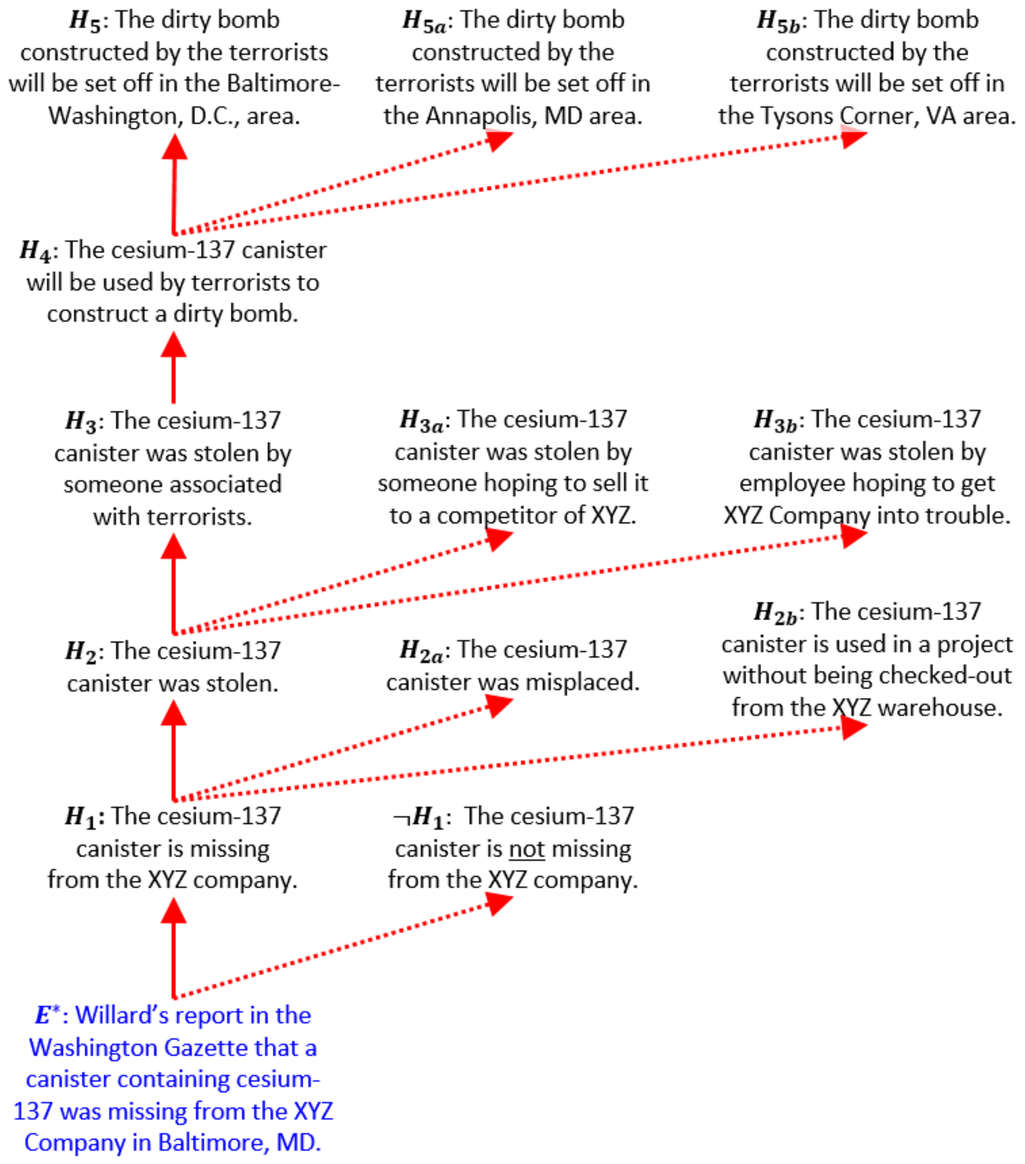
H_5: The dirty bomb constructed by the terrorists will be set off in the Baltimore-Washington, D.C., area.
H_{5a}: The dirty bomb constructed by the terrorists will be set off in the Annapolis, MD area.
H_{5b}: The dirty bomb constructed by the terrorists will be set off in the Tysons Corner, VA area.
H_4: The cesium-137 canister will be used by terrorists to construct a dirty bomb.
H_3: The cesium-137 canister was stolen by someone associated with terrorists.
H_{3a}: The cesium-137 canister was stolen by someone hoping to sell it to a competitor of XYZ.
H_{3b}: The cesium-137 canister was stolen by employee hoping to get XYZ Company into trouble.
H_{2b}: The cesium-137 canister is used in a project without being checked-out from the XYZ warehouse.
H_2: The cesium-137 canister was stolen.
H_{2a}: The cesium-137 canister was misplaced.
H_1: The cesium-137 canister is missing from the XYZ company.
$\neg H_1$: The cesium-137 canister is not missing from the XYZ company.
E^*: Willard’s report in the Washington Gazette that a canister containing cesium-137 was missing from the XYZ Company in Baltimore, MD.

Figure 10. Hypotheses generation.

Mavis asks herself a series of questions and the answers she provides leads to the abduction steps shown in the figure.

Why is it possible that the cesium-137 canister is missing? Missing radioactive materials is not a rare event. Washington Gazette is usually a credible source and there has been no denial of this report from the STEMQ Company. So, Mavis reasons as if hypothesis H_1 is true: The cesium-137 canister is missing.

What might have happened tothe cesium-137 canister?

1. The cesium-137 canister was stolen.
2. The cesium-137 canister was misplaced.
3. The cesium-137 canister is used in a project without being checked-out from the STEMQ warehouse.

Why am I inferring $\boldsymbol{H_2}$ *from* $\boldsymbol{H_1}$(that the cesium-137 canister was stolen)?

1. My standpoint objective makes me naturally suspicious when radioactive materials go missing. Things of value that are missing might have been stolen, thus theft is a plausible explanation.
2. The records show that the STEMQ Company has never lost any cesium-137 in the past.
3. There has been no denial of this report from the STEMQ Company.

Why is she favoring the stolen hypothesis over the others? Because of her standpoint *objectives and the fact that theft is plausible*; also, because substances such as cesium-137 are very rarely misplaced or used without this being known. Now Mavis reasons as if hypothesis H_2 is true: The cesium-137 canister was stolen.

Who might have stolen the cesium-137 canister?

1. The cesium-137 canister was stolen by someone associated with terrorists.
2. The cesium-137 canister was stolen by someone hoping to sell it to a competitor of XYZ.
3. The cesium-137 canister was stolen by employee hoping to get STEMQ Company into trouble.

Why am I inferring $\boldsymbol{H_3}$ *from* $\boldsymbol{H_2}$(that the cesium-137 canister was stolen by someone associated with terrorists)?

1. My experience and standpoint allow me to strongly suspect that terrorists might have a hand in the theft of substances like cesium-137. Terrorist groups themselves do not have the facilities necessary to produce radioactive materials, and these substances cannot be purchased on the open market.
2. This is possible but is outside of Mavis's responsibility.
3. This is also possible but is outside of Mavis's responsibility.

It would seem preposterously unlikely that terrorists would have a benign use for radioactive substances. From her experience she easily infers $\boldsymbol{H_4}$, that the terrorists who stole the cesium-137 plan to use it in the construction of a dirty bomb. The reason, of course, is that cesium-137 is a radioactive material. If dispersed in an explosion it could cause great panic and render the surrounding area uninhabitable for decades. Then a very reasonable guess is hypothesis $\boldsymbol{H_5}$, that the terrorists will set off the dirty bomb they will have constructed somewhere in the Baltimore-Washington D.C., area. One reason is that the STEMQ Company, from which the cesium-137 canister was stolen, is in Baltimore. But this is a quite vague hypothesis and she knows that she will later have to refine by saying where in this area the bomb will be set off, when will it be set off, and by whom.

The chain of inferences from shows clearly the possibility that a dirty bomb will be set off in the Baltimore-Washington D.C., area. However, we cannot conclude that until we analyze all the alternative hypotheses and show that those on the chain from E* to $\boldsymbol{H_5}$ are actually the most likely ones. But to analyze all these alternative hypotheses and make such an assessment we need additional items of evidence. How can we get them? As represented in the middle of Figure 6, we put each hypothesis at work to guide us in the collection of additional evidence. This process is illustrated in the next section.

Hypotheses in Search of Evidence

Let us first consider the hypothesis H_1: The cesium-137 canister is missing from the STEMQ company, shown again at the top of Figure 11. The question is: *Assuming that this hypothesis is true, what other things should be observable?* What are the necessary conditions for an object to be missing from a warehouse? It was in the warehouse, it is no longer there, and no one has

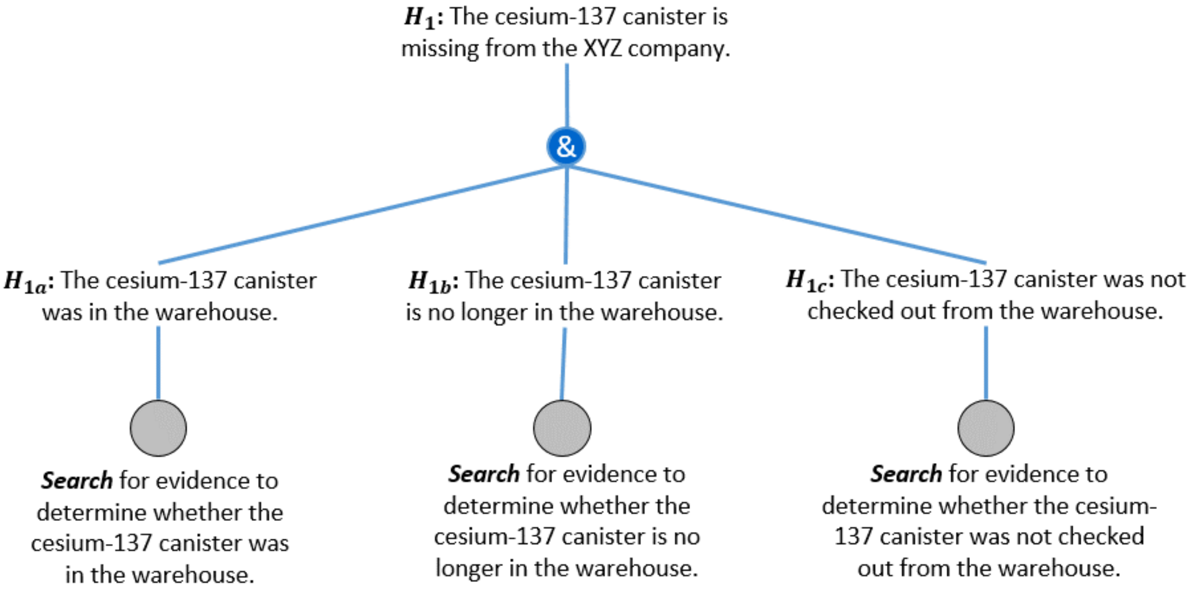
$\boldsymbol{H_1}$: The cesium-137 canister is missing from the XYZ company.
&
$\boldsymbol{H_{1a}}$: The cesium-137 canister was in the warehouse.
$\boldsymbol{H_{1b}}$: The cesium-137 canister is no longer in the warehouse.
$\boldsymbol{H_{1c}}$: The cesium-137 canister was not checked out from the warehouse.
Search for evidence to determine whether the cesium-137 canister was in the warehouse.
Search for evidence to determine whether the cesium-137 canister is no longer in the warehouse.
Search for evidence to determine whether the cesium-137 canister was not checked out from the warehouse.

Figure 12. Hypothesis-driven evidence collection.

checked it out. This suggests the decomposition of the hypothesis H_1 into the conjunctive of three simpler hypotheses, as shown in Figure 11. This clearly indicates that you should look for evidence that indeed the cesium-137 canister was in the warehouse, that it is no longer there, and that no one has checked it out.

Guided by the evidence collection tasks at the bottom of Figure 11, Mavis contacst Ralph, the supervisor of the STEMQ warehouse, who provides the information shown in Table 1.

Table 1. Information obtained through the collection tasks in Figure 11.

I1 Ralph: Contacted about the cesium-137 canister Ralph, the supervisor of the warehouse, who has a good reputation for honesty, reports that the cesium-137 canister is registered as being in the warehouse, that no one at the STEMQ Company had checked it out, but it is not located anywhere in the hazardous materials locker. He also indicates that the lock on the hazardous materials locker appears to have been forced.

When we are given testimonial information, or descriptions of tangible items, the information might contain very many details, dots, or trifles. Some of the details might be interesting and relevant evidence, and others not. What we always have to do is to parse the information to extract the information that we believe is relevant in the inference task at hand. Consider, for example, the information provided by Ralph's testimony from Table 1. It provides us with the dots or items of evidence from Table 2 that are relevant to assessing the hypothesis H_1: The cesium-137 canister is missing from the STEMQ Company.

Table 2. Dots or items of evidence obtained from Willard and Ralph.

E1 Washington Gazette: Willard's report in the Washington Gazette that a canister containing cesium-137 was missing from the STEMQ warehouse in Baltimore, MD.

E2 Canister registered: Ralph, who has a reputation for honesty, reports that the cesium-137 canister is registered as being in the warehouse.

E3 Not in locker: Ralph, who has a reputation for honesty, reports that the cesium-137 canister […] is not located anywhere in the hazardous materials locker.

E4 Not checked-out: Ralph, who has a reputation for honesty, reports that […] no one at the STEMQ Company had checked it out.

E5 Forced lock: Ralph, who has a reputation for honesty, reports that lock on the hazardous materials locker appears to have been forced open.

Evidentiary Testing of Hypotheses

To assess the hypotheses, we first need to attach each item of evidence to the hypothesis to which it is relevant, as shown in Figure 12. Then we need to establish the relevance and the credibility of each item of evidence which will result in the inferential force of that item of evidence on the corresponding hypothesis, as explained below.

Here is, for example, how we can assess the relevance of E2 Canister registered: If the canister is registered in the warehouse, as Ralph claims, then almost certainly the canister was in the warehouse. There may be some exceptions, such as a mistake in the records, but we would expect them to be extremely rare. We next have to assess the credibility of E2. In general, the credibility of an item of evidence depends on the credibility of its source. The source of E2 is Ralph who is the supervisor of the warehouse and has a reputation for honesty. We will therefore assess his credibility as very likely. Based on the credibility and the relevance of evidence, Cogent determines the inferential force of E2 on H_{1a} as very likely (i.e., the minimum of almost certain and very likely), as shown in Figure 12. The inferential forces of the other items of evidence are determined in a similar way.

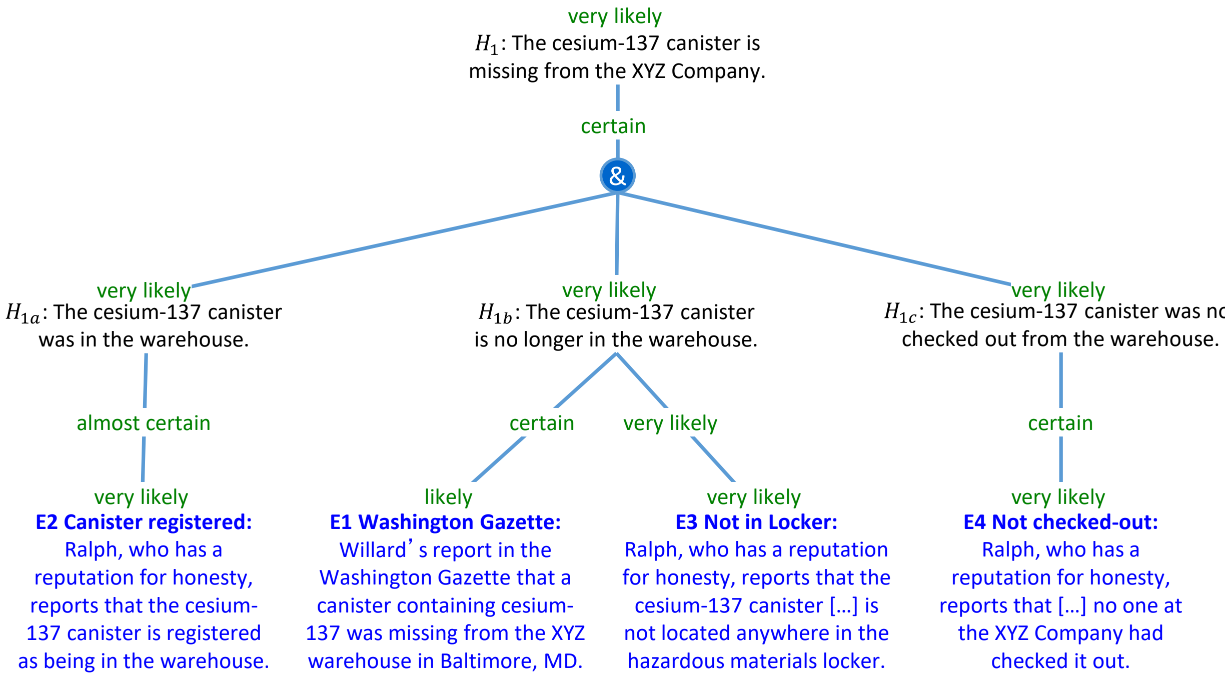
very likely
H_1: The cesium-137 canister is missing from the XYZ Company.
certain
&
very likely
H_{1a}: The cesium-137 canister was in the warehouse.
very likely
H_{1b}: The cesium-137 canister is no longer in the warehouse.
very likely
H_{1c}: The cesium-137 canister was nc checked out from the warehouse.
almost certain
certain
very likely
certain
very likely
E2 Canister registered:
Ralph, who has a reputation for honesty, reports that the cesium-137 canister is registered as being in the warehouse.
likely
E1 Washington Gazette:
Willard's report in the Washington Gazette that a canister containing cesium-137 was missing from the XYZ warehouse in Baltimore, MD.
very likely
E3 Not in Locker:
Ralph, who has a reputation for honesty, reports that the cesium-137 canister […] is not located anywhere in the hazardous materials locker.
very likely
E4 Not checked-out:
Ralph, who has a reputation for honesty, reports that […] no one at the XYZ Company had checked it out.

Figure 13. Evidence-based hypotheses assessment.

Notice in Figure 12 that there are two items of evidence that are relevant to the hypothesis H_{1b}. In this case, the probability of H_{1b} is the result of the combined inferential force of these two items of evidence.

We then need to assess the relevance of the AND argument of the top hypothesis. In this case the conjunction of the three sub-hypotheses represents a sufficient condition for the top hypothesis. This means that, if the sub-hypotheses are true, then the top hypothesis is also true. Thus, the relevance of this argument is certain. Based on the probabilities of the three hypotheses and on the relevance of the AND argument, Cogent determines that it is very likely that the cesium-137 canister is missing from the STEMQ warehouse.

Notice that this is a process of multi-INT fusion of evidence since, in general, the assessment of a hypothesis involves combining information from different types of evidence.

The analysis continues with the next abductive step, and the assessment of the corresponding alternative hypotheses using the same processes of hypotheses in search of evidence and evidentiary testing of hypotheses. The next chapters of this book include exercises for completing this analysis which will further illustrate the synergistic integration of analyst's imagination with computer's critical reasoning.

In conclusion, the computational theory of intelligence analysis presented in this volume, as well as its current implementation in Cogent, provide a framework for integrating the art and science of intelligence analysis, to cope with its astonishing complexity. However, while the computational theory and Cogent guide you through the intelligence analysis steps, and also automates many of them, it requires you to continuously exercise your imagination. Therefore, in the next chapter we return to this all important capability to describe useful heuristics for marshaling your thoughts and evidence. Using such heuristics in conjunction with a cognitive assistant like Cogent is the approach we advocate for coping with the astonishing complexity of "connecting the dots."

1.6 Review Questions

1. Characterize each of the questions below with respect to the number of answers it can have.
   Q_1: Will the president select General Martin to be the country's next defense minister?
   Q_2: Which of the country's four-star generals is the president likely to nominate as the country's next defense minister?
   Q_3: Why did the president select General Martin to be the next defense minister?

2. A terrorist incident occurred two weeks ago in an American city involving considerable destruction and some loss of lives. After an investigation, two foreign terrorist groups have been identified as possible initiators of this terrorist action: an Al Qaeda-affiliated Group A from Yemen, and a Taliban Group B from Pakistan. What are some hypotheses we could entertain about this event?

3. You might have reason to suspect that Iran is now supplying IEDs to a Taliban group in Afghanistan. Since there are other possible sources for these weapons you will have more than one major hypothesis about possible suppliers of these IEDs. What are some of these other hypotheses?

4. Consider the hypothesis that Iran is now supplying IEDs to a Taliban group in Afghanistan. What evidence we might find concerning this hypothesis?

5. Consider the hypothesis that the Al Qaeda-affiliated Group A from Yemen was the one

involved in the terrorist incident. What evidence we might find concerning this hypothesis?

6. Sometimes we have evidence in search of hypotheses or possible explanations. For example, consider the dog-tag containing the name of one of our soldiers who has been missing since the end of our conflict with Country Z. This tag was allegedly given to a recent visitor in Country Z who then gave it to us. One possibility is that this soldier is still being held as a prisoner in Country Z. What are some other possibilities?

7. Sometimes we have hypotheses in search of evidence. Suppose our hypothesis is that Person X was involved in the terrorist incident. So far, all we have is evidence that he was at the scene of the incident an hour before it happened. If this hypothesis were true, what other kinds of evidence might we be able to observe about X?

8. True or false: Source A provides information on subject B. If source A is a longstanding enemy of subject B, the credibility of this information, all other things being equal, should be increased.

9. True or false: The relevance of evidence is an assessment of the extent to which the evidence may be believed.

10. Inferential force is an assessment that takes into account:
    a) the credibility of evidence
    b) the relevance of evidence
    c) both the credibility and relevance of evidence

11. Consider the hypothesis that Countries A and B are about to engage in armed conflict. Here is a report you have just obtained; it says that there has just been an attempt on the life of the president of Country B by an unknown assailant. Why is this report, if credible, relevant evidence on the hypothesis that Countries A and B are about to engage in armed conflict?

12. A car bomb was set off in front of a power sub-station in Washington DC on 25 November. The building was damaged but, fortunately, no one was injured. From the car's identification plate, which survived, it was learned that the car belonged to Quick Car Rental Agency. From information provided by Quick, it was learned that the car was last rented on 24 November by a man named M.

    a) Construct an argument from this evidence (E*) to the hypothesis that person M was involved in this car-bombing incident.

    b) Suppose that we have determined that evidence E* is believable and therefore we think that M indeed rented a car on November 24. We need additional evidence to assess F, which states that M drove the car on November 25. As discussed in Section

1.5 and illustrated in Section 0, we can use this hypothesis to guide us in collecting new evidence. Employ this approach to find the needed evidence.

13. Defendant Dave is accused of shooting a victim Vic. When Dave was arrested sometime after the shooting, he was carrying a 32 cal. Colt automatic pistol. Let H be the hypothesis that it was Dave who shot Vic. A witness named Frank appears and says he saw Dave fire a pistol at the scene of the crime when it occurred; that's all Frank can tell us.

    a) Construct a simple chain of reasoning that connects Frank's report to the hypothesis H that it was Dave who shot Vic.

    b) The chain of reasoning that connects Frank's report to the hypothesis that it was Dave who shot Vic shows only the possibility of this hypothesis being true. What are some alternative hypotheses?

    c) In order to prove the hypothesis that it was Dave who shot Vic, we need additional evidence. As discussed in Section 1.5 and illustrated in Section 0, we need to use this hypothesis to guide us in collecting new evidence. Employ this approach to find the needed evidence.

    d) Our investigation has led to the discovery of additional evidence. By itself, each evidence item is hardly conclusive that Dave was the one who shot Vic. Someone else might have been using Dave's Colt automatic. But Frank's testimony along with the fact that he was carrying his weapon, and with the ballistics evidence puts additional heat on Dave. Assess the probability of the hypothesis that Dave was the one who shot Vic.

14. Justify the assessments of relevance of E1 Washington Gazette, E3 Not in Locker, and E4 Not checked-out shown inTable 2.

15. Justify the assessments of credibility of E1 Washington Gazette, E3 Not in Locker, and E4 Not checked-out shown in **Error! Reference source not found.**.

16. True or false: The relevance of the evidence “Willard, who is an unverified source, said that a canister containing cesium-137 is missing from the STEMQ warehouse” to the hypothesis “The canister is no longer in the warehouse” is certain.

17. If the credibility of an item of evidence is low and the relevance of this evidence to a hypothesis is high, the inferential force of this evidence is:
    a) high
    b) medium
    c) low

18. True or false: In the problem about the missing cesium canister, Ralph's reputation for honesty must be taken into account when assessing the relevance of the information provided by Ralph.

19. In the problem about the missing cesium canister, the credibility of information provided by Ralph (very likely) was assessed at a higher level than the information in Willard's article in the Washington Gazette (likely) because:
    a) Ralph has a reputation for honesty. His position at the STEMQ warehouse is not pertinent to the information's credibility.
    b) Ralph has first-hand access to the information. His reputation for honesty is not pertinent to the information's credibility.
    c) Ralph has a reputation for honesty and has first-hand access to the information.

20. True or false: If Willard had a brother who was fired by the STEMQ warehouse four years ago, the credibility of the information in Willard's report would increase.

21. Consider the following hypothesis and items of evidence:

    Hypothesis: The canister is no longer in the warehouse.

    E1: Willard's report in the Washington Gazette that a canister containing cesium-137 was missing from the STEMQ warehouse in Baltimore, MD.

    E2: Ralph, who has a reputation for honesty, reports that the cesium-137 canister […] is not located anywhere in the hazardous materials locker.

    True or false: *E1* should have lower relevance to the hypothesis than *E2*.

22. The argument that "The cesium-137 canister is missing from the STEMQ warehouse" needs to establish:
    a) only that the canister is not in the warehouse;
    b) only that the canister was in the warehouse but is no longer there;
    c) that the canister was in the warehouse but is no longer there and was not checked out from the warehouse.

2 MARSHALING THOUGHTS AND EVIDENCE FOR IMAGINATIVE ANALYSIS

2.1 Investigation and Evidence Marshaling

2.1.1 Sherlock Holmes and Investigation or Discovery

If you have read any Sherlock Holmes mysteries, you know that Holmes had several foils in the form of rather incompetent police investigators such as Inspector Lestrade and Inspector Gregory. Confidently believing that they had a case solved, Lestrade or Gregory had obviously overlooked details that were observed and then imaginatively analyzed by Holmes. In one case, *The Boscombe Valley Mystery* (Baring-Gould W.S., 1967, Vol. II, page 148), Holmes tells his colleague Dr. Watson:

> *"By an examination of the ground I gained the trifling details which I gave to that imbecile Lestrade, as to the personality of the criminal. Watson asks: But how did you gain them? Holmes replies: You know my method. It is founded on the observance of trifles."*

As an intelligence analyst you are confronted daily with hundreds, perhaps thousands of trifles or details. We could also refer to Sherlock Holmes' observed trifles as one form of dot that we must try to connect. Taken alone, an individual trifle may mean very little. But some of them, taken in combination, may suggest new and important hypotheses or possibilities that should be taken very seriously. *The trick is to be able to identify which combinations of trifles to examine carefully and which ones to ignore*. This is where alternative schemes for selecting and marshaling trifles, together with our thoughts about them, become all important. As we noted in Section 1.3.4, it would not make any sense to examine all possible combinations of trifles, even if we could do so. Considerable imagination is required, both in deciding which trifle combinations to examine and in generating new and productive hypotheses from the trifle combinations you have identified.

Sherlock Holmes also seemed particularly adept at asking questions as his investigations unfolded. The process of discovery involving the generation of new hypotheses or possibilities rests crucially on the questions we ask. One noted logician claims that Sherlock Holmes imaginative feats of skill were largely due to his skill at inquiry, the asking of questions (Hintikka, 1983, pp 170-178). If we do not ask appropriate questions, as our intelligence investigations unfold, we stand little chance, without an abundance of luck, of generating hypotheses that stand some chance of containing truth.

2.1.2 Mycroft Holmes and Evidence Marshaling

It is here that we must introduce Sherlock Holmes' older brother Mycroft Holmes. We do not hear much about Mycroft since he only appears in two of the Holmes mystery stories, *The Greek*

Interpreter (Baring-Gould W.S., 1967, Vol. I, pp. 590 – 605) and *The Bruce-Partington Plans* (Baring-Gould W. S., 1967, Vol. II, pp. 432 – 452). On first impressions, Mycroft appears as a minor civil servant working as an auditor for various British government departments in Whitehall. But Sherlock says that this would be dramatically misleading since, as Sherlock admits, Mycroft's investigative and inferential capabilities are far greater than his own. Mycroft's true role was obviously kept a closely guarded secret. We get the clearest account of Mycroft's capabilities in *The Bruce-Partington Plans.* Sherlock says that Mycroft is in fact the most indispensable man in the country. One reason is that Mycroft has the tidiest and most orderly brain with the greatest capacity for storing facts of anyone living. Further, the conclusions of every governmental department are passed to Mycroft who serves as a central exchange or a clearinghouse that makes out a balance. In examining these various inputs, Mycroft can focus on them all and say how each input would influence the others. In Mycroft's brain everything is pigeon-holed and can be accessed instantly. Sherlock says that again and again Mycroft's word has decided national policy and that, on occasion, Mycroft has been <u>the</u> British government. So, one way to describe Mycroft's major capability is to say that he had superlative skills in marshaling masses of evidence and in generating correct conclusions from this marshaled evidence. Sherlock said of Mycroft, "All other men are specialists, but his specialism is omniscience."

Perhaps the most frequently overlooked element of intelligence analysis concerns the manner in which we marshal or organize our thoughts and our evidence as we proceed with some analytic task. But such oversight causes no end of difficulties since how skillful we are in marshaling our <u>existing</u> thoughts and evidence greatly influences how skillful we will be in generating or discovering <u>new</u> ideas (in the form of possible hypotheses) and <u>new</u> lines of inquiry and evidence. Skillful evidence marshaling is not only necessary during the discovery-related processes of intelligence analysis, it forms the very basis for the later task of constructing defensible and persuasive arguments on hypotheses we entertain. In short, developing useful strategies for marshaling thoughts and evidence during intelligence analysis is absolutely crucial.

Concern about means for marshaling our thoughts and our evidence arises for two major reasons. The first is the fact that marshaling methods can assist us in being more imaginative during the process of discovery as we are attempting to determine what has happened or what will happen in some situation of interest to us. Appropriate marshaling strategies can assist us in generating productive new hypotheses and new lines of inquiry and evidence. Second, evidence marshaling strategies are key ingredients of the task of determining what our evidence means and in constructing defensible and persuasive arguments on hypotheses we are considering. This is where the marshaling of ideas in addition to evidence becomes so important. We must generate chains of reasoning whose ingredients consist of ideas we have in showing how we believe the evidence we have is linked to, or is relevant on, hypotheses we are considering. This is why we say that what we are marshaling are our <u>thoughts</u> and our <u>evidence</u>. It is quite obvious that marshaling thoughts and evidence is a major task in the process of “connecting the dots.”

Concern about thought and evidence marshaling is not of course limited to intelligence analysis. As indicated in Section 1.3, many noted persons in science and mathematics (including several Nobel Laureates) have emphasized the importance of combining ideas during the process of discovery in which new ideas and new lines of evidence are being generated. All of them emphasize the fact that new ideas frequently result from particular combinations of evidence and ideas we already have. The trouble is that we can never look through all possible combinations of information that we have, even if we could do so. Indeed, this would be the attempt to look through everything in the hope of finding something. Most combinations of trifles would be meaningless anyway; just a few combinations might lead to startling and productive new ideas. How do we decide which combinations of trifles or details to examine? Here is where the process of imaginative reasoning is so necessary and where evidence marshaling becomes so important. Our imaginative reasoning begins to be applied by the questions we ask of and about the evidence we already have. Questions we ask about our evidence help us to determine the three major credentials of evidence already mentioned: relevance, credibility, and inferential force. Questions we ask of our existing evidence allow us to generate new hypotheses and new lines of inquiry and evidence. We cannot productively ask these appropriate questions unless we have marshaled or organized our existing thoughts in meaningful ways.

In most intelligence analyses the events of concern are singular, unique, or one-of-a-kind. What this means is that there are very few if any statistical records available to allow us to predict events that are of concern to our nation's defense. There were no statistics available that would have allowed us to forecast the tragic events that took place on September 11, 2001. In many discussions in the field of artificial intelligence concerning the process of discovery, it is claimed that all discovery amounts to is having sophisticated methods of search. In such discussions, something crucial is left out, namely the process of inquiry, the asking of questions. *Having productive search methods is necessary but not sufficient during the process of discovery in many fields, including intelligence analysis*. The reason is that, absent any relevant statistical or other prior records, we may have nothing to search as some intelligence analysis task begins. We only begin to have relevant evidence to search through when we begin to ask questions about the situation of concern. No discovery process, in any discipline, can proceed in the absence of someone asking important questions. The issue then is: *How do we become more skillful in forming the questions we ask of and about our existing evidence and about the situation(s) of interest to us?* Answers to this question are supplied in part by the strategies we employ in marshaling our thoughts and our evidence in different ways.

2.2 Abductive (Imaginative) Reasoning

2.2.1 Peirce on Abductive Reasoning

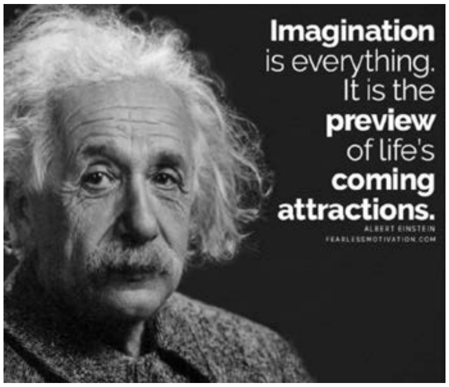

Until the time of Peirce, most persons interested in discovery and investigation supposed that the discovery-related arms of the arch in Figure 5 involved some form of inductive reasoning that proceeds from particulars (in the form of observations) to generalities (in the form of hypotheses). But inductive reasoning is commonly associated with the process of justifying or trying to prove existing hypotheses based on evidence. The question remains: Where did these hypotheses come from? Pondering such matters, Peirce relied on a figure of reasoning he found in the works of Aristotle. The reasoning proceeds as follows:

- If H were true, then E, F, and G would follow as a matter of course.
- But E, F, and G have been observed.
- Therefore, we have reason to believe that H might possibly be true.

Peirce was unsure about what to call this form of reasoning. At various points in his work he called it abduction, retroduction, and even just hypothesis (Pierce, 1998; 1901). The essential interpretation Peirce placed on the concept of abduction is illustrated in Figure 14.

He often used as a basis for his discussions of abduction the observation of an anomaly in science. Let us suppose that we already have a collection of prior evidence in some investigation and an existing collection of hypotheses $H_1, H_2, \dots, H_n$. To varying degrees, these *n* hypotheses explain the evidence we have so far. But now we make an observation E^* that is embarrassing in the following way: We take E^* seriously, but we cannot explain it by any of the hypotheses we have generated so far. In other words, E^* is an anomaly. Vexed by this anomaly, we try to find an explanation for it. In some cases, often much later when we are occupied by other things, we

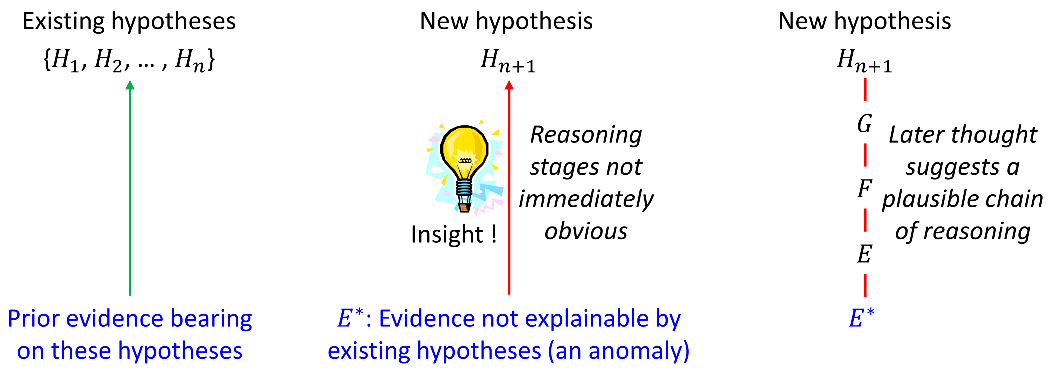
Existing hypotheses
$\{H_1, H_2, \dots, H_n\}$
Prior evidence bearing on these hypotheses
New hypothesis
H_{n+1}
Insight !
Reasoning stages not immediately obvious
E^*: Evidence not explainable by existing hypotheses (an anomaly)
New hypothesis
H_{n+1}
G
F
E
E^*
Later thought suggests a plausible chain of reasoning

Figure 14. Peirce's interpretation of abductive reasoning.

experience a "flash of insight" in which it occurs to us that a new hypothesis H_{n+1} could explain this anomaly E^*. It is these "flashes of insight" that Peirce associated with abduction. Asked at this moment to say exactly how H_{n+1} explains E^*, we may be unable to do so. However, further thought may produce a chain of reasoning that plausibly connects H_{n+1} and E^*. The reasoning might go as follows:

- I have evidence E^* that event E happened.
- If E did happen, then F might be true.
- If F happened, then G might be true.
- And if G happened, then H_{n+1} might be true.

It is possible that the chain of reasoning might have started at the top with H_{n+1} and ended at E^*. This is why we have shown no direction on the links between E^* and H_{n+1} in Figure 14.

But our discovery-related activities are hardly over just because we have explained this anomaly. Our new hypothesis H_{n+1} would not be very appealing if it only explained anomaly E^*. Figure 15 shows the next steps in our use of this new hypothesis. We first inquire about the extent to which it explains the prior evidence we collected before we observed E^*. An important test of the suitability of the new hypothesis H_{n+1} involves asking how well this new hypothesis explains other observations we have taken seriously. This new hypothesis would be especially valuable if it explains our prior evidence better than any of our previously generated hypotheses. But there is one other most important test of the adequacy of a new hypothesis H_{n+1}:

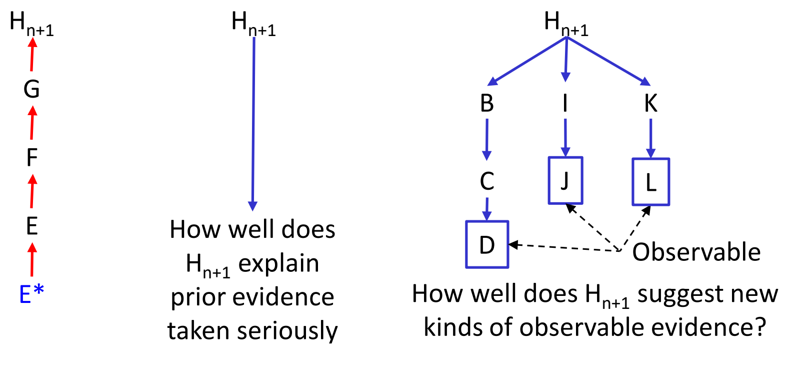
H_{n+1}
G
F
E
E*
H_{n+1}
How well does H_{n+1} explain prior evidence taken seriously
H_{n+1}
B
I
K
C
J
L
D
Observable
How well does H_{n+1} suggest new kinds of observable evidence?

Figure 15. Putting an abduced hypothesis to work.

> *How well does this new hypothesis suggest new potentially observable evidence that our previous hypotheses did not suggest?*

If H_{n+1} would be true, then B, I, and K would also be true; and if B would be true, then C would be true. Now if C would be true, then we would need to observe D.

In the illustrations Peirce used, which are shown in Figure 14 and Figure 15, we entered the process of discovery at an intermediate point when we already had existing hypotheses and evidence. In other contexts we must of course consider abductive reasoning from the beginning of an episode of fact investigation when we have no hypotheses and no evidence bearing on them. Based on our initial observations, by this process of abductive or insightful reasoning, we may generate initial guesses or hypotheses to explain even the very first observations we make.

Such hypotheses may of course be vague, imprecise, or undifferentiated. Further observations and evidence we collect may allow us to make an initial hypothesis more precise and may of course suggest entirely new hypotheses.

These strategies of abductive reasoning, insight, and discovery described by Peirce seem almost identical to those of Sherlock Holmes, Arthur Conan Doyle's fictional character from his many mystery stories. Holmes did not, of course, describe his investigative reasoning as abductive. Instead he said his reasoning was "backward," moving from his observations to possible explanations for them. In spite of the similarity of Peirce's and Holmes's (Conan Doyle's) views of discovery-related reasoning, there is no evidence that Peirce and Conan Doyle ever shared ideas on the subject. A very informative and enjoyable collection of papers on the connection between Peirce and Sherlock Holmes appears in a work of Umberto Eco and Thomas Sebeok (1983).

2.2.2 Umberto Eco on Abductive Reasoning

One of Eco's chapters in *The Sign of Three: Dupin, Holmes, Peirce* is entitled *"Horns, Hooves, Insteps: Some Hypotheses on Three Types of Abduction"* (Eco and Sebeok, 1983, pp.198-220). In this chapter, Eco actually mentions *four* types of abductive reasoning that we will illustrate in the context of intelligence analysis. What distinguishes the first three of Eco's forms of abductive reasoning concerns this question: *How creative is the new idea or hypothesis being generated?* Generated ideas or hypotheses may vary in the extent to which they actually say something new. Eco first restates the figure of reasoning that Peirce drew from Aristotle (Anderson and Twining, p.443). This restatement supplies us with a form of reasoning, according to Peirce, for the upward arm of the arch of knowledge illustrated in Figure 5:

- A surprising fact E is observed (a result).
- If H were true, then E would follow as a matter of course (a rule).
- Hence, we have reason to suspect that H is true (a case).

Peirce's abduction in this reasoning pattern involves inferring a case from a result, the case being an element of a rule. In the following account of Eco's four species of abduction we will use an evidence E^* to indicate that event E occurred. E^* might not be entirely credible evidence of event E.

2.2.3 Overcoded Abduction

The first, and least creative form of abductive reasoning is said by Eco to be "overcoded" because it is based on an association or contiguity that has been observed in the past. What may sound like a new idea is often the result of applying already existing knowledge of what Eco calls "prior contiguities." In such instances, we exploit already existing experience-based rules or generalizations in order to determine what an evidence item might mean in some new situation.

The form of abductive reasoning in intelligence analysis in this case goes as follows:

- We observe evidence E^* that event E occurred.
- Based on our prior knowledge of contexts in which things like event E have occurred, we say: "Whenever something like H has occurred, then something like E has also occurred." Rephrased, we might say: "If H were true, then E would follow as a matter of course.''
- Thus, there is reason to suspect that H may explain the occurrence of evidence E^*. In other words, evidence E^* points to H as a possible explanation for its occurrence.

Thus, in this case we have exploited knowledge of the past contiguity of events E and H to account for evidence E^*. Eco says that this form of abduction often occurs automatically or semi-automatically and that the degree of creativity in the generation of H in this case is minimal. For example, finding that the detonator of an improvised explosive device (IED) is similar to those previously employed in Iraq that are known to be of Iranian origin (E^*) can easily be explained by the hypothesis (H) that Iran is supplying IEDs to a Taliban group in Afghanistan. One reason why we might regard the reasoning in this case as abductive rather than inductive is that other things may be associated with E about which we have no present awareness. For example, we don't know how long this detonator was where we found it. It might have been there before the explosion. Further investigation may reveal that this detonator was planted at the scene by someone (i.e., E^* is not authentic evidence). What we are grading at this point is just the plausibility of H and not its probability. Eco goes a bit further with overcoded abduction by mentioning that such abductions can involve multiple items of evidence and related generalizations based on knowledge of past associations. For example, we might observe as evidence that person P did E, person Q did F, and person R did G. We may puzzle over what appear to be unrelated facts until someone notes that, based on experience, whenever things like H have occurred, then events like E, F, and G have also happened. The abduction of H here is still overcoded since reliance is placed on knowledge of a past association or contiguity.

2.2.4 Undercoded Abduction

This involves situations in which we may have alternative past associations or contiguities to draw upon. In other words, there may be several possible experience-based meanings we could attach to an observed evidence. Eco says this represents a higher degree of creativity than overcoded abduction since which rule to apply may not be so obvious. Here is an abstract example of the reasoning in undercoded abduction:

- We observe evidence E^* that event E occurred.
- G_1: If H were true, event E might follow;
- G_2: If J were true, event E might follow;
- G_3: If K were true, event E might follow.
- We decide that G_2 is most plausible based on what we know so far, thus we conclude that

J is the most plausible explanation of evidence E^*.

This may sound like “inference to the best explanation" (Josephson and Josepson, 1984), but we will certainly not argue that hypothesis J is the "best" explanation until we have canvassed all possible explanations for E^*, something we may have trouble doing. One problem with saying that abductive reasoning is inference to the best explanation is that we may not have any settled criterion for saying what is the "best" explanation. Nor do we often have assurance that we have canvassed all possible explanations. It seems better to say that, of the possibilities we have considered so far, J seems most plausible as an explanation for E^*. As an example, consider a terrorist incident that occurred two weeks ago in an American city involving considerable destruction and some loss of lives. Based on past experience, we know that several terrorist groups have employed this form of action in the past. To further the investigation at this point, we decide to focus the investigation on one of these groups. Eco notes that our choice of possible explanations for some form of evidence might involve more than one possibility. In a long listing of possibilities, we may not feel forced to choose only one as an explanation of evidence E^*. For example, we may have an experience-based generalization that reads: If H or J were true, event E would follow as a matter of course. Our conclusion would then be that H or J seem most plausible in explaining E.

2.2.5 Creative Abduction

In many situations, especially in intelligence analysis, there may be no single or alternative experience-based generalizations to support an abductive inference. In such instances, possible explanations for any form of observation must be generated de novo. Here is where real guesswork comes in, and we encounter situations that made Sherlock Holmes so famous. Very good examples of Holmes's creative abductions (which Holmes called "deductions") appear in a work by Sebeok and Urniker-Sebeok (1983) entitled *You Know My Method.* This work also contains some marvelous examples of the abductions made by Dr. Joseph Bell, the Edinburgh surgeon whom Doyle seems to have taken as a model for Sherlock Holmes. Here, in abstract form, is Eco's creative abductive reasoning:

- I observe evidence E^* that event E occurred.
- Having no prior experience-based generalizations to draw upon, I guess that if H were true, event E would follow as a matter of course.
- Thus, I have a hunch that H might be true.

We leave Eco for a moment to dwell upon a matter Peirce recognized while discussing the abductive generation of new ideas. Novelty alone is not enough as far as creative reasoning is concerned. We might easily generate some truly novel explanations for evidence E^* that have nothing else going for them apart from novelty. As summarized so carefully in the work by Rescher (1978), Peirce was also concerned about the actual productive value of new hypotheses as well as about the efficiency with which the abductive process is undertaken. We return again to Figure 15 concerning the process of putting a new hypothesis to work. This new guess or hunch we have as a possible explanation of evidence E^* is useful to the extent that it may help explain other evidence we already have and, in addition, informs us about what new evidence to search for. In his account of creative abduction, Eco recognized that we may make more than one guess, or have more than one hunch, about how to explain evidence E^*. How do we choose which of these guesses or hunches to act upon? Plausibility is certainly one important criterion. As shown in Figure 14, at the time we form a guess or hunch that some new hypothesis H_{n+1} might be true, the linkage between our evidence E^* and this new hypothesis might not be so apparent. Later thought may allow us to form a chain of reasoning from E^* to this new hypothesis. The plausibility of this new hypothesis will depend upon how defensible and persuasive is this chain of reasoning. If there are disconnects or non sequiturs in this chain of reasoning, our new hypothesis will not be taken seriously. However, in many works on the process of discovery in various disciplines, an *aesthetic* criterion is often discussed in situations in which more than one hypothesis may seem entirely plausible. For example, we have plausible guesses H_{n+1}, J, and K for clue E^*, all of which seem entirely reasonable. We argue that hypothesis H_{n+1} offers the most "elegant" or "beautiful" explanation of evidence E^*. This aesthetic criterion appears in many of Sherlock Holmes's accounts of why he favored one hypothesis over others he had entertained.

2.2.6 Meta-Abduction

The fourth form of abduction mentioned by Eco is actually not an increased gradation of creativity. Suppose we have guess or hunch H as an explanation of evidence E^*. We say: "From event E we reason that F might be true, from F we reason that G might be true, and from G we reason that H might be true." This additional form of abduction recognizes the fact that, in fields such as intelligence analysis, we must frequently act upon creative hunches or guesses, such as H, without always being able to verify that F and G are true. Eco says that this is one reason why there are more popular works regarding detective mysteries than there are regarding medical diagnoses. He goes on to say: "Detectives are rewarded by society for their impudence in betting by meta-abduction, whereas scientists are socially rewarded for their patience in testing their abductions" (Eco and Sebeok, 1983).

Figure 16 shows the relationship between Peirce's conception of abduction as the insightful reasoning that leads from an anomaly E^* to the generation of a new possibility H_{n+1}, and argument consruction from evidence E^* to H_{n+1} that itself has merit in suggesting new possibilities. In the construction of an argument we must first infer event E from clue E^*. Just because we have evidence of E does not mean that E happened. The remaining stages of this argument call for abductive reasoning, by which we imagine how event E might be linked to hypothesis H_{n+1}. We imagine that event E could mean event F, event F could mean event G, and event G could mean event H. There are two major benefits of such argument construction. The first is that, if our argument is plausible, we have linked evidence E^* with our guess or hunch that H_{n+1} is true. Second, we have identified two new lines of potential evidence on events F and G, which would be more direct evidence of H_{n+1} than our initial evidence E^*. So, there is definite heuristic value in argument construction, which also can involve abductive reasoning.

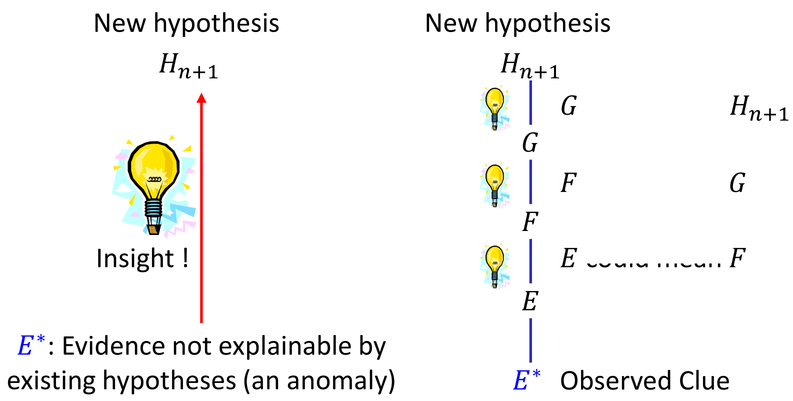

Figure 16. Abductive reasoning and argument construction.

2.2.7 Hypothesis Generation through Multi-step Abduction

Automatic hypothesis generation through abductive reasoning is computationally-intensive because there are so many hypotheses that can be generated (Josephson and Josephson, 1994; Schum, 2001a; Walton, 2005; Forbus, 2015; Langley, 2019). If we were to perform a single-step abduction, from evidence E to a hypothesis of interest that would explain it, we would obtain a huge number of hypotheses represented as dots at the top of Figure 17.

We would then need to investigate each of these competing hypotheses to find the most likely explanation. Now contrast this process with multi-step abduction. From E, one may abduce F_i, F$_j$, and F$_k$. At this point, we would search for evidence relevant to these three hypotheses and we would assess them based on the discovered evidence concluding, for example, that F_i is the most likely. Then we would continue the abduction from F_i, abducing G_l, G$_m$, and G$_n$, assessing these hypotheses, and concluding, for

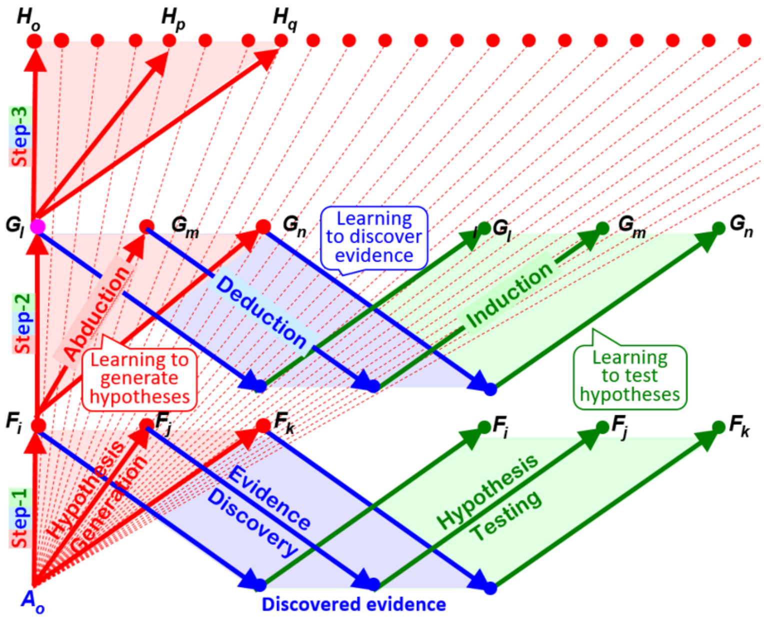

Figure 17. Hypothesis generation through multi-step abduction.

example, that G_l is the most promising. Finally, from G_l, we would abduce the hypotheses of interest H_o, H_p, and H_q, and assess them. This approach to hypothesis generation based on spiral hybrid reasoning, where small abductive, deductive, and inductive steps feed each other, significantly reduces the hypothesis space.

2.3 An Imaginative Analyst

We illustrate in this section the multi-step, multi-type abduction supported by deduction and induction.

2.3.1 Analyst's Standpoint

You have met the intelligence analyst Mavis in Section 1.5. Her work in a counter-terrorism office requires her to assist in monitoring the activities of organizations or individuals known or suspected of engaging in activities leading to episodes of violence against American citizens that could be labeled acts of terror. Some of these organizations or individuals are entirely domestic in origin such as the many hate groups in existence having strong antagonism against the government or against persons in various ethnic, racial, or religious groups. But some of these organizations or individuals have origins in foreign countries where there is a strong level of hatred against our entire democratic system. The many jihadist organizations in Pakistan, Afghanistan, and in other places supply ready examples. These anti-American jihadist organizations may employ foreign nationals who have come to America, legally or illegally, and who are already skilled in the construction and use of a variety of weapons. But these jihadist groups may also recruit American citizens and, in many cases, have provided them with weapons training in countries such as Pakistan and Afghanistan and then have encouraged them to return to America. In either case, we might refer to these persons as "sleepers", who await orders to engage in terrorist activities in homeland America. They may of course initiate terror actions on their own.

Mavis's standpoint (Anderson et al., 2005, pp. 115-117) depends of the situation she encounters, but in this case is the same with the one provided in Section 0. She has been following the activities of a person, Henry P., who is a leading figure in a domestic hate group called "*Aryan Militia*" [a fictitious hate group; we know of no domestic hate group with this name]. This group has its major activities in Northern Virginia, but it has operations in other areas of Virginia, and has connections with other hate groups in other states. Just yesterday, Mavis is shown a photo of Henry at the Sunny Valley restaurant in Leesburg, VA. The photo was taken two days ago by an FBI investigator. The photo shows Henry in company with another man. Who this other man is excites Mavis's curiosity greatly and leads her to draw a very surprising conclusion. What follows is a construction of Mavis's train of thought as she begins to explore the meaning of the evidence she has just received.

In our construction of Mavis's train of thought, we imagine that she informs us about stages or steps of her reasoning she forms based on the photo evidence she has received from the FBI. What we will consider is the argument or chain of reasoning she has constructed, over some period of time, that consists of a sequence of links she believes to form a logical connection of her evidence with her major conclusion. Careful attention to these links will be very important in illustrating a wide variety of evidential and inferential matters we need you to consider.

2.3.2 Induction

We begin with a question for you the reader: If you were Mavis, where would you begin in thinking about the photo evidence you just received? One very sensible and necessary place to begin is by considering the credibility of this evidence. Such consideration of credibility forms the very foundations of all arguments from evidence.

The photo evidence is an example of a type of evidence we have labeled tangible evidence, since Mavis can examine it for herself to see what events it reveals. Suppose the photo is crystal-clear and in perfect focus. But Mavis also needs to consider whether the photo is authentic: Is this photo what the FBI claims it to be?

In assessing the credibility of evidence you must always note the following. There is a difference between having evidence for some event and the actual occurrence of this event. Just because we have evidence for this event does not entail that the event did occur; there are various reasons depending on the source of the evidence and what kind of evidence we have. These matters are discussed in detail in Chapter 4 of this book. Here is what we have in Mavis's present situation. She has photo evidence, that we label E^*, that Henry of the Aryan Militia domestic hate group, was in company with another man in the Sunny Valley restaurant in Leesburg, VA two days ago. Then, we label as event E, that it was actually Henry of the Aryan Militia domestic hate group who was in company with another man in the restaurant in Leesburg, VA two days ago. So, we have evidence E^* about event E. Mavis's inference here concerns whether event E did occur, as the photo evidence E^* suggests. Here is where Mavis's concern about the authenticity of the photo evidence arises.

First, Mavis has often received evidence from the FBI and has never had any reason to question its authenticity. She does not believe photo evidence E^* is a fake, or that the photo was actually taken at a different time or place. Also, Mavis has had considerable experience with Henry and believes the photo does show Henry. So, based on this prior experience, Mavis believes that event E did occur. *Her reasoning here is inductive in nature since it is based entirely in her past experiences.* Figure 18 shows the first link in the chain of reasoning she constructs. This figure does show one other important element of Mavis's inductive inference, namely that it is also necessarily probabilistic in nature. She acknowledges that the FBI might have been wrong about when or where the photo was taken and that she might have been incorrect in identifying Henry. So, she allows for the fact that event E might not be true, that we label as $\neg E$. However, Mavis believes E to be so probable that she rests all further stages of her inference on the truth of E.

<u>Question for the reader</u>: Can you think of any other reasons why the photo evidence might not be authentic and event E not being true?

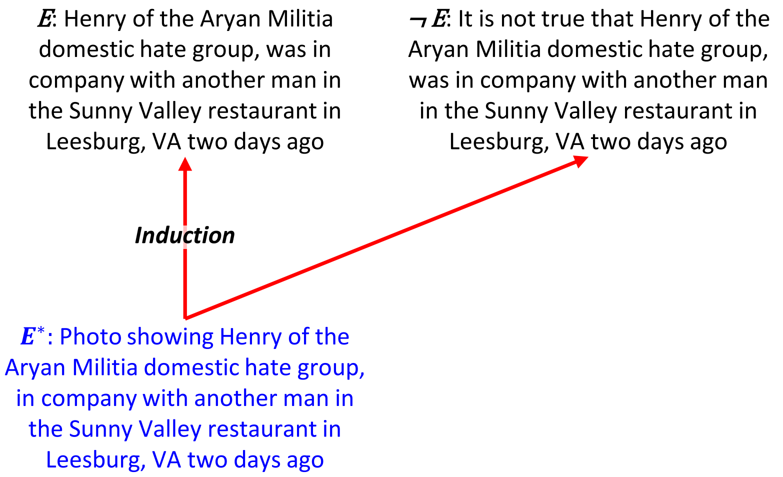
E: Henry of the Aryan Militia domestic hate group, was in company with another man in the Sunny Valley restaurant in Leesburg, VA two days ago
¬E: It is not true that Henry of the Aryan Militia domestic hate group, was in company with another man in the Sunny Valley restaurant in Leesburg, VA two days ago
Induction
E*: Photo showing Henry of the Aryan Militia domestic hate group, in company with another man in the Sunny Valley restaurant in Leesburg, VA two days ago

Figure 19. Credibility foundation of the argument.

2.3.3 Undercoded Abduction

Now, Mavis considers the matter that aroused her curiosity: Who was this other man in company with Henry, at the Sunny Valley restaurant in Leesburg two days ago? Suppose the FBI photo shows a side view of this man's face. Mavis believes she has recently seen other photos of this same man seen in company with Henry. Checking her photo files, Mavis is quite astonished to see the strong possibility that the other man in the Leesburg photo is Rizwan Kayani, an American of Pakistani origin and Islamic faith, whose several recent trips to Pakistan have been of concern. Mavis knows that there is fairly credible evidence that Kayani was in contact with several groups associated with the Taliban. Kayani's present residence is in Falls Church, VA. Mavis's file also indicates that Kayani has used the alias "George Wilson." Both photos Mavis examined show that Kayani could easily pass for someone of Anglo-Saxon origin instead of a Middle Eastern origin.

But Mavis knows that there are other men who resemble the man in the FBI photo, including two men she knows who live in the Leesburg, VA area and who have had suspected associations with Henry in the past. One is Sgt. Bill T. [US Army], and the other is Ralph W., a car salesman in Leesburg, VA. So, Mavis considers the hypotheses in Figure 20 about the other man with Henry. Mavis tests two of these hypothesis A_2 and A_3, and

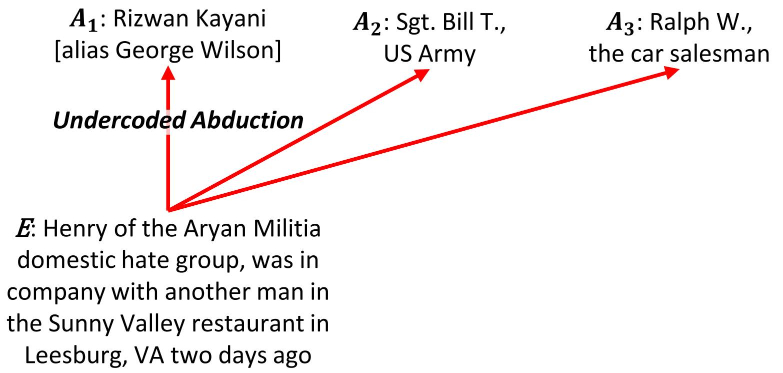
A1: Rizwan Kayani [alias George Wilson]
A2: Sgt. Bill T., US Army
A3: Ralph W., the car salesman
Undercoded Abduction
E: Henry of the Aryan Militia domestic hate group, was in company with another man in the Sunny Valley restaurant in Leesburg, VA two days ago

Figure 20. Undercoded Abduction.

rejects both of them. First, suppose Mavis learns that A_2 cannot be true, since she finds evidence that Sgt. Bill is currently in Iraq and has been there for six weeks. Second, she calls Ralph, the car salesman, who says he has not been in the Sunny Valley restaurant in Leesburg for five years since he hates the food and dislikes the owner. So, by process of elimination, Mavis settles on Rizwan Kayani as being the man in company with Henry.

Now, Mavis knows her inference regarding A_1 can certainly not be conclusive since she has only ruled out two other persons; think of how many other men there are who resemble Riswan Kayani. So, probably the best Mavis can say is that she has made a reasonable guess about who this other man was who was with Henry. We will show how her guess here is what Umberto Eco called undercoded abduction, the second grade of creativity (see Section 2.2.5). The links in Mavis's chain of reasoning are shown in Figure 20.

2.3.4 Overcoded Abduction

The next question Mavis asked is: How does it happen that Henry, a leader in Aryan Militia, was in company with a person having Islamic origins and possible jihadist intentions. The Aryan Militia organization has often announced its hatred of all persons of Islamic faith and says that these persons should be immediately expelled from America; those who remain should be exterminated. One explanation occurring to Marsha, is that Henry really does not know who he was talking to at the Sunny Valley restaurant in Leesburg, VA. Using his alias "George Wilson", Kayani may have made contact with Henry announcing views in strong sympathy with the views of Aryan Militia. This sets up the next link on Mavis's chain of reasoning. Here come some more guesses on Mavis's part. She considers the three hypotheses at the top of Figure 21. In testing these three hypotheses, Mavis first has evidence that Rizwan Kayani has used the alias "George Wilson" successfully in the past on several occasions. Why should Kayani change this alias on the present occasion, especially if he believes Henry will check up on his identity? So, this makes B_1 more likely than B_2. Now, as far as B_3 is concerned, Mavis can think of no reason why Kayani would reveal his true identity to a person, like Henry, already well-known to have a strong hatred for persons of Islamic faith, particularly those who are jihadists. So, Mavis abductively guesses B_1 over B_2 and B_3. But this may be an instance of a mixture of abductive and inductive inference, of the sort Charles S. Peirce recognized, and that Umberto Eco called overcoded

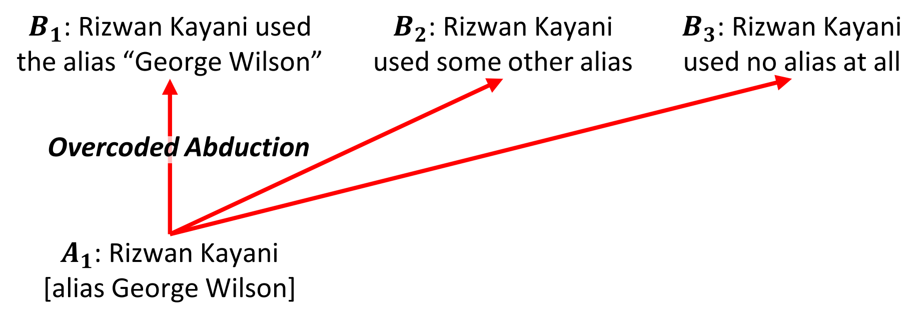
B_1: Rizwan Kayani used the alias "George Wilson"
B_2: Rizwan Kayani used some other alias
B_3: Rizwan Kayani used no alias at all
Overcoded Abduction
A_1: Rizwan Kayani [alias George Wilson]

Figure 21. Overcoded Abduction.

abduction, Eco's least creative species of abduction. The alias "George Wilson" worked in the past, Kayani guesses that it will work in the present.

2.3.5 Creative Abduction

At this stage Mavis tells us that she believes that Kayani made contact with Henry, using the alias "George Wilson". Here Mavis would naturally ask why should Kayani [Wilson] make this connection with Henry, a known leader of the hate group Aryan Militia. Here Mavis's curiosity becomes very focused. The most plausible reason occurring to her is that Kayani [Wilson] wishes to exploit the Aryan Militia organization in some way. Then, her next related question is: What could jihadists, such as Kayani [Wilson], get from Aryan Militia in such exploitation? Mavis is well aware of the fact that domestic hate groups such as Aryan Militia have acquired large stocks of all sorts of weapons. She also knows that some of these weapons have been stolen from various military bases here in America; some have been purchased from military personnel who have themselves stolen weapons of various sorts. And, some of these weapons may have been purchased or obtained from other sources such as the producers of weapons by organizations both here and abroad. So, a plausible yet surprising inference here is that Kayani [Wilson] wishes to acquire weapons from domestic hate groups that could be used here in America by jihadists planning acts of terror. The four possibilities Mavis considers are shown in Figure 22.

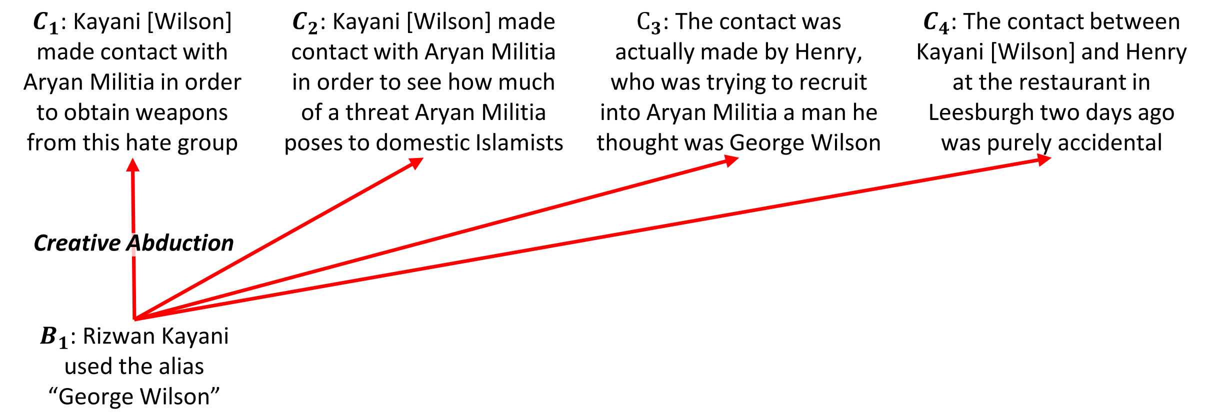
C_1: Kayani [Wilson] made contact with Aryan Militia in order to obtain weapons from this hate group
C_2: Kayani [Wilson] made contact with Aryan Militia in order to see how much of a threat Aryan Militia poses to domestic Islamists
C_3: The contact was actually made by Henry, who was trying to recruit into Aryan Militia a man he thought was George Wilson
C_4: The contact between Kayani [Wilson] and Henry at the restaurant in Leesburgh two days ago was purely accidental
Creative Abduction
B_1: Rizwan Kayani used the alias "George Wilson"

Figure 22. Creative Abduction.

Question for the reader: Could you imagine other possibilities?

Mavis chooses C_1 after attempting to assess the probability of all four of these hypotheses. First, she comes close to ruling out C_2, reasoning that domestic Islamists are already quite aware of the public statements of the Aryan Militia group for their hatred of persons of Islamic faith. However, she does allow that Kayani [Wilson] might wish to learn about any specific act of terror Aryan Militia, and related hate groups, are planning to inflict upon local domestic Islamists. As far as C_3 is concerned, Mavis has no present evidence regarding who arranged the contact between Kayani [Wilson] and Henry. However, Mavis does know that Aryan Militia is very selective of persons it tries to recruit. Regarding C_4, Mavis is able to rule this out by contacting the owner of the Sunny Valley restaurant who said that Henry, a regular customer, came and left with a stranger two days ago.

Now we are going to argue that Mavis's guess that C_1 occurred is an example of what Eco called a creative abduction, the highest grade of creativity. It is true that there is nothing creative about the thought that domestic jihadists need weapons to use against their fellow Americans. But what is creative is the idea that the jihadist group to which Kayani [Wilson] belongs should try to obtain weapons from a domestic hate group like Aryan Militia, that especially hates all Islamists. To our [and Mavis's] knowledge, this is an idea that has no precedent.

2.3.6 Undercoded Abduction

Mavis then considers the question: "What kinds of weapons would Kayani [Wilson] wish to acquire on behalf of jihadist groups he may be associated with"?

If you, the reader, were Mavis, what possible weapons would you consider?

Here are some things that Mavis does consider. First, jihadists are on record of using weapons that can kill many people to promote great terror, and they have done so in many places including America. Second, jihadists are relishing the success they had on 9/11/2001 using hijacked airliners as explosive devices in their attacks on the World Trade Center and the Pentagon. Third, jihadists have attempted, on several occasions, to bring down other domestic airliners, using explosive devices brought aboard these aircrafts. Fourth, these attempts have all failed. Fifth, the jihadists know that the increased security procedures in place at American airports and other airports around the globe have made it difficult to bring explosive devices aboard domestic airliners.

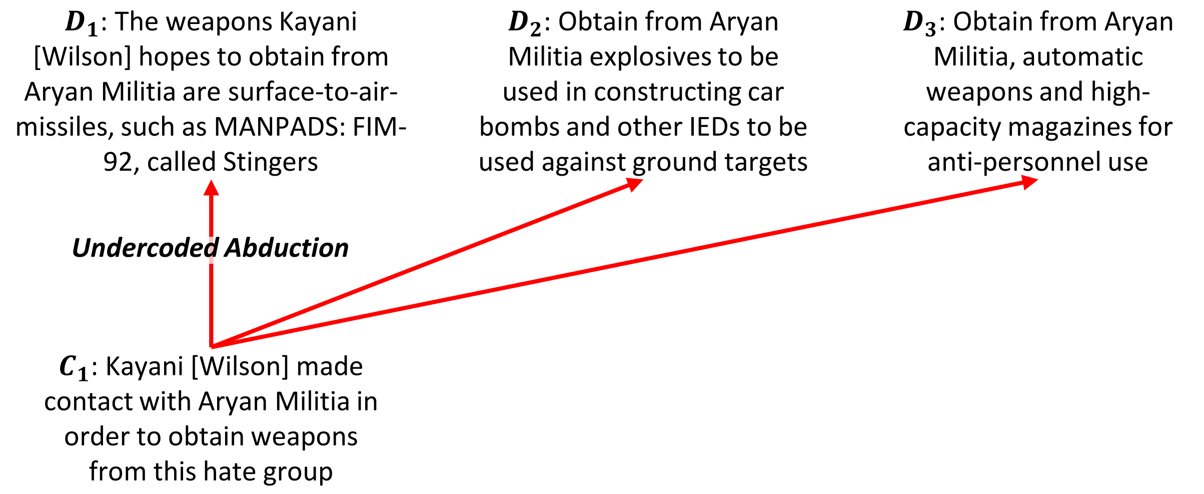
$\boldsymbol{D_1}$: The weapons Kayani [Wilson] hopes to obtain from Aryan Militia are surface-to-air-missiles, such as MANPADS: FIM-92, called Stingers
$\boldsymbol{D_2}$: Obtain from Aryan Militia explosives to be used in constructing car bombs and other IEDs to be used against ground targets
$\boldsymbol{D_3}$: Obtain from Aryan Militia, automatic weapons and high-capacity magazines for anti-personnel use
Undercoded Abduction
$\boldsymbol{C_1}$: Kayani [Wilson] made contact with Aryan Militia in order to obtain weapons from this hate group

Figure 23. Undercoded Abduction.

So, Mavis reasons that domestic airliners could be brought down without bringing any explosives aboard aircraft but by using surface-to-air missiles when airliners are taking off or landing. She also knows that jhadists have had experience, and considerable past success, in using MANPADS [Man-Portable Air Defense Systems], some of which were provided by America for use against Soviet invaders of Afghanistan (1979 – 1989). The version of these missiles was the FIM-92, called the Stinger. But Mavis also knows that jihadists have had great experience in using other kinds of weapons. Mavis chooses D_1 and so her reasoning chain is the one shown in Figure 24. All of these three possibilities would result in great loss of American lives, which is one of the jihadist's objectives. However, Mavis guesses that bringing down domestic airliners would be the highest priority terror-inducing activity in which the jihadists could be engaged. Mavis must know somehow that Aryan Militia has some MANPADs and would be willing to sell some of them. Mavis's guess of D_1 here is an example of Eco's undercoded abduction.

2.3.7 Deduction

Mavis has now told us so far about a sequence of four abductive inferences she has made based on her inductive inference concerning the authenticity of the photo evidence she has received from the FBI. She then tells us about two deductive inferences she found it necessary to make before she reached her final conclusion that was based on a fifth abductive inference we will mention in a minute. Here's what Mavis has told us, based on answers she has obtained to further questions she asked.

Mavis has wondered whether the Aryan Militia hate group has any MANPADS and whether they would be willing to sell any of them. Mavis says she obtained some valuable information from an FBI source who has been an informer within the Aryan Militia hate group. This informant first told Mavis that the Aryan Militia group does have about a dozen Stinger missiles that they bought from a person who stole them from a military supply depot in Maryland. The informant then said that the Aryan Militia group was in serious financial trouble and facing bankruptcy as a result of losing a law suit in which Henry was named as the main defendant. The informant then said further that Henry said he had been offered $50,000 each for Stinger missiles from a stranger he had met recently. Also, Henry said that this was an attractive offer since Aryan Militia had only paid $10,000 each for the Stinger missiles they purchased from the man who had stolen them. On the basis of this information, Mavis deduces the following event F: Henry of Aryan Militia sells around 5 stolen Stingers to the man he believes is George Wilson. Then, based on event F, Mavis further deduces event G: Kayani [Wilson] distributes the Stinger missiles he purchased from Henry of Aryan Militia to members of a local jihadist group of sleepers.

2.3.8 Undercoded Abduction

Mavis is now ready to generate and select the major hypothesis in her efforts. She has to guess which targets will be hit by the jihadists who will have the Stingers in their possession. The top of Figure 24 shows the possibilities she considers that face the jihadist Kayani [Wilson].

Now, all of these hypotheses would involve major losses of American lives. So Mavis has to consider what Kayani [Wilson] would consider in making his choice. Widely scattered targets would certainly enhance the terror generated by these Stinger attacks. American people will say that there is no safe place to avoid domestic airliner attacks. This makes H_3 and H_4 attractive. However, there is a risk of detection involved in transporting these missiles to widely scattered places. In addition, Mavis thinks that Kayani [Wilson] would also consider having to rely upon "sleeper" jihadists he may not know to be competent and reliable, even if he could find sleepers at these widely scattered targets. Mavis assumes that Kayani [Wilson] knows sleepers in the Washington-Baltimore area he can trust who will be competent and reliable; this makes H_1 and H_2 especially attractive since the Stingers would not have to be transported over great distances. Mavis has Kayani [Wilson] choosing H_1 over H_2 because BWI might be far enough away from Washington, since Baltimore could be considered a different location. Mavis guesses H_1, which we will say is an instance of undercoded abduction.

We have now shown you an extended example of hypothesis generation during the process of discovery and how this process required different species or creativity levels of abductive reasoning. We argued that discovery rarely, if ever, rests upon a single glorious episode of abductive reasoning. Instead, discovery involves mixtures of the three forms of abduction Eco

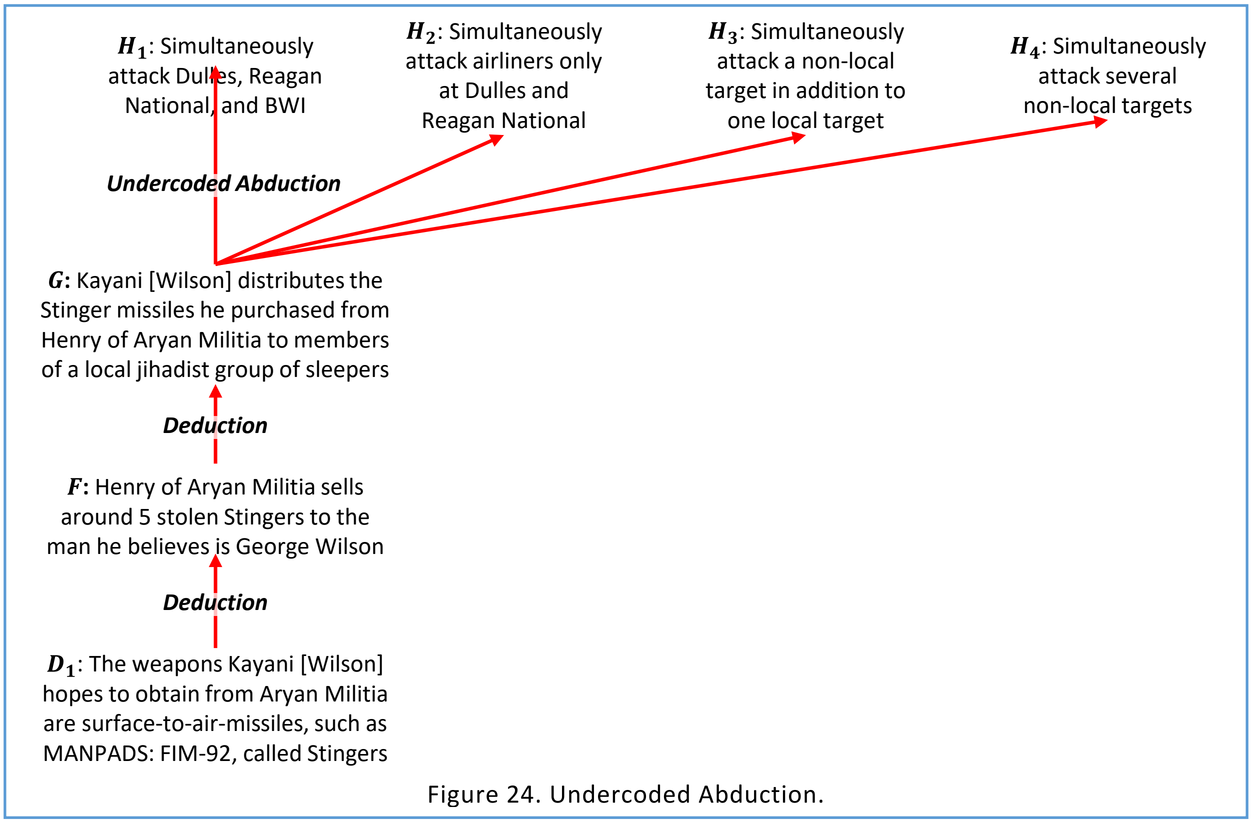
H_1: Simultaneously attack Dulles, Reagan National, and BWI
H_2: Simultaneously attack airliners only at Dulles and Reagan National
H_3: Simultaneously attack a non-local target in addition to one local target
H_4: Simultaneously attack several non-local targets
Undercoded Abduction
G: Kayani [Wilson] distributes the Stinger missiles he purchased from Henry of Aryan Militia to members of a local jihadist group of sleepers
Deduction
F: Henry of Aryan Militia sells around 5 stolen Stingers to the man he believes is George Wilson
Deduction
D_1: The weapons Kayani [Wilson] hopes to obtain from Aryan Militia are surface-to-air-missiles, such as MANPADS: FIM-92, called Stingers

Figure 24. Undercoded Abduction.

named. Further, discovery involves mixtures of different species of abduction taken together with instances of inductive and deductive reasoning. We developed, a stage at a time, the chain of reasoning analyst Mavis generated. We did this to show the evidential and inferential problems she encountered at each stage of the bottom-up reasoning she employed. As you see, Mavis generated several levels of interim hypotheses, before she could consider her major hypothesis in the situation she faced in trying to make sense out of an item of evidence she obtained. Our argument is that the combination of different ideas in our example allows us to capture more of the true complexity of discovery processes. Mavis's conclusion from a single photograph is certainly surprising. She would certainly not report this conclusion to colleagues and superiors without considering how she would defend her conclusion. Our example shows how an analyst like Mavis could defend an imaginative conclusion like the one she reached.

2.3.9 Conclusion

We all wish to have concise definitions of the terms used in any form of analysis. Many attempts have been made to supply labels to the mental processes according to which new ideas and new evidence are generated during the process of intelligence analysis and similar discovery-related activities in other disciplines. As we have noted, the term "abduction" has now been applied with great frequency to this generation process and has been defined in so many different ways that it begins to resemble a "wild card". Most current definitions of abductive reasoning such as "acts of insight" or “inference to the best explanation" do not seem to capture the true complexity of this vital reasoning activity when it is examined in contexts as rich as intelligence analysis. During intelligence analysis, there are many observations that need to be explained, and so we have the necessity of generating multiple hypotheses to explain these different events. The thoughts represented by these hypotheses, as well as the evidence upon which they are based, need to be marshaled or colligated in various ways. In the emerging "science of complexity" (Mitchell Waldrop, 1992), it has been said that the human brain is the very “cathedral of complexity" in the known universe: Pursuing this metaphor, one of the most interesting and complex *services* taking place in this "cathedral" concerns the generation or creation of new ideas and new evidential tests of them. Attempts to understand more of the complexity of these services require that we attend such services in many cathedrals. These services allow us a glimpse of the true complexity of the mental activities in investigation or discovery, upon which so much depends in our lives.

2.4 Marshaling “Magnets” or Attractors

Apart from convenience, it probably does not matter very much how we organize the clothes in our closet, the food in our pantry, or the books on our shelves. But in intelligence analysis and other inference tasks it does matter <u>very much</u> how we organize our <u>thoughts and our evidence</u>. How well we organize or marshal our ideas and evidence helps determine what new evidence and hypotheses we will generate and what conclusions we will draw. Different ways of organizing thought and evidence may lead us to:

- Ask different questions of and about our evidence.
- Discover different evidence and hypotheses.
- Draw different conclusions.

The processes of discovering evidence and hypotheses, and then using evidence as a basis for drawing conclusions, involves many different mental tasks. It is clear that there is no single way of organizing thoughts and evidence that will meet the demands of all of these tasks. We now examine various ways in which evidence you gather might be organized in different ways, each of which serves a useful purpose. Figure 25 is a simple picture of what different evidence marshaling strategies should help you accomplish; it involves the metaphor of a "magnet" or an attractor.

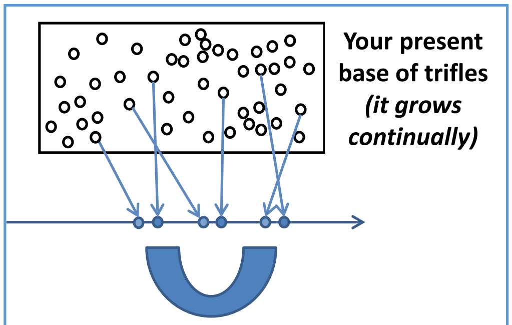

Figure 25. Marshaling magnet attracting interesting combinations of "trifles."

What we should like to have are conceptual "magnets" that could attract interesting and useful combinations of "trifles", as Sherlock Holmes called them. In Figure 25 the magnet has attracted several trifles that, together, may allow us to generate a new possibility or hypothesis, ask a new and important question, or generate some new potential evidence. It happens that different ways in which we marshal our thoughts and evidence can in fact serve like the "magnet" shown in this figure. On occasion, you may be able to generate a new idea, possibility, or hypothesis from a single trifle. More often, however, new ideas, new questions, and new possible evidence will only emerge in your mind from combinations of two or more trifles. Remember, it makes no sense to try to examine every possible combination of trifles or details you collect. The purpose of the evidence marshaling magnets we will mention is to help you decide which combinations of trifles would be valuable for you to identify and evaluate. Different ways of organizing your thoughts and your evidence may be heuristically valuable in suggesting questions you might ask of your data.

> *A **heuristic** is simply a rule of thumb that aids you in any discovery, inference, learning, or decision problem.*

You need such heuristics as aids in deciding which combinations of data you might most profitably examine. Discussed below are some specific marshaling strategies, or "magnets" you might consider. The sequence in which you employ these magnets depends upon the nature of the analysis tasks you will actually encounter.

In an intelligence analysis you may begin to accumulate trifles or details at a very rapid rate. Some of these details will be provided for you as a matter of course (e.g., your daily message traffic). Other details you will obtain in response to questions you ask. Some of the details will be tangible in nature; others will be items of testimony received from human sources.

Which details should you keep and which should ignore? If you could answer this question, your task would be much simpler. However, unless you are clairvoyant you cannot know for sure, at least in the early stages of an analysis, which details will become important relevant evidence and which will not. This task may get easier as your analysis proceeds and you begin to form a collection of plausible hypotheses on which certain details can become relevant evidence. Stated in other words, as your analysis proceeds, you may be able to ask better questions and obtain relevant details more efficiently.

Before we consider some specific marshaling operations or “magnets”, it seems wise to acknowledge that analysts will already have some schemes for organizing the information they acquire. But these schemes may only exist for the purpose of archiving their information in orderly ways. Such archiving, by itself, has little heuristic value but it may facilitate the operation of marshaling operations that do have heuristic value in suggesting new hypotheses and new lines of inquiry and evidence.

2.4.1 Credibility Magnet

One rather obvious way to archive trifles or details is in terms of the sources from which you received them. For example, we could organize evidence by "INTs": HUMINT, SIGINT, IMINT, MASINT, etc. In some cases, particularly regarding HUMINT, security matters will arise and you may not always be able to precisely identify a human source; you may only have an alias of some sort (such as a code word). It may also be useful to keep separate tracks of tangible evidence and testimonial evidence. If you organize your data in terms of their sources, there are two kinds of information you need to record, if you have it. The first, of course, are the trifles or details you receive from a source (whatever it is). In other words, you record what the source tells you. The second are details you have about this source. Details about a source become important as you begin to assess the credibility of what the source has given you. Credibility assessment will be discussed in detail in Chapter 4. Suppose a source has provided you with some tangible evidence. You should record whatever information you have about the authenticity and accuracy of this tangible item. If the detail exists in the form of a testimonial assertion from some human source, you should record whatever information you have regarding this source's competence, veracity, objectivity, and observational sensitivity. This forms the basis for the first marshaling magnet.

Keeping careful accounts of the trifles or details we have about our intelligence sources will provide a basis for our decision concerning whether or not to believe what a source is reporting to us. Suppose a source is discovered to be faulty in some way. If you have marshaled together all the trifles you have obtained about this source, you may begin to question the extent to which the items you have received from this source might not be credible. This can often pay huge inferential dividends. For example, suppose we now believe that a source S of HUMINT was lying to us in his report that event E occurred. We ask Why did S choose this lie about event E in preference to other lies he might have told us? Answers to this question may suggest very interesting new possibilities. Second, if we have at hand (i.e., marshaled) all the other events source S has reported, we might now begin to question the credibility of other reports S has provided. What we know about a source of evidence (testimonial or tangible) is often at least as interesting as what this source tells us. We can characterize this Credibility Magnet in the following way:

> ***Credibility Magnet*** *is a magnet that attracts trifles we have concerning the credibility of our sources of intelligence information.*

Now the issue becomes: What questions should we be asking about our sources? Section 4 discusses in details these questions that depend on the type of evidence. Section 4.5 discusses an often-overlooked credibility assessment problem. This problem also arises because of the many things that are done to information before it reaches the attention of analysts who will attempt to draw conclusions from this information. It is common to refer to this matter as involving the chain of custody through with information has passed before an intelligence analyst receives it. Many things might have been done to this information including translations, interpretations, editing, processing, and summarizing. A number of persons or automated processes might have been involved. The issue is: How authentic is the information received by an analyst? To what degree is this information an accurate and complete account of what an original source has reported? We have recorded our thoughts on these credibility matters involving chains of custody of intelligence information in (Schum et al., 2009).

Table 3 shows some trifles we have received from several sources and examples of the credibility information we might have about them.

Table 3. Illustration of the application of the credibility magnet.

Trifles from Sources	Trifles about Sources
An extract from a HUMINT report from Source A.	We have not been able to communicate with Source A since he gave us the information in his HUMINT report.
An extract from a copy of a document said to have been obtained by Source B from the files of a potential adversary.	An expert's assessment that there is a 30% chance that the document we received from Source B is a fake.
An extract from a table compiled by Source C showing the range of a certain missile.	Information that Source C has made errors in the past in calculating the range of missiles.
A fragment of the door panel of a car said to contain the explosive device that destroyed an overseas embassy.	Information that there were other cars parked in the vicinity of the explosion that were also destroyed.

2.4.2 Chronology Magnet

Many of the trifles or data we receive are "time-stamped" in terms of the time at which events in the data are alleged to have occurred. We are just as interested in when events have occurred as we are in their occurrence. The timing of events can be a most valuable heuristic source of possible hypotheses and new lines of inquiry and evidence. There is an inference here; sometimes we are misled about the time some reported event has occurred. An event chronology is simply an ordering of events according to the times at which we believe them to have occurred (new evidence might cause us to change our minds about this temporal ordering). Having some idea about the order in which events might have occurred also gives us some basis for establishing causal patterns that may be very important in reaching any final conclusions. Such chronologies can also be very useful in efforts to predict events that may happen in the future. So, our Chronology Magnet has this purpose:

> ***Chronology Magnet*** *is a magnet that attracts inferred times at which reported events have occurred and allows inferences about the temporal ordering of these events.*

There is nothing new about event chronologies; you may already be naturally constructing them. But there are some difficulties here; event chronologies can get very messy if we have many events to record. One strategy is to form parallel event chronologies, each of which records events belonging to a certain class.

Figure 26 shows an example involving an intelligence analysis in which there may be many "actors." An actor may be either a person or an entire group of persons such as a terrorist organization or an army battalion.

These parallel chronologies record the times at which we believe the events associated with each actor have occurred. Evidence about the events associated with each actor may have come from the actor or from another source. It is not difficult to observe how event chronologies serve as important "magnets" that can stimulate the process of inquiry.

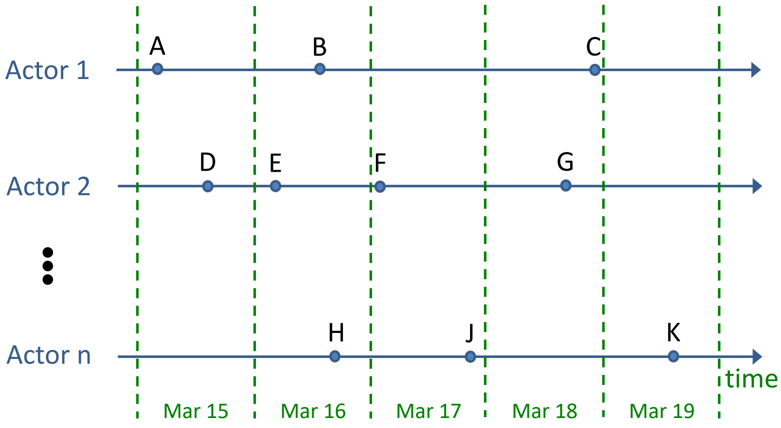

Figure 26. Parallel event chronologies.

Consider the three events associated with Actor 1: A, B, and C. We might ask, "If Actor 1 did A, and then sometime later did B, what was he doing in the interval?" We might also ask such questions as, “If Actor 2 did D and E between the time Actor 1 did A and B, did Actor 2's actions have any effect on what Actor 1 did during this interval?” By such means we are stimulated to examine particular combinations of trifles and to try to discover them if they are not now at hand. We add here that Indications and Warnings (I&W) assignments involve use of evidence chronologies. Such marshaling allows us to provide timely warnings of future events involving matters of national security.

H_1		H_2		H_3	
favoring	disfavoring	favoring	disfavoring	favoring	disfavoring
E_6^*	E_{10}^*	E_1^*	E_9^*	E_1^*	E_2^*
E_8^*		E_3^*			E_4^*
		E_4^*			
		E_5^*			
		E_7^*			

Figure 27. Evidence marshaling by hypotheses.

Example 3. Here are three events not in any order:

Event A = Person P drank three double martinis.
Event B = Person P left the base parking lot.
Event C = Person P was involved in a car accident.

Consider each possible sequence and see what story it tells. If you were person P, you would rather have the following sequence of events B → C → A than either of the two sequences B → A → C or A → B → C. In these two sequences P faces a charge of driving while intoxicated which he does not face in the first sequence B → C → A.

2.4.3 Question Magnet

In the search for hypotheses or possibilities that will account for all of the evidence you believe to be relevant, various specific issues and questions will arise. Here is an example.

Example 4. You are trying to determine whether or not a certain terrorist organization is planning an act of destruction at a certain location in the near future. You now have some evidence that this group is obtaining materials for the construction of explosive devices from a particular source. A variety of questions come to mind, such as:

What type of explosives might be used?
How are they to be detonated?
Where are they to be detonated?
When are they to be detonated?

Questions you ask serve as most important "magnets" for attracting combinations of data from your records. We can define this Question Magnet as follows:

> ***Question Magnet*** *is a magnet that attracts trifles representing possible answers to any question that comes to mind as an intelligence analysis proceeds.*

Remember that the process of inquiry is a most vital ingredient of productively imaginative intelligence analysis. Each question you ask not only serves as a device for attracting existing trifles, but also serves as a device for generating new trifles you do not presently have. Keeping track of all the trifles attracted by a certain question can be most valuable in generating new hypotheses and new potential evidence.

2.4.4 Hypothesis Magnet

Suppose we now have a large assortment of trifles and we identify ten of them as being evidence we regard as relevant to three hypotheses we are now considering. One obvious method of organizing these ten items of evidence is in terms of the hypothesis each item seems to favor (see Figure 27). This marshaling method is closely tied to what is frequently said to be the method of "competing hypotheses" in which we attempt to judge each hypothesis on its merits.

In the situation from Figure 27 evidence items 6, 8, and 10 are relevant to hypothesis H_1; items 1, 2, and 4 are relevant to hypothesis H_2; and the rest are relevant to hypothesis H_3. Notice that, under each hypothesis, we group the evidence into favoring and disfavoring. This form of marshaling is often useful since it allows us to observe which hypothesis seems to have the most evidence in its favor. We can define this Hypothesis Magnet as follows:

> ***Hypotheses Magnet*** *is a magnet that uses generated hypotheses to attract information items that could become relevant evidence in their favor or against.*

Marshaling evidence by hypotheses has another useful feature that concerns the completeness or sufficiency of the evidence we have. Suppose someone says, “Hypothesis H_2 (in Figure 27) has the most evidence favoring it, so we should conclude that this hypothesis is the one we ought to advocate.” But, another person very wisely says, “Before we decide on Hypothesis H_2, or any other, we ought to ask how many questions there are that remain unanswered by all the evidence we have.” This question of evidential completeness or sufficiency is so often overlooked and can result in some dramatic inferential miscarriages. Marshaling by hypotheses, as well as marshaling by argument, to be presented next, allow us to survey what we have and we don’t have in the way of evidence on every hypothesis we are considering.

In some instances, evidence will say something about one hypothesis but say nothing at all about other hypotheses. For example, evidence E_{10}^* in Figure 27 goes against hypothesis H_2, but does not favor or disfavor hypotheses H_1 or H_3. Similarly, E_6^* favors H_2 but it is not relevant to the other hypotheses. As an example, consider two organizations, A and B, which may now be distributing narcotics in a certain city. We have evidence that A has contacts with foreign suppliers of narcotics. This tends to raise our suspicions about A but it tells us nothing about B.

In other instances, evidence may favor more than one hypothesis. For example, E_1^* in Figure 27 favors both H_2 and H_3. As a concrete example, suppose that a person X belongs to two terrorist groups. Evidence that X was at the scene of a terrorist incident might favor hypotheses concerning either or both of these groups being responsible.

We may also encounter evidence we regard as relevant but is not well explained by any hypotheses we are currently considering. One way to describe this evidence as unexplainable by any existing hypothesis being considered is to say that it is an anomaly. *We might consider disregarding this anomalous evidence completely; but to do so would be the height of foolishness.* Perhaps this anomaly means that there is a hypothesis we have not yet considered that could explain this anomaly and, perhaps, better explain all the other evidence we have. This is one frequently observed way in which new hypotheses and new evidence are generated. Pondering upon an anomaly we are led to generate a new hypothesis and new evidential tests of all of our hypotheses.

2.4.5 Argument Magnet

Here is an important and useful refinement of marshaling by hypotheses. Consider again the situation in Figure 27 showing the generated hypotheses and the collection of relevant evidence bearing upon each of them. Suppose further that you are preparing to argue that H_2 is the most likely hypothesis from among those you have considered. In your defense of hypothesis H_2, you expect to be required to produce specific evidence-based arguments about why you favor hypothesis H_2 over the hypotheses H_1 and H_3. You think about this problem carefully and decide that your evidence suggests three major lines of argument on hypothesis H_2; these major lines of argument (or sub-hypotheses) are H_{A1}, H_{A2}, and H_{A3}. A natural form of evidence marshaling would be to organize your existing evidence under each of these sub-hypotheses, as shown in Figure 28. In this case the Marshaling Magnet has the following function: Argument Magnet is a magnet that attracts trifles that will form relevant evidence on major arguments for some hypothesis being entertained.

This form of marshaling is both useful and necessary in your efforts to construct defensible and persuasive arguments as far as hypothesis H_2 is concerned. Such marshaling helps you to see what additional evidence you will need to construct stronger and more complete arguments in defense of hypothesis H_2. Remember the discussion from the previous section concerning the selection of hypothesis H_2 because it has the most evidence favoring it (see Figure 27)? The advice there was that before we decide on Hypothesis H_2, we ought to ask how many questions there are that remain unanswered by all the evidence we have. The argument in Figure 28 shows that we have no evidence bearing on sub-hypothesis A_3, and we should not reach any conclusion before we are also able to assess A_3.

As you construct your arguments from this array of marshaled evidence, you can obtain a better idea about how strong and sufficient your arguments favoring H_2 will be at this time. There are specific ways in which you can construct chains of reasoning from the evidence you have to each of these main lines of argument in your analysis. The very first step in the process is a careful marshaling of evidence under each of your main lines of argument. Finally, your marshaling efforts may allow you to see that your main lines of argument are not sufficient and that you may need additional lines of argument. Your hope is that your main lines of argument will be necessary and sufficient or, at least, sufficient to show that hypothesis H_2 is true.

2.4.6 Eliminative Magnet

In many situations we use evidence not to support hypotheses but to try to eliminate them on the basis of a variety of different questions we ask and hope to have answered by evidence. This is precisely the method Sherlock Holmes says he used in solving his cases. The hypothesis that survives our best attempts to eliminate any of our hypotheses is the one we should take seriously. Now, one of the most embarrassing things that can happen to an analyst is to eliminate some hypothesis that later proves to be true. Critics will say that this analyst "snatched defeat from the jaws of victory." So, when we say we are eliminating some hypothesis (we should never do so completely), we ought to make sure that we have exhausted all reasons for keeping this hypothesis alive. To do this is to protect ourselves from the "hindsight critic" who will say, post mortem, you had the truth in your grasp and you let it slip away. By implication, this critic is saying that he would have kept it alive. He can make himself look good of course, since he now knows what did happen. So what an analyst should do is to keep track of all of the questions he/she asked and all of the evidence obtained, in response to these questions, which argued against the hypothesis you have chosen to eliminate. Unless we do this, someone can always later say: Why did you reject this hypothesis that now seems so obvious?

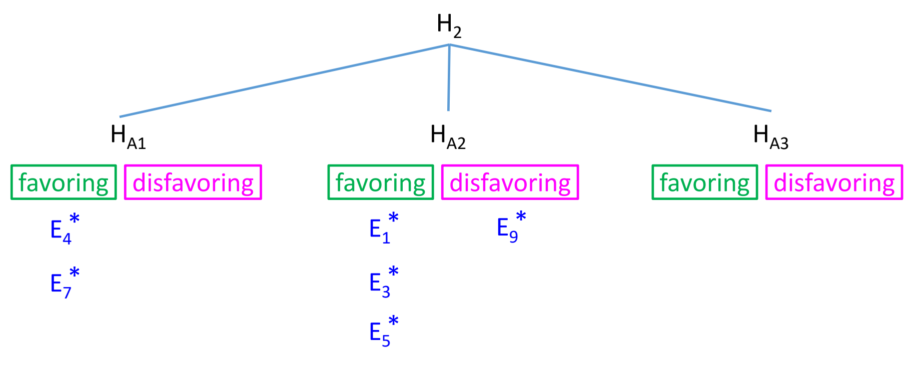

Figure 28. Evidence marshaling by arguments.

Suppose you have decided to eliminate hypothesis H_3 from consideration (see Figure 28). You will want to have evidence showing why you have chosen not to keep H_3 alive. So the Eliminative

Marshaling Magnet has the following interpretation:

> ***Eliminative Magnet*** *is a magnet that attracts trifles representing evidence relevant in showing why some hypothesis can be safely eliminated.*

Suppose you have eliminated hypothesis H_3 early in your analysis because someone tells you that it is too improbable for anyone to believe. Though you don't have much evidence yet, you decide to forget about this hypothesis. You gather further evidence and draw a conclusion that hypothesis H_2 is true and you report your conclusion to higher authorities. Later, it turns out that H_3 was true after all and you are reprimanded. You are reminded that you dismissed hypothesis H_3 without attempting to gather other evidence that might have been in its favor. In short, your evidence was not complete enough for you to rule out H_3.

Example 5. As an example, consider that someone has been leaking classified information from your organization. Person Y is presently a suspect. You decide to eliminate Y as a possibility. What general factors should you consider before you decide to rule out Y as a possibility?

It seems you could only rule out Y if all these conditions are met:

(i) Y would never have had access to this information.
(ii) Y never attempted to obtain this information from someone who did have access to it.
(iii) Y had no reason to wish or need to obtain this information.

2.4.7 Scenario Magnet

The event chronologies mentioned above simply allow us to list interesting events in the order in which we believe they occurred. *One of the most heuristically-valuable exercises is to construct stories or scenarios about what we believe might have happened, might be happening, or might happen in the future.* An example of a scenario is that in Figure 29.

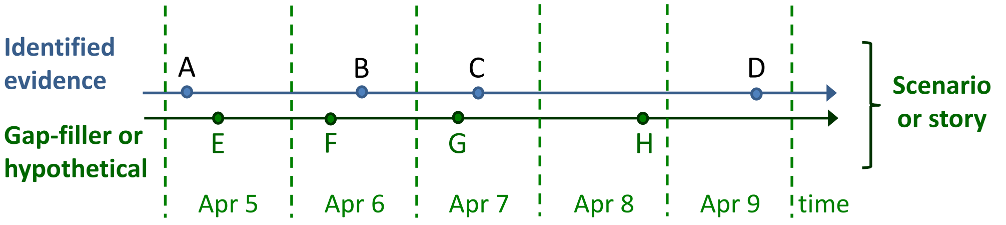

Figure 29. Evidence marshaling through the construction of scenarios or stories.

Like any stories, the ones we construct always consist of a mixture of evidence and fiction or fancy. It is the fanciful elements of stories we tell that are most valuable in generating new hypotheses and new evidence. Such stories also provide yet another heuristic for focusing attention on specific combinations of data we have on file. Figure 30 shows a simple picture of a story or scenario constructed for heuristic purposes.

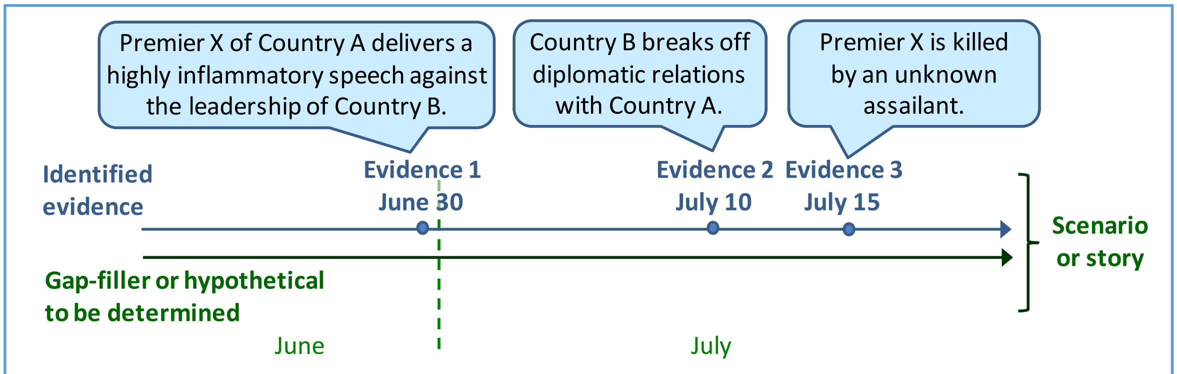
Premier X of Country A delivers a highly inflammatory speech against the leadership of Country B.
Country B breaks off diplomatic relations with Country A.
Premier X is killed by an unknown assailant.
Identified evidence
Evidence 1
June 30
Evidence 2
July 10
Evidence 3
July 15
Scenario or story
Gap-filler or hypothetical to be determined
June
July

Figure 30. Evidence about a sequence of events.

Suppose we have evidence that events A, B, C, and D have occurred at the times indicated in Figure 29. But to tell a coherent story or to construct a coherent scenario about what these events mean, we need to fill in the gaps with evidence we do not now have. Suppose, for example, part of our story involves saying that the occurrence of A led to the occurrence of B; and we have evidence that both of these events occurred. But, we then think that A, by itself, could not have given rise to B, unless events E and F had also happened. So, we fill in the gap between A and B by these two hypotheticals or "gap fillers" in order to tell this part of our story. We do the same thing at other points, such as filling in the B - C gap with G and filling in the C - D gap with H. Here is the payoff: Each new gap-filler we identify alerts us to examine our existing data base to see if such data exist. If they do not, then we are stimulated to try to discover them. Naturally, we may discover that these hypothesized events did not occur. If this happens on enough occasions we have to change our story. We may also be led to question the credibility of the evidence that suggested this story in the first place. The Story or Scenario Magnet has the following interpretation:

> ***Scenario Magnet*** *is a magnet that attracts a temporally-ordered sequence of trifles forming relevant evidence about events that will form the basis for a story or scenario about what has happened in some situation of interest.*

Example 6. From any collection of information arranged in chronological order, a virtual infinity of different stories might be told. The smaller the number of evidence items, the more possible stories there are. As an illustration, suppose we now have just the three items of evidence from Figure 30 whose temporal ordering we have reason to believe. What possible scenario does this sequence of events suggests?

The top part of Figure 31 shows a possible scenario: Premier X was killed by a member of a group in his own country that regarded his leadership as reckless. The group will argue that its action prevented a war between countries A and B. This Scenario 1 suggests the two gap fillers shown above the time line. These gap fillers open new lines of investigations, suggesting that we should look for evidence that:

- Members of a certain group in Country A were enraged by what they regarded as a reckless speech by X.
- Members of this group within A determine that X has to go.

But Scenario 1 is not the only scenario consistent with the available sequence of events. Another one is Scenario 2 from the bottom part of Figure 31: Hoping that their own interests would be better served by a leadership change in Country A, the leaders of Country B decide to take matters into their own hands.

Scenario 2 suggests other gap fillers and guide us to look for evidence that:

- The leadership in Country B reacts strongly against the leadership in Country A.
- The leadership in B initiates a plot against the life of Premier X in Country A.

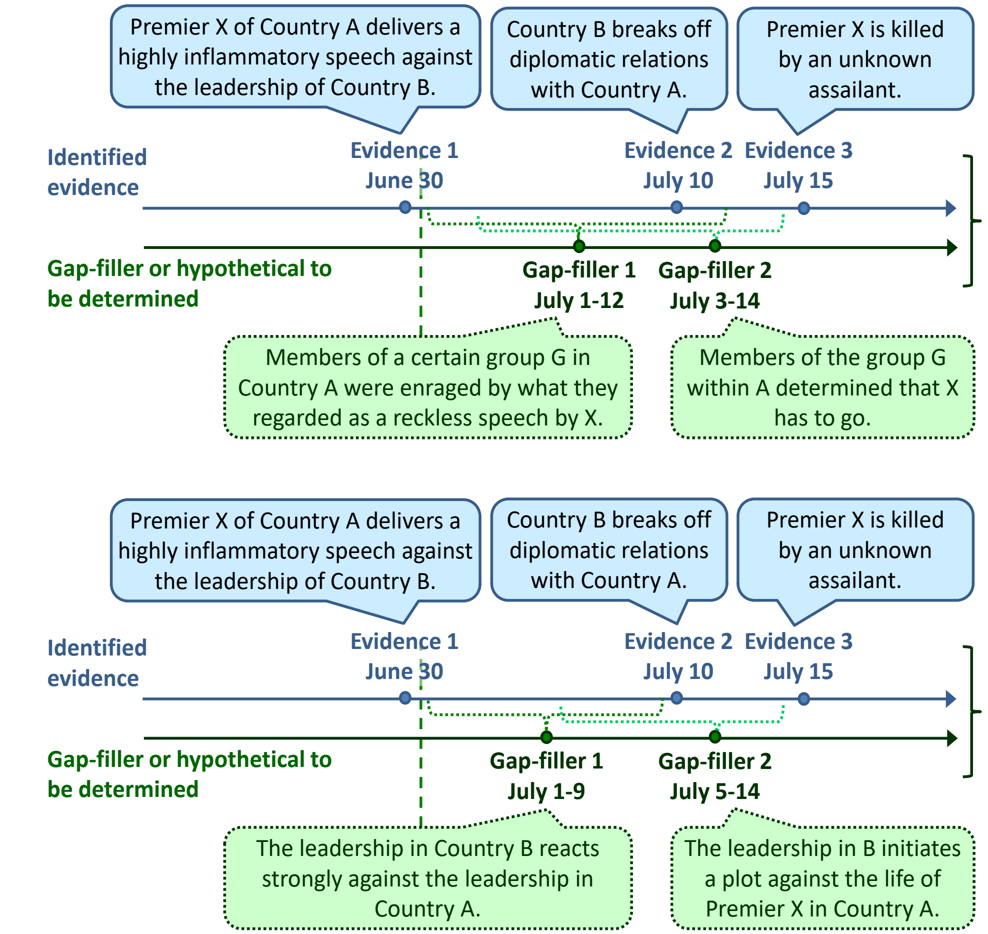

Premier X of Country A delivers a highly inflammatory speech against the leadership of Country B.
Country B breaks off diplomatic relations with Country A.
Premier X is killed by an unknown assailant.
Identified evidence
Evidence 1 June 30
Evidence 2 July 10
Evidence 3 July 15
Gap-filler or hypothetical to be determined
Gap-filler 1 July 1-12
Gap-filler 2 July 3-14
Members of a certain group G in Country A were enraged by what they regarded as a reckless speech by X.
Members of the group G within A determined that X has to go.
Scenario 1: Premier X was killed by a member of a group in his own country that regarded his leadership as reckless. The group will argue that its action prevented a war between countries A and B.
Premier X of Country A delivers a highly inflammatory speech against the leadership of Country B.
Country B breaks off diplomatic relations with Country A.
Premier X is killed by an unknown assailant.
Identified evidence
Evidence 1 June 30
Evidence 2 July 10
Evidence 3 July 15
Gap-filler or hypothetical to be determined
Gap-filler 1 July 1-9
Gap-filler 2 July 5-14
The leadership in Country B reacts strongly against the leadership in Country A.
The leadership in B initiates a plot against the life of Premier X in Country A.
Scenario 2: Hoping that their own interests would be better served by a leadership change in Country A, the leaders of Country B decide to take matters into their own hands.

Figure 31. Two different scenarios hypothesized from the events inFigure 29.

2.5 The Case of General Alpha: Facing Insurgency?

2.5.1 Context and Background

In a certain part of the world of strategic importance to the U.S., a country named Orange [population around 2.5 million] has, for five years, been in the throes of a dictatorship that has decidedly unfriendly inclinations toward the U.S. We presently have no diplomatic relations with Orange. In addition, this dictatorial government has adopted quite ruthless methods to maintain its hold on the populace of Orange. Further, the ruler in Orange, General Alpha, has made life unpleasant for a bordering country, Green, friendly to U.S. interests, by fomenting repeated border disputes. It seems clear that Gen. Alpha has every intention of extending his dominance to include this bordering country Green. Gen. Alpha's support has come mainly from the military forces in Orange and from individuals in the two major urban areas in Orange who have profited by cooperating with Gen. Alpha.

The majority of Orange's population lives in rural areas, many of which are quite remote from the two Orange urban centers, cities A and B. Opposition to Gen. Alpha among the general urban and rural population in Orange has been held in check due to very ruthless measures taken by Gen. Alpha's military against persons Alpha has called "undesirables." Many of these "undesirables" have either been publicly executed or now languish in military prisons. The military in Orange, numbering about 75,000, is controlled exclusively by Gen. Alpha and seems well equipped for ground fighting. In addition, Gen. Alpha has a small air force consisting mainly of propeller-driven aircraft and a few helicopters. Information about events in Orange is generally hard to obtain. International news organizations have been barred from entering Orange since Gen. Alpha wrested control of Orange about five years ago. Gen. Alpha's government controls the press, radio, and the single television station that is allowed to broadcast.

Quite recently, an opposition group, named the Blues, has come into existence and is led by two persons who have managed to avoid capture by Gen. Alpha's military. Person X was a high-ranking military leader in country Orange before its government was overthrown by Gen. Alpha. Person Y is a former professor of political science at the only university in Orange and who also held a high-ranking political position in Orange before the government was overthrown. Both X and Y have remained in hiding in a remote region of Orange and have, from time to time, been granted sanctuary in neighboring country Green. When they were last in Green, about a week ago, X and Y told our foreign intelligence operatives that the stage was set for an insurgent operation against Gen. Alpha and his ruthless regime in country Orange. Furthermore, they indicated that their insurgency operation would be launched sometime in the next month. Of course they asked for military and other assistance from the U.S.

2.5.2 Your Task

Suppose you are a member of a team of military capabilities analysts who have been assigned "The Case of General Alpha." Your Commander is concerned about this insurgency and has asked your Intelligence team to assess the probability that the insurgent group will overthrow General Alpha. As analysts, you are charged with reaching a conclusion about whether the Blue

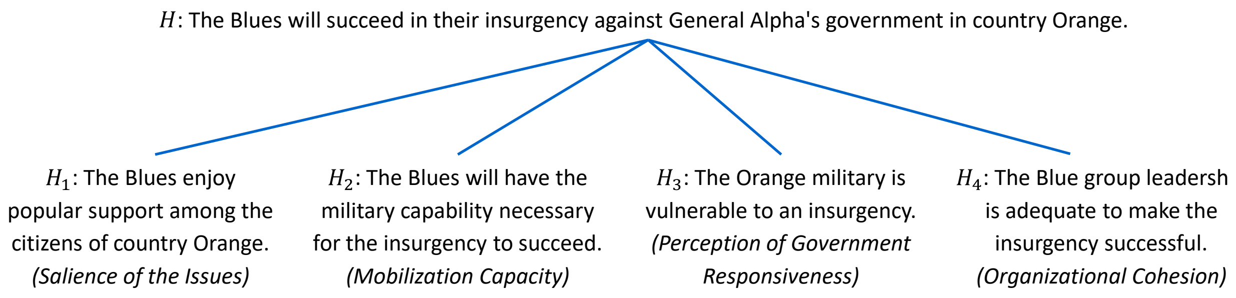
H: The Blues will succeed in their insurgency against General Alpha's government in country Orange.
H_1: The Blues enjoy popular support among the citizens of country Orange. (Salience of the Issues)
H_2: The Blues will have the military capability necessary for the insurgency to succeed. (Mobilization Capacity)
H_3: The Orange military is vulnerable to an insurgency. (Perception of Government Responsiveness)
H_4: The Blue group leadersh is adequate to make the insurgency successful. (Organizational Cohesion)

Figure 32. Major lines of arguments for the top hypothesis.

insurgency will succeed in forcing General Alpha and his associates from power in country Orange.

This case study is designed to give you some experience in marshaling or organizing a mass of evidence and then constructing arguments from this evidence that bear on a certain hypothesis regarding possible explanations for the evidence. We have simplified this case study considerably by supplying you with a major hypothesis to defend and four main lines of argument on this hypothesis. We know that supplying you with a hypothesis and main lines of argument on it is a luxury that you will rarely, if ever, have in actual intelligence analyses. However, we have made this case study simpler as an introduction to the more complex case studies to follow. The major hypothesis H is consistent with your commander's stated requirement, that the Blues will succeed in their insurgency against General Alpha's government. The alternative hypothesis is $\neg H$, where $\neg H_1$ means that the Blues will not succeed.

The next step is to imagine an argument for H. You reason that H would be true if the Blues have the military and leadership capabilities; that the Blues have the support of Orange citizens; and that the Orange military under General Alpha is vulnerable. You decide on four major lines of argument (sub-hypotheses) on hypothesis H, as shown in Figure 32.

2.5.3 The Evidence at Hand

Following is a list of 31 items of evidence that have been gathered so far. As you well know, evidence does not come to us already marshaled or organized. The following listing of evidence might simply be the order in which you received this evidence. Nor does evidence come to us with its relevance, credibility, and force credentials already established. You have to establish them.

Table 4. Collected evidence.

E1 SIGINT: [From one of our own SIGINT analysts]: Radio communications among clandestine groups operating in rural and urban areas of country Orange has increased over the past few days.

E2 SIGINT: [From the same SIGINT analyst as in E1 SIGINT]: Radio communications sites are scattered widely throughout the rural areas of country Orange.

E3 HUMINT: [From a military officer in country Green]: Country Green is now assembling weapons near its border for use by the Blue insurgent group in country Orange.

E4 IMINT: [As interpreted by one of our own image analysts from an overhead photo taken two days ago]: Weapons are being assembled in the area described in E3 HUMINT above.

E5 SIGINT: [Recording of Orange military radio transmission two days ago]: A car bomb was set off near an Orange military installation in city B; twelve Orange soldiers killed and three vehicles destroyed.

E6 IMINT: [As interpreted by one of our own image analysts from an overhead photo taken yesterday]: Observation of building destruction and disabled vehicles at the Orange military installation noted in E5 SIGINT above.

E7 HUMINT: [From an alleged Orange country refugee R1, now in sanctuary in country Green]: General Alpha's recently increased taxes levied on all agricultural and mineral products have forced many farm and mine owners out of business.

E8 HUMINT: [From the same source as E7 HUMINT]: Gen. Alpha's military forces have begun the takeover of farms and mines in rural areas of country Orange.

E9 HUMINT: [From an alleged Orange country refugee R_2, now in sanctuary in country Green]: Says that person X in the Blues, the former military commander in country Orange, has no support among his former soldiers now living in Orange.

E10 HUMINT: [From an alleged defector from Orange military, now in sanctuary in country Green]: The number of counterinsurgency exercises conducted by the Orange military has increased during the past month.

E11 IMINT: [As interpreted by one of our own image analysts from an overhead photo taken several weeks ago]: There are few developed roads in rural areas of country Orange.

E12 HUMINT: [From another alleged Orange country refugee R_3, now in sanctuary in country Green]: This source says that person X does enjoy widespread support among his former soldiers still living in Orange.

E13 HUMINT: [From R_3; the same source as in Evidence 12)]: R_3 says that R_2, who provided Evidence 9) is no refugee but a member of the Orange military forces.

E14 HUMINT: [From an interview in country Green with person Y]: Even the remotest areas of country Orange have clandestine news sources operated by the Blues.

E15 IMINT: [As interpreted by one of our own image analysts]: Many small Orange military forces are scattered throughout the rural areas of country Orange.

E16 OPSOURCE: By far the majority of persons living in both rural and urban areas belong to the religious group headed by leader L.

E17 SIGINT: [From one of our own COMINT interpreters]: Orange military C^3I capabilities are limited by the age and condition of their communications equipment.

E18 IMINT: [From an overhead taken two weeks ago and interpreted by one of our own image analysts]: Evidence of heavy weapons firing ranges in operation in a remote area of Orange.

E19 HUMINT: [From interview of person X while in sanctuary in country Green]: X says these heavy weapon firing ranges in the E18 IMINT) photo are operated by the Blue group.

E20 SIGINT: [Clandestine radio announcement recorded two weeks ago]: Major religious leader L was arrested at his home in city A in Orange and is now held in custody by the Orange military.

E21 HUMINT: [From source R_3, who also provided us with E12 HUMINT and E13 HUMINT]: Says he heard the announcement in E20 SIGINT while he was still in country Orange.

E22 IMINT: [From an overhead taken two weeks ago and interpreted by one of our own image analysts]: Evidence of physical training areas in another remote region of country Orange.

E23 HUMINT: [from interview of person X while in sanctuary in country Green]: X says that these physical training sites are operated by his Blue group.

E24 HUMINT: [From interview of person Y while in sanctuary in country Green]: Y says his family has been held hostage by Gen. Alpha's military forces.

E25 SIGINT: [Recording of clandestine radio broadcast two weeks ago]: Report of a successful ambush of Orange military vehicles in a remote area of Orange.

E26 IMINT: [From an overhead taken two weeks ago and interpreted by one of our own image analysts]: Burned military vehicles observed in area described in E25 SIGINT.

E27 DOC: [Provided by a recent visitor in Orange]: A leaflet, said to be widely circulated in cities A and B describing the arrest and detention of religious leader L.

E28 HUMINT: [Another Orange refugee R_4 now in sanctuary in country Green]: Says that Y's agricultural policies, while he was in the Orange government five years ago, were very acceptable to Orange farmers.

E29 HUMINT: [From a recent visitor in country Orange]: Religious leaders in country Orange are privately urging their followers to resist Gen. Alpha's government.

E30 HUMINT: [From an alleged defector from country Orange military]: Reports that Gen. Alpha has requested emergency military assistance in the form of tracked vehicles and tactical missiles from an international arms supplier.

E31 HUMINT: [From an interview with person X while he was in sanctuary in country Green]. Person X says that he believes that the Blues are the only organized insurgency group operating in country Orange.

2.5.4 Marshaling the Items of Information to Determine their Relevance as Evidence

The first thing we ask you to do in thinking about this case is to regard the 31 items just listed as items of information or, if you like, data. But we have listed these 31 items as being evidence. The trouble is that no items of information or data become evidence until their relevance is established by defensible arguments to hypotheses we are trying to prove or disprove. In this case concerning General Alpha, your major or ultimate hypothesis is H which you decomposed into four sub-hypotheses: H_1, H_2, H_3, and H_4 as described above, each of which provides a touchstone for determining the relevance of the 31 items of data.

You will employ here the *argument magnet* (see Section 2.4.5) and determine which of the items of information in Table 4 is relevant to which of the four sub-hypotheses. Now, the task of judging the relevance of each of these 31 data items as evidence is no easy task. Consider, for example, determining which of the items are relevant to H_3, concerning the vulnerability of Gen. Alpha's military:

- Items E5 and E6 suggest that Alpha's military forces are vulnerable to attacks in urban areas.
- Items E10, E17, and E30 suggest that members of Gen Alpha's military are aware of their vulnerability to attack.

We determined that the following items are relevant to H_3: E5, E6, E10, E11, E15, E17, E25, E26, E30. As part of Exercise 33 you will determine which items are relevant to H_1, H_2, and H_4.

2.5.5 Beginning the Construction of Defensible Chains of Reasoning

What we have is an existing collection of evidence and we must judge how this evidence bears on these sub-hypotheses. This is an abductive reasoning process that proceeds from the bottom-up, from the evidence to the sub-hypotheses. Reasoning upward, these collections of evidence suggest further sub-hypotheses under sub-hypothesis H_3:

- E11, E15, E25, and E26 show vulnerability of Alpha forces in rural areas,
- E5 and E6 show vulnerability of Alpha forces in urban areas, and
- E10, E17, and E30 show that Alpha forces know they are vulnerable to attack.

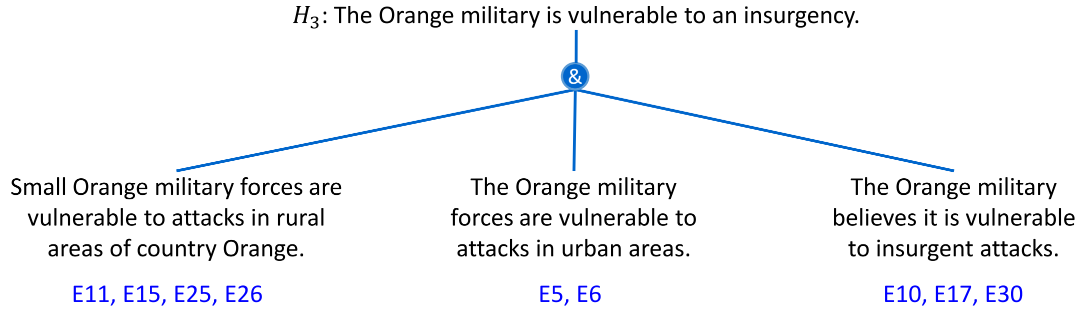
H_3: The Orange military is vulnerable to an insurgency.
&
Small Orange military forces are vulnerable to attacks in rural areas of country Orange.
E11, E15, E25, E26
The Orange military forces are vulnerable to attacks in urban areas.
E5, E6
The Orange military believes it is vulnerable to insurgent attacks.
E10, E17, E30

Figure 33. Evidence marshaling using the argument magnet.

2.5.6 Continuing the Construction of Defensible Chains of Reasoning

Here comes the most difficult and challenging, but also the most interesting, part of our argument construction process. So far, we have not completed the abductive and bottom-up part of careful argument construction in defense of relevance of our evidence. As an illustration, consider what we have just said concerning H_3 and its supporting evidence (i.e., E11, E15, E25, and E26) showing the vulnerability of Alpha forces in rural areas of country Orange. What is left out so far are arguments saying why these four items of evidence, taken together, allow us to infer this possible vulnerability of Alpha forces in rural areas. A possible argument structure based on these four items of evidence is shown in Figure 34.

The reader should note that each of the hypotheses in this structure is a source of doubt. Also, what is left out in this example are the foundation stages of every argument, namely the credibility of the source(s) of our evidence, whether the evidence is testimonial from human sources or tangible evidence from objects or sensors of various kinds. Similarly, we develop the arguments in Figure 35 and Figure 36.

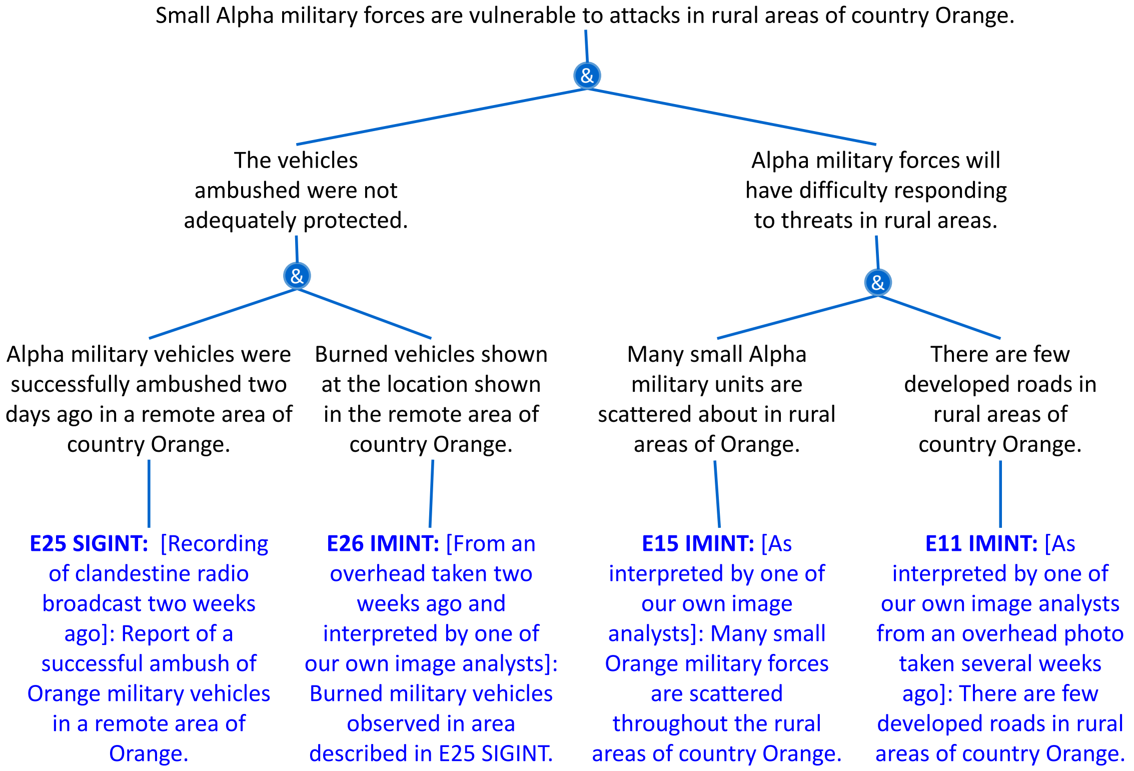

Small Alpha military forces are vulnerable to attacks in rural areas of country Orange.

&

The vehicles ambushed were not adequately protected.

Alpha military forces will have difficulty responding to threats in rural areas.

&

Alpha military vehicles were successfully ambushed two days ago in a remote area of country Orange.

Burned vehicles shown at the location shown in the remote area of country Orange.

&

Many small Alpha military units are scattered about in rural areas of Orange.

There are few developed roads in rural areas of country Orange.

E25 SIGINT: [Recording of clandestine radio broadcast two weeks ago]: Report of a successful ambush of Orange military vehicles in a remote area of Orange.

E26 IMINT: [From an overhead taken two weeks ago and interpreted by one of our own image analysts]: Burned military vehicles observed in area described in E25 SIGINT.

E15 IMINT: [As interpreted by one of our own image analysts]: Many small Orange military forces are scattered throughout the rural areas of country Orange.

E11 IMINT: [As interpreted by one of our own image analysts from an overhead photo taken several weeks ago]: There are few developed roads in rural areas of country Orange.

Figure 35. Possible argument for the vulnerability of Alpha forces in rural areas of Orange.

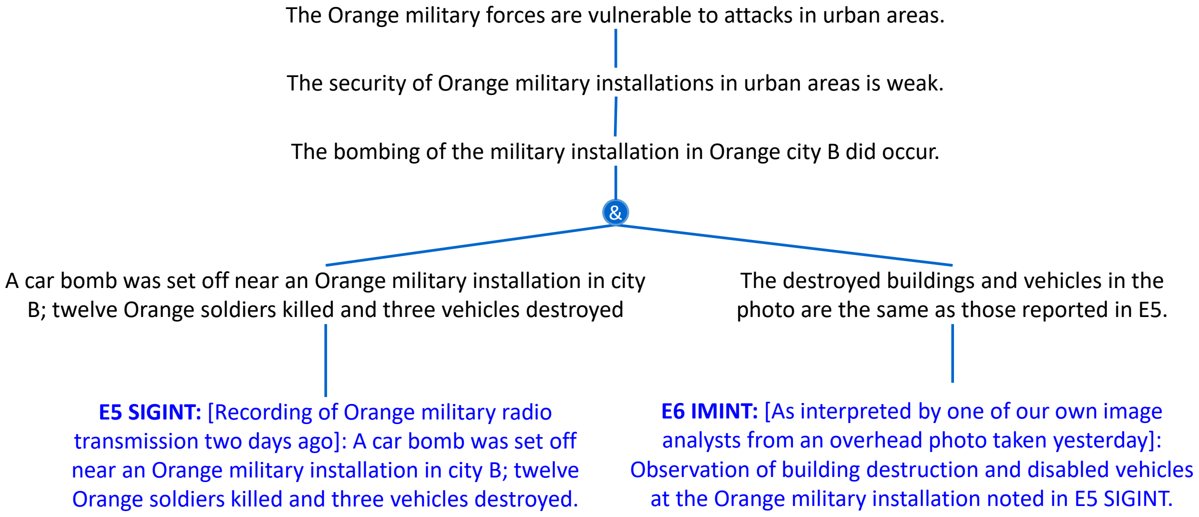

The Orange military forces are vulnerable to attacks in urban areas.

The security of Orange military installations in urban areas is weak.

The bombing of the military installation in Orange city B did occur.

&

A car bomb was set off near an Orange military installation in city B; twelve Orange soldiers killed and three vehicles destroyed

The destroyed buildings and vehicles in the photo are the same as those reported in E5.

E5 SIGINT: [Recording of Orange military radio transmission two days ago]: A car bomb was set off near an Orange military installation in city B; twelve Orange soldiers killed and three vehicles destroyed.

E6 IMINT: [As interpreted by one of our own image analysts from an overhead photo taken yesterday]: Observation of building destruction and disabled vehicles at the Orange military installation noted in E5 SIGINT.

Figure 34. Possible argument for the vulnerability of Alpha forces in urban areas of Orange.

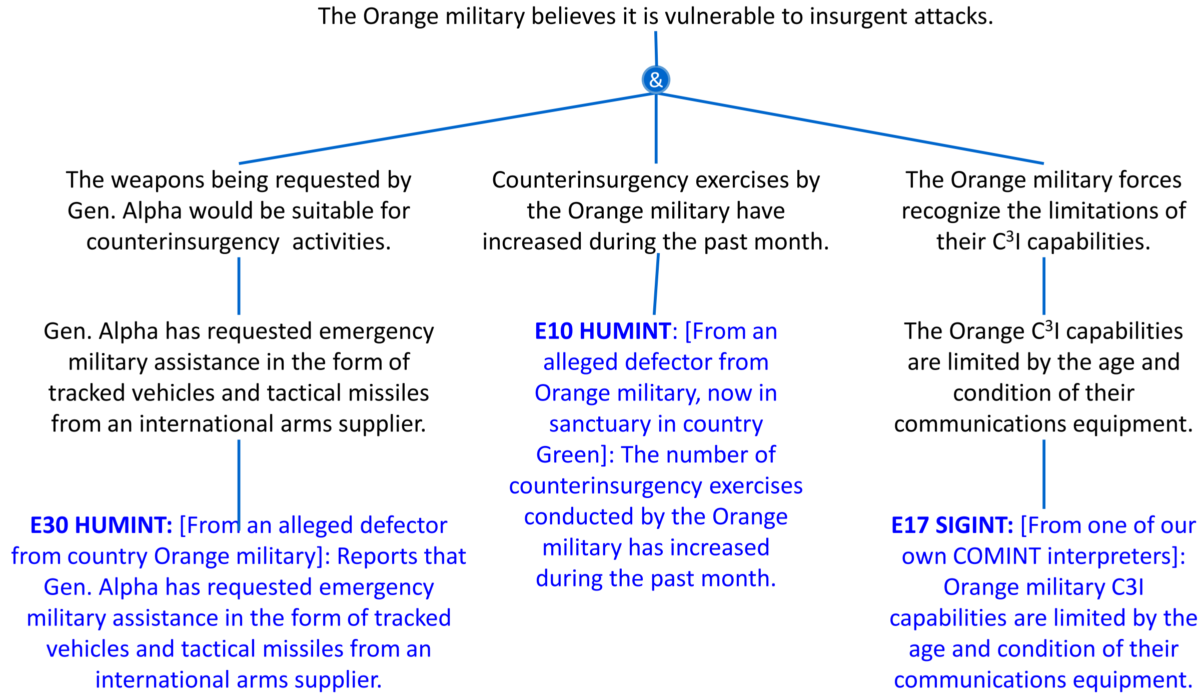
The Orange military believes it is vulnerable to insurgent attacks.
&
The weapons being requested by Gen. Alpha would be suitable for counterinsurgency activities.
Counterinsurgency exercises by the Orange military have increased during the past month.
The Orange military forces recognize the limitations of their C3I capabilities.
Gen. Alpha has requested emergency military assistance in the form of tracked vehicles and tactical missiles from an international arms supplier.
E10 HUMINT: [From an alleged defector from Orange military, now in sanctuary in country Green]: The number of counterinsurgency exercises conducted by the Orange military has increased during the past month.
The Orange C3I capabilities are limited by the age and condition of their communications equipment.
E30 HUMINT: [From an alleged defector from country Orange military]: Reports that Gen. Alpha has requested emergency military assistance in the form of tracked vehicles and tactical missiles from an international arms supplier.
E17 SIGINT: [From one of our own COMINT interpreters]: Orange military C3I capabilities are limited by the age and condition of their communications equipment.

Figure 36. Possible argument showing that Alpha forces know they are vulnerable to attack.

2.5.7 Final Words about General Alpha

As we mentioned at the beginning of this case of insurgency against General Alpha, we have used this case to illustrate how necessary it is to marshal evidence in preparation for the all-important task of constructing defensible arguments concerning the relevance of evidence on matters we wish to try to prove or disprove. We considered a collection of 31 items of evidence. This number would of course be very small relative to the number of potential evidence items you would enconter in actual episodes of intelligence analysis. But even this small number shows the great complexity of tasks you face in drawing defensible conclusions from this evidence.

Notice that we have given you a collection of 31 items that were all relevant to your required conclusion. Another way of saying this is to say that we have found an inferential home for every one of these evidence items in arguments necessary for supporting your conclusion. But this will never be a feature of actual intelligence situations you face. So much of the information or data you encounter will not be relevant to the inferential task you now face. You might be tempted to regard this irrelevant data as being just “noise”. However, you should also be alert to the possibility that patterns of apparently irrelevant evidence might be carefully contrived by an adversary to lead you away from considering hypotheses that do capture an adversary’s true intentions. As you certainly know, intelligence analyses take place in a hostile world in which potential or real adversaries will do all they can to mislead you. We consider these possibilities in other cases we have included in this book.

2.6 Review Questions

23. From any collection of information arranged in chronological order, a virtual infinity of different stories might be told. Suppose we have the following three items of evidence whose temporal ordering we have reason to believe:
    - Person Y agreed on March 4 to supply us with information about the military in his country.
    - On August 1, Source Y supplied us with a HUMINT report saying that the commanding general of the military was planning to launch a coup attempt against the elected leadership in his country on August 15.
    - On August 18, the leadership in this country announced that the commanding general of its military, along with several members of his staff, were being held in prison.

24. Think of all the different stories that might be told about why the event predicted in Y's HUMINT did not come to pass.

25. The use of index cards and shoeboxes to organize incoming intelligence information is old hat. Are there any computer-based methods you have tried? Have they been helpful in allowing you to generate hypotheses for any analysis you have been working on?

26. Here is a HUMINT source S who tells us that a person P has been assembling explosive devices in his garage. What kinds of questions should you be asking about and of source S? Have another look at the examples shown in Table 3 concerning questions of and about our sources,

27. In Section 2.4.2 we presented an example of the importance of event ordering to Person P who we assume does not wish to have a certain event ordering happen as he left work today; the event of concern involves his having consumed three double martinis. Can you think of other cases in which event ordering is so important?

28. What other questions seem natural to ask about the terrorist organization described in Section 2.4.3?

29. Hypotheses become most useful "magnets" for attracting productive combinations of evidence to consider as we illustrated in Section 2.4.4. Here we consider instances of hypotheses in search of evidence we mentioned earlier. As an example, suppose we form the hypothesis that S is a credible source of information about an important event E. This source might either tell us that event E occurred or it did not occur. What evidence would

we need to justify our hypothesis that S is credible?

30. Consider the argument "magnet" described in Section 2.4.5. Here we must consider arguments favoring or disfavoring of sub-hypotheses we are considering. Consider again testing your hypothesis that S is a credible source. What arguments should you be prepared to offer in support of this hypothesis?

31. Consider the situation shown in Section 2.4.6 in which we are concerned with the leakage of classified information from an intelligence office. You have been charged with investigating the activities of a person Y who is suspected of being the leaker. As a result of your investigation you report that person Y can be eliminated from consideration. At some time later, the classified documents are found on a laptop belonging to Y, and Y admits to having been the leaker. You are then confronted by your boss who says, "You managed to seize defeat from the jaws of victory, how could you have been so foolish? You had Y and you let him go. You made all of us look bad and I am considering demoting you." What defense can you offer your boss and perhaps preserve your position?

32. We hope you will appreciate the many heuristic virtues of telling yourself stories or constructing scenarios based on evidence you have gathered. From the same collection of available evidence, you may be able to tell an array of different stories depending on the "gap-fillers" or hypothetical events you include. Every different story you can tell suggests different hypotheses and new lines of evidence you might consider. Look again at the two different stories we tell based on the same evidence as shown in Figure 31. What is another different story you could tell?

33. Determine which items from Table 4 are relevant to which of the following sub-hypotheses:
    H_1: The Blues enjoy popular support among the citizens of country Orange.
    H_2: The Blues will have the military capability necessary for the insurgency to succeed.
    H_4: The Blue group leadership is adequate to make the insurgency successful.

34. Develop an argument structure showing how the evidence supports the hypothesis H_1:
    The Blues enjoy popular support among the citizens of country Orange.

35. Develop an argument structure showing how the evidence supports the hypothesis H_2:
    The Blues will have the military capability necessary for the insurgency to succeed.

36. Develop an argument structure showing how the evidence supports the hypothesis H_4:
    The Blue group leadership is adequate to make the insurgency successful.

3 EVIDENCE

3.1 What is Evidence?

Marcus Aurelius

Everything we hear is an opinion, not a fact. Everything we see is a perspective, not the truth.

In his marvelous book *Evidence, Proof, and Facts: A Book of Sources*, Law Professor Peter Murphy (2003, p.1) provides the following definition to the term evidence:

"In its simplest sense, evidence may be defined as any factual datum which in some manner assists in drawing conclusions, either favorable or unfavorable, to some hypothesis whose proof or refutation is being attempted."

He notes that this term is appropriate not only in law, but in any field in which conclusions are reached from any relevant datum. Thus, physicians, scientists of any ilk, historians, and persons of any other conceivable discipline, as well as ordinary persons, use evidence every day in order to draw conclusions about matters of interest to them.

Difficulties arise when a variety of terms are used as synonyms for the term evidence: *data*, *items of information*, *facts*, and *knowledge*. When examined carefully, there are some valid and important distinctions to be made among these terms that are not always apparent, as will we now discuss.

3.1.1 Evidence, Data, and Information

Data are un-interpreted signals, raw observations, or measurements, such as the 137 or STEMQ.

Information is data equipped with meaning provided by a certain context, such as "cesium-137", "STEMQ Company", "or "STEMQ customer."

There are untold trillions of data and items of information in existence that will almost certainly never become evidence in most inferences. As we will discuss, *items of information only become evidence when their relevance is established regarding some matter to be proved or disproved.*

3.1.2 Evidence and Fact

Now consider the term *fact*; there are some real troubles here as far as its relation to the term evidence is concerned. How many times have you heard someone say, "I want all the facts before I draw a conclusion or make a decision?" Or, "I want to know the facts in this matter?" The first question is easily answered; we will never have all the facts in any matter of inferential interest. Answers to the second question require a bit of careful thought.

Here is an example of what is involved:

> Clyde tells us that the lock had been forced. Now we regard it as fact that Clyde gave us this information. But whether the lock had been forced *is only an inference and is not a fact*. This is precisely why we need to carefully distinguish between an *event* and *evidence* for this event.

Here is what we have:
Clyde has given us evidence E*, saying that event E occurred, where E is the event that the lock had been forced. Whether this event E did occur or not is open to question and depends on Clyde's credibility. If we take it as *fact* that event E did occur, just because Clyde said it did, we would be overlooking the *credibility* foundation for any inference we might make from his evidence E*. Unfortunately, it so often happens that people regard the events reported in evidence as being facts when they are not. Doing this suppresses all uncertainties we may have about the source's credibility. We have exactly the same concerns about the credibility of tangible evidence. For example, we have been given a tangible object or an image as evidence E* that we believe reveals the occurrence of event E. But we must consider whether this object or image is authentic and it is what we believe it to be. In any case, the events recorded in evidence can only be regarded as facts if provided by perfectly credible sources, something we almost never have. As another example, any information we find on the Internet should be considered as only a claim by its source rather than as fact, that is, as *evidence* about a potential fact rather than a *fact*.

3.1.3 Evidence and Knowledge

Now consider the term *knowledge* and its relation with evidence. Here is where things get interesting and difficult. As you know, *the field of epistemology is the study of knowledge, what we believe it may be, and how we obtain it.* Two questions we would normally ask regarding what Clyde just told us are as follows:

- Does Clyde really know what he just told us, that the lock had been forced?
- Do we ourselves then also know, based on Clyde's testimony, that the lock had been forced?

Let's consider the first question. For over two millennia, some very learned people have troubled over the question: What do we mean when we say that person A *knows* that event B occurred? To apply this question to our source Clyde, let's make the assumption that will simplify our answering this question. Let's assume that Clyde is a *competent* observer in this matter. Suppose we have evidence that Clyde is a professional locksmith and he actually himself looked at the lock. This is a major element of Clyde's credibility.

Here is what a standard or conventional account says about whether Clyde knows that the car did not stop at the red-light signal. First, here is a general statement of the standard account of knowledge: *knowledge is justified true belief*. Person 𝒜 knows that event B occurred if:

- Event B did occur, [True]
- 𝒜 got non-defective evidence that B occurred, [Justified], and
- 𝒜 believed this evidence [Belief]

This standard analysis first says that event B must have occurred for 𝒜 to have knowledge of its occurrence. This is what makes 𝒜's belief true. If B did not occur, then 𝒜 could not know that it occurred. Second, 𝒜's getting non-defective evidence that B occurred is actually where 𝒜's competence arises. 𝒜 could not have gotten any evidence, defective or non-defective, if 𝒜 was not where B could have occurred. Then, 𝒜 believed the evidence 𝒜 received about the occurrence of event B; and 𝒜 was justified in having this belief by obtaining non-defective evidence of B's occurrence.

So, in the case involving Clyde's evidence, Clyde knows that the lock had been forced if:

- The lock had been forced,
- Clyde got non-defective evidence that the lock had been forced,
- Clyde believed this evidence.

If all of these three things are true, we can state on this standard analysis that Clyde knows that the lock had been forced.

Before we proceed, we must acknowledge that this standard analysis has been very controversial in fairly recent years and some philosophers claim to have found alleged paradoxes and counterexamples associated with it. Other philosophers dispute these claims. Most of the controversy here concerns the justification condition; what does it mean to say that A is justified in believing that B occurred? In any case, we have found this standard analysis very useful as a heuristic in our analyses of the credibility of testimonial evidence.

But now we have several matters to consider in answering the second question: Do we ourselves also *know*, based on Clyde's testimony, that the lock had been forced? The first and most obvious fact is that we do not know the extent to which any of the three events just described in the standard analysis are true. We cannot get inside Clyde's head to obtain necessary answers about these events. Starting at the bottom, we do not know for sure that Clyde believes what he just told us about the lock being forced. This is a matter of Clyde's *veracity* or *truthfulness*. We would not say that Clyde is being truthful if he told us something he did not believe.

Second, we do not know what sensory evidence Clyde obtained, on which to base his belief, and whether he based his belief at all on this evidence. Clyde might have believed that the lock had been forced either because he expected or desired this to be true. This involves Clyde's *objectivity* as an observer. We would not say that Clyde was objective in this observation if he did not base his belief on the sensory evidence he obtained in his observation.

Finally, even if we believe that Clyde was an objective observer who based his belief about the lock on sensory evidence, we do not know how good this evidence was. Here we are obliged to consider Clyde's sensory *sensitivities or accuracy in the conditions under which Clyde made his observations*. Here we consider such obvious things as Clyde's visual acuity. But there are many other considerations such as, "Did Clyde only get a fleeting look at the lock?" For a variety of such reasons, Clyde might simply have been mistaken in his observation; the lock had not been forced.

So, what it comes down to is that the extent of our knowledge about whether the lock had been forced, based on Clyde's evidence, depends on these attributes of Clyde's credibility: his competence, veracity, objectivity, and observational sensitivity. We will have much more to say about assessing the credibility of sources of evidence, and how Cogent can assist you in this difficult process, in Section 4 of this book.

As we have noted on several occasions in the preceding sections of this book, evidence has three major properties or credentials: relevance, credibility, and inferential force or weight.

3.2 The Credentials of All Evidence

As we have noted on several occasions in the preceding sections of this book, evidence has three major properties or credentials: relevance, credibility, and inferential force or weight. Let's now make sure that we understand what these three evidence credentials mean and why they are so important.

3.2.1 Relevance

We know of no better definition of the term relevance, applied to evidence, than the one provided in the *Federal Rules of Evidence for United States Courts* (Mueller and Kirkpatrick, 2009). These Federal Rules of Evidence (FREs) govern the offering and admissibility of evidence introduced in the U.S. courts, and they are the result of centuries of experience in the Anglo-American system of common law. But these rules are not set in stone and can be revised in light of new experience and insight. One of these numbered rules, FRE-401, defines relevant evidence as follows:

> *"Relevant evidence" means evidence having any tendency to make the existence of any fact that is of consequence to the determination of the action more probable or less probable than it would be without the evidence. (FRE-401, Federal Rules of Evidence, 2009 Edition)*

We need to parse this definition in order to see what it says and how it applies to intelligence analysis. First, consider the phrase, "any fact that is of consequence to the determination of the action." This basically refers to the matters to be proved or disproved; i.e., it refers to the hypotheses being considered. In law, the hypotheses will involve some charge in a criminal case or a complaint lodged in a civil case. For example, in a murder case the determination of the action would involve such hypotheses as: "The defendant being charged was the one who unlawfully killed the victim." In intelligence analysis one hypothesis that might be considered is: "At least one terrorist action will be initiated against homeland USA in the next year." In both the law and intelligence examples just mentioned, these would be major or upper-level hypotheses. But in constructing arguments bearing on these major hypotheses, there will be any number of lower-level hypotheses involving the sources of doubt we recognize in our arguments. Thus, evidence is relevant if it bears either directly or indirectly on any of our recognized sources of doubt.

Now consider the phrase "more probable or less probable than it would be without the evidence." What this says is that relevant evidence causes us to change our beliefs, one way or the other, about some hypothesis being tested. But notice that it does not tell us how much we should change our beliefs in these hypotheses. Another way of stating this is to say that relevant evidence has <u>some inferential force or weight</u>, but it does not say how much force or weight it should have. One very good reason why FRE-401 does not say how much force or weight evidence should have is that assessing the force or weight of evidence is a very complex matter involving the probabilistic strength of our arguments based on evidence that concern both relevance and credibility matters. In law and in intelligence analysis, reasonable people may disagree strongly about how forceful a given item of evidence, or some collection of evidence, is on hypotheses of interest.

Defending the relevance of evidence is no easy task especially when we have masses of it to consider, as we do in intelligence analysis. There is a problem associated with what we termed <u>evidential synergism</u> that we discussed in Section 1.3.4. Suppose we have an item of information about event A that, by itself, seems irrelevant or useless as far as an analysis problem we are currently facing and we contemplate discarding it. But someone notices that we also have other items about events B and C that would make item A seem relevant when we take all these three items together. The point here is that the relevance of an item of information often depends on what other items of information we have. Another way of stating this is to say that items of information interact or are dependent in ways that influence their relevance on hypotheses being considered. As we also emphasized, the detection of these interactions producing evidential synergisms depends on how devoted the analysts are to sharing information across agencies and offices in the intelligence services.

> *<u>Relevance</u> answers the question: So what? How is this item linked to any hypothesis or possible conclusion of interest in an intelligence analysis?*

Example 7. Suppose your intelligence analysis involves inferences about possible forms of terrorist actions that could be taken against targets here in the United States. Here are two items of information you have just received:

> Item #1: Professor David A. Schum was the owner of a green 2000 Toyota Corolla vehicle carrying Virginia license plate # TSL-782.

The first question you would ask about Item #1 is: "So what? What conceivable bearing does this bit of information have on any possible conclusions I could reach about possible terrorist incidents here in the United States?" Unless David Schum was associated with any terrorist organization, you would be justified in saying that this datum is totally irrelevant in your present analysis. The concept of relevance concerns our attempts to answer "So what?" questions regarding items of information we have. Now consider this second datum:

> Item #2: The STEMQ company in Baltimore, MD, manufactures devices for sterilizing medical equipment. A person named Willard reported in a *Washington Post* article that a canister containing powder, including approximately 3000 curies of cesium-137, had gone missing from the company's warehouse in Baltimore. There were indications that the storage area where this powdered cesium-137 was located showed signs of forcible entry.

Asking the same question, "So what?" an analyst would give a different answer to this second item than the one given to Item #1. Cesium-137 has all sorts of uses in medicine and in industry. It is used to sterilize food, to manufacture thickness and moisture density gauges, and it is used for various diagnostic purposes in medicine in addition to the sterilization purposes just mentioned. Unfortunately, cesium-137 could be put to other uses including the construction of a dirty bomb. Just a few ounces of this powdered material set off by a conventional explosive in a bomb could contaminate an entire city for decades. Indeed, it is known that the Chechen Mujahidin placed a canister of cesium in a Moscow park in the hopes of spreading radiation throughout the city. Fortunately, this canister was discovered by the police before it was set off. So, Item #2 seems relevant in inferences about what kind of actions terrorist groups may be contemplating here in the United States.

The arguments concerning the relevance of evidence involve chains of reasoning. Consider an item of relevant evidence and the argument or chain of reasoning an analyst has constructed that links this item to hypotheses being considered. Here starts the initial root of uncertainty in intelligence analysis. In forming this chain of reasoning, each identified link in a relevance chain of reasoning involves a proposition that might be true or not; i.e., it is a source of doubt or uncertainty. The argument being constructed thus forms a chain of sources of doubt the analyst believes to be interposed between the evidence and what the analyst is trying to prove or disprove from it. These links or sources of doubt are laid out in a logically consistent order in which one link is inferable from its predecessor.

Where do these links (sources of doubt) come from? They come from the analyst's imagination based upon his/her experience and knowledge of the analytic problem area.

Here we have the roots of disagreements among analysts concerning the relevance of evidence and disagreements about uncertainties in conclusions that may be reached. Different analysts may construct different arguments from the same evidence and thus perceive different sources of uncertainty. Even if they agree about the links in an argument, they may disagree about how strong they are. It is also true, of course, that analysts may generate different possible hypotheses from the same evidence.

A relevance argument having been constructed, it is now time for its evaluation; here is where the necessity for critical reasoning arises. The analyst must ask the following question:
Is the chain of reasoning I have just constructed that shows the relevance of this evidence item logically coherent? Does it contain any disconnects or non sequiturs?

Here is where the necessity for the defensibility of an analytic argument arises. Should anyone else take an analyst's argument seriously as far as the relevance of this evidence is concerned? As we will see in a moment, this question also involves the extent to which anyone else should take the analyst's assessments of uncertainty seriously. An absolutely crucial element in intelligence analysis is the defensibility of arguments constructed by analysts from the evidence available to possible conclusions to be proved or disproved. Of course, it is also true that an argument must be persuasive and that not all defensible arguments are persuasive. One fairly certain way of failing to make an argument persuasive is to have it revealed that it is not defensible on logical grounds. But there is an important point here that needs further elaboration: *There is no such thing as a uniquely correct or perfect argument from evidence*.

Here is a chain of reasoning an analyst has just constructed to establish the relevance as evidence of an item of information. Someone, another analyst or perhaps another critic, may be able to find what this person believes to be a missing link or a link that is improperly stated. Perhaps this fellow analyst or critic discovers a disconnect in the analyst's initial relevance argument. What no one, be they analyst or critic, can say is that they have the "correct" argument justifying the relevance of this item of evidence. Someone can come along later and discover inadequacies in the revised argument proposed by the other analyst or critic. *What can be done is to have someone identify defects in an argument. But what no one can do is to say that the argument being proposed is the final or ultimately true argument that anyone could propose*. What this says is that someone can correct defects in an argument without ever being able to say that they have the only argument that could ever be made regarding the relevance of evidence. This also accounts for the fact that there will always be some disagreement about the uncertainty that is assessed, combined, and reported among analysts themselves and among persons for whom the analysis was performed. Different persons will construct different arguments from the same bodies of evidence.

The fact that there is no such thing as the ultimately true argument has an important bearing on the many examples we provide of various evidential and inferential issues. We will never say that the arguments we construct in providing examples are the only correct or true ones. You, the reader, may perceive other reasoning routes from the evidence we present and may indeed see other possible conclusions. Our major hope of course is that you will not see any disconnects or non sequiturs in our arguments.

3.2.2 Credibility

Here is an item of evidence E_1^*, from source I, that all of us believe is relevant on hypotheses we are considering. We strongly believe that we can defend the relevance of event E, as reported in our evidence E_1^*, by the argument we have constructed that we all agree is very strong and free of disconnects or non sequiturs. But the crucial question remains:

> *How certain are we that event E did occur just because source I said it did?*
> *Source I's credibility is the major issues here.*

This is why we have said that source credibility considerations form the very foundation for all arguments we make from evidence to our hypotheses. We illustrated this fact in Credibility foundation of argument. in Section 1.3.2. However strong our relevance argument may be, if it rests on a weak credibility foundation, it will falter. Perhaps the best example of inferential calamities known to us in the open-source literature concerns the credibility of the human source called "Curveball" (Bruce, 2008).

The Federal Rules of Evidence contain many rules associated with testing the credibility of witnesses and the authenticity, reliability, and accuracy of tangible evidence (Mueller and Kirkpatrick, 2009). But there are no rules concerning how we ought to grade credibility and how strong it should be. As far as witness credibility is concerned, we have consulted over a hundred works on evidence in law that contain most valuable accumulated strategies for supporting or undermining human testimonial credibility (i.e. HUMINT credibility), and the credibility of various forms of tangible evidence. These strategies have been accumulated over six hundred years of experience in our courts and concern the competence, veracity, objectivity, and observational sensitivity of witnesses. These accounts of credibility matters allow us to generate experience-tested questions to ask about the credibility of HUMINT sources and of the credibility-related attributes of tangible evidence from a variety of sources. These questions are discussed in Chapter 4. We add here that the Cogent system knows what questions you should try to answer regarding the credibility of different forms of evidence.

A final matter here concerns the correct use of the term reliability. So often we hear of persons being described as reliable sources. There is a difficulty associated with using this term with reference to human sources such as for HUMINT. The trouble is that the term reliability is most often used to indicate how consistent or repeatable some process is. You say your car is reliable to the extent to which it will take you where you wish to go for some period of time in the future. A test of any kind is reliable to the extent that it gives the same result over again if it is repeated. But there is so much more to the credibility of a human source than mere consistency; we have discussed how attributes of the credibility of sources of HUMINT concern their competence, veracity, objectivity, and observational sensitivity. What counts most is the credibility of what these sources tell us. The term reliability does not capture any of these attributes. Where we use the term reliability is with reference only to our sources of tangible evidence such as sensors of various kinds. In our work, we will use the term credibility with reference to our sources since what matters in all cases is whether an analyst can believe the information he or she has received.

3.2.3 Force or Weight of Evidence

In a very general sense, the force or weight of evidence indicates how strong the evidence is in favoring or disfavoring hypotheses we are considering. But this is as far as we can go, since there is considerable controversy about what the terms force and weight mean, and especially how these concepts should be assessed and combined. As we noted in Section 1.3.6, there are only two uncontroversial statements we can make about the force or weight of evidence. First, it has vector-like properties indicating the direction and the strength with which evidence favors or disfavors hypotheses we are considering. Second, the force or weight of evidence is always graded in probabilistic terms. This second statement is actually the greatest source of controversy since a variety of careful scholars in probability, who have given their days and nights to the study of evidence and probability, cannot agree about how force or weight of evidence should be assessed and combined. In Chapter 5 we will review such alternative probability methods.

Probabilistic judgments can be expressed numerically in several ways, and also in terms of words, as will be discussed in Section 5.7. Speaking of numerical judgments of probability, a very wise and devoted scholar, Professor Glenn Shafer (1988, pp. 5 - 9), has correctly noted:

> *"Probability is more about structuring arguments than it is about numbers. All probabilities rest upon arguments. If the arguments are faulty, the probabilities however determined, will make no sense."*

We add the same concern about verbal assessments of probability, such as "very probable", "probable", "unlikely", and so on. If these arguments are not defensible, no one will take seriously any numerical or verbal assessments we make concerning the force or weight of our evidence. This is just one reason why we consider the construction and defense of our relevance and credibility arguments so carefully in the sections to follow.

So, we have found it much easier to provide definitions and meaning to the relevance and credibility credentials of evidence than we have been able to do for the force or weight credential. All we can say at this point is that the force of evidence depends on the strength of our credibility and relevance arguments, as shown in Figure 37 But as we proceed, will provide an assortment of examples about assessing the force or weight of evidence that will be useful in intelligence analysis. Remember that all conclusions in intelligence analysis, in common with all other contexts in which conclusions are based on evidence, must always be hedged probabilistically by some means. This is because, as we will again discuss in Section 3.5, our evidence is always incomplete, usually inconclusive, frequently ambiguous, commonly dissonant, and comes to us from sources having any gradation of credibility shy of perfection. As we will observe, alternative views of probabilistic reasoning capture some of these important considerations, but no single view captures them all.

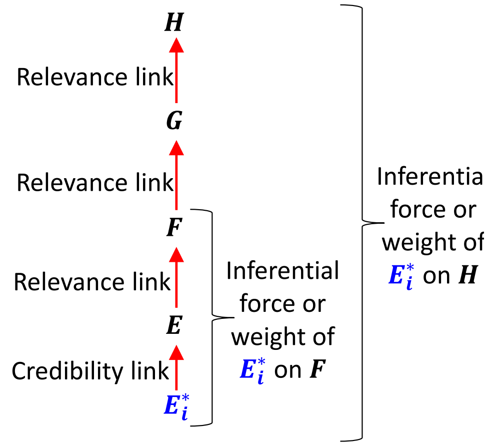
H
Relevance link
G
Relevance link
F
Relevance link
E
Credibility link
E_i^*
Inferential force or weight of E_i^* on F
Inferentia force or weight of E_i^* on H

Figure 37. Credentials of evidence.

3.3 Types of Evidence

3.3.1 Substance-blind Classification of Evidence

We can classify all evidence, regardless of its substance or content, into just a few categories of recurrent forms of evidence, as shown in Figure 38. This classification is called substance-blind because is based on its inferential properties rather than upon any feature of its substance or content. Knowledge of these substance-blind forms and combinations of evidence pays great dividends. Such knowledge informs us and Cogent how to evaluate the credibility of evidence, based on its type. It allows us to more easily assess evidence coming from different sources and to compare the evidence and conclusions reached from it in different intelligence analyses at different times.

Here is an important question we are asked to answer regarding the individual kinds of evidence we have:

> *How do you, the analyst, stand in relation to this item of evidence? Can you examine it for yourself to see what events it might reveal? If you can, we say that the evidence is tangible in nature. You can examine it and apply your own senses in a determination of what the evidence may be telling you. We might say that in assessing tangible evidence your own senses provide a direct interface with events of interest.*

As we will discuss momentarily, there are many forms of tangible evidence. But suppose instead you must rely upon other persons, assets, or informants, to tell you about events of interest. Their reports to you, about these events, are examples of testimonial evidence. You yourself were not privy to the occurrence or nonoccurrence of these events and so you make inquiries of these assets who may have observed these events. It is of vital interest to know how these persons obtained the information they report. And, as we will observe, the human sources of testimonial evidence can express varying degrees of uncertainty about what they have observed.

3.3.2 Tangible Evidence

The adjective "tangible" comes from the Latin *tangere*, meaning "to touch"; and *tangibilis*, meaning "that may be touched". In a very strict interpretation, tangible items are those we can touch or examine directly by our tactile senses. But this term is now used with reference to any item that can be seen, heard, tasted, smelled, or touched. For many years in the field of law, even today in some works, tangible evidence was/is referred to as "real" evidence, as opposed to testimonial evidence. The essential distinction has been that evidence that is "real" can be examined directly by the fact finders [jurors or judges] themselves. In intelligence analysis the "fact finders" will be intelligence analysts. They draw conclusions based on their own direct sensory examinations of various things. Testimonial evidence, on the other hand, is a report by another person such as an external witness, an asset, or an informant who has allegedly had some direct sensory experience with events of interest. One way to describe this difference is to say that the persons drawing conclusions from "real" evidence have a direct sensory interface with the events possibly revealed in such evidence. But in testimonial evidence it is the witness

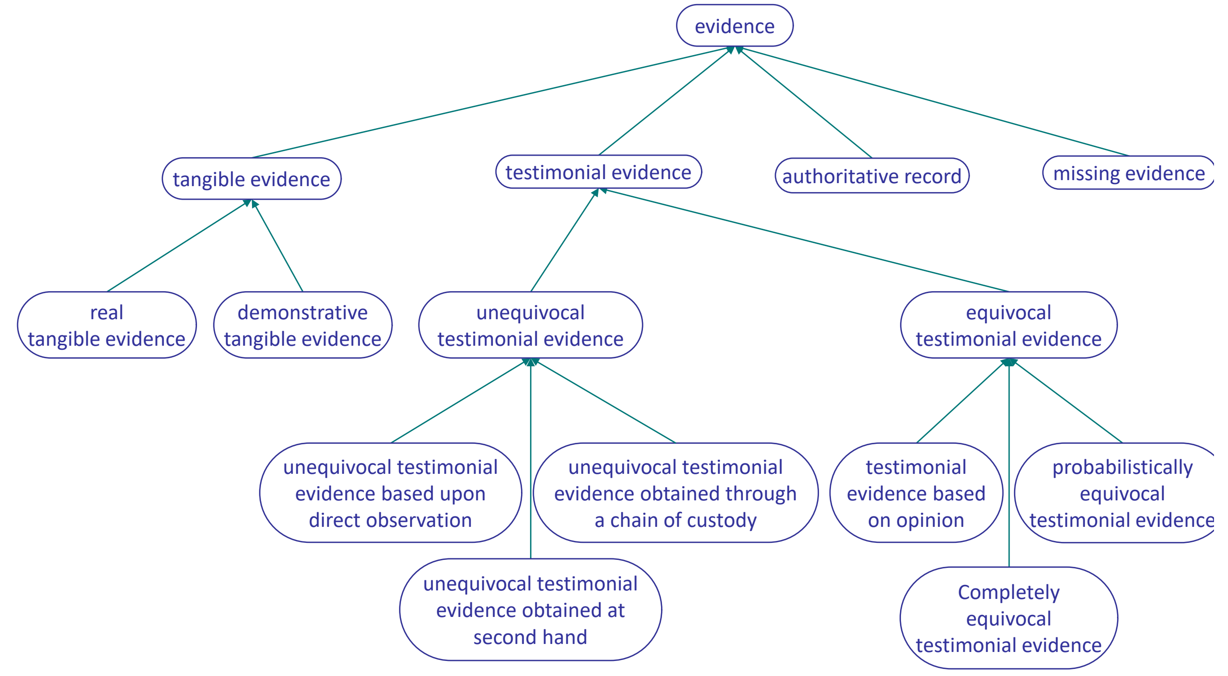
evidence
tangible evidence
testimonial evidence
authoritative record
missing evidence
real tangible evidence
demonstrative tangible evidence
unequivocal testimonial evidence
equivocal testimonial evidence
unequivocal testimonial evidence based upon direct observation
unequivocal testimonial evidence obtained through a chain of custody
unequivocal testimonial evidence obtained at second hand
testimonial evidence based on opinion
probabilistically equivocal testimonial evidence
Completely equivocal testimonial evidence

Figure 38. Substance-blind classification of evidence.

or asset, not the fact finders or analysts, who allegedly had sensory interface with the events of interest. The term "real" evidence is still preserved in an excellent recent work we will refer to again in this account of tangible evidence (Lempert et al., 2000, p. 1148).

Not everyone has been pleased with use of the term "real" evidence. This term may suggest that testimonial evidence has no status in reality and may somehow always be inferior. The term "real" evidence applied only to tangible objects may suggest that testimonal evidence is "unreal". Wigmore attempted to resolve this matter by using the term "autoptic proference" with reference to anything observable at trial by fact finders, including the testimony of witnesses that can be directly heard by the fact finders (Wigmore, 1937, p 11). The term "autoptic proference" translates as "a self-observable thing being offered" and refers to any evidence that can be perceived by fact finders through one of their senses. Wigmore used this term instead of the term "real" evidence that he considered ambiguous. The term "autoptic proference" is not in frequent use today and occurs only in works that dwell upon Wigmore's analytic and synthetic methods for constructing complex arguments or inference networks (Anderson et al., 2005, p 380). In this work tangible evidence is defined to be "evidence that can be directly examined by persons drawing conclusions to see what event(s) this evidence reveals"; testimonial evidence was defined as: "evidence provided by a human source to the evaluator of the evidence" (Anderson et al., 2005, p 386).

There is another term used in law that we might find very useful in our deliberations about tangible evidence in intelligence analysis. Lempert, Gross, and Leibman use the term "exhibits" with reference to tangible evidence of various sorts. They claim that all evidence provided at trial consists of either exhibits or oral testimonial assertions by witnesses (Lempert et al., 2000, p. 1145 - 1146). But Lempert, Gross, and Leibman provide a further very useful distinction involving tangible exhibits; this distinction appears in many other works on evidence in law (Lempert et al., 2000, p. 1146 - 1148). They use the term "real evidence" only with reference to the actual things themselves that may be evidence in the process of proof. Examples on law include the weapon used in a homicide, drugs taken from the defendant, blood or body tissues taken from the victim, documents signed by the parties in a dispute, and police surveillance photos of a robbery in progress. For such "real" evidence the Latin phrase "res ipsa loquitur" applies: the evidence speaks for itself. But other forms of tangible evidence are said to be "demonstrative"; examples include diagrams, maps, scale models, charts, statistical or other tabled measures, sound or video recordings, or simulations of various kinds. These items are not the things themselves but only representations or illustrations of things. What is always at issue for an item of demonstrative evidence is the extent to which it fairly and accurately represents the real object or thing at the relevant time. However, it happens that the credibility attributes for real and demonstrative evidence happen to be virtually the same ones: authenticity, accuracy, and reliability. This is one reason why in many instances the distinction between "real" and "demonstrative" evidence is not always emphasized.

There is an assortment of tangible items we might encounter, and that could be examined by an intelligence analyst. Both IMINT and SIGINT provide various kinds of sensor records and images that can be examined. MASINT and TECHINT provide various objects such as soil samples and weapons that can be examined. COMINT can provide audio recordings of communications that can be overheard and translated if the communication has occurred in a foreign language. We also note that documents are tangible evidence that can be examined by an analyst. These documents might have been captured or revealed by a human asset; but they also include any document obtained from open sources whatever they may be: newspapers, books, websites, etc. We also list as tangible items tabled measurements of any kind including statistical records; charts showing various kinds of scientific or technological relations; and maps and diagrams or plans of various kinds. Some of these kinds of items might be included in TECHINT or MASINT sources.

One thing we are obliged to note is that the analyst observing these tangible items may need the assistance of other analysts who have expertise in explaining the analyst what a tangible item reveals. For example, an expert in photo analysis may assist another analyst by showing what an image has revealed to us. Analysts whose assistance is required in such cases play the role of “expert witnesses”, so common in both criminal and civil trials. All of this highlights the fact that intelligence analysis is so often a cooperative venture involving teams of analysts each of whom may have particular knowledge and skills.

Here are several examples and questions involving evidence that is tangible and that the analyst can examine personally to see what events it reveals.

Following are some intelligence-relevant examples of tangible evidence. In these examples we will preserve the distinction between real and demonstrative evidence mentioned above. We will also distinguish between two major sources of tangible evidence: (1) HUMINT sources and (2) other sources such as sensing devices of various sorts. In the process, we will mention tangible evidence supplied by various INTs [such as IMINT, COMINT, MASINT, TECHINT]. We do so with awareness that these INTs keep changing.

The following are examples of real tangible evidence supplied by human sources.

- Here is an Iraqi national we have code-named "Rambo" who supplies us with the detonator of an IED he says he took from a disassembled weapon of this sort in a cache of weapons he found outside the Iraqi town of Penjwin near the Iraq-Iran border.
- Here is a listing, in the Pashtun language, of intended targets to be struck by the Taliban this month that was found on the body of a slain Taliban member by a Special Forces unit

operating near Gardez in Afghanistan. This listing has been translated into English by an Afghani national who works for us.

- Here is a listing of names and addresses of members of a Jihadist organization operating near Paris, France. This listing was provided by a source code-named "Sunshine" who says she copied this list from a laptop computer belonging to a member of this organization.

The following are examples of demonstrative tangible evidence provided by human sources.

- Here is a table showing a breakdown of arms supplied by the Iranian Islamic Revolutionary Guards Corp [IRGC] to Iraqi insurgent militia units during the preceding six-month period. This table was supplied for us by a source named "Lightning" who is an alleged defector from the IRGC.
- Here is a hand-drawn map taken from a member of a Shiite militia member who was captured by a Marine unit during a raid on the town of Al Kufa, south of Baghdad. According to the captured Iraqi, this map shows the location of other insurgent militia forces operating in the vicinity that were called upon to participate this attack.
- Here is a diagram showing the elements of a dirty bomb constructed with a core of semtex surrounded by a layer of cesium-137, together with a detonator. This diagram was given us by a source named "Moonshine" in Lublln, Poland.

The following are examples of real tangible evidence provided by non-human sensing devices.

- Here is a real-time IMINT video recording of the traffic observed entering and leaving an underground location two days ago in Karaj in Iran.
- Here is a COMINT recording of a phone conversation between a high-ranking Iraqi government official and msn known to be a member of the Iranian IRGC. This recording was taken last Thursday at 3 AM Baghdad time.
- Here is a satellite photo of military aircraft shown at an Iranian air force base outside of Esfahan in iran.

The following are examples of demonstrative tangible evidence provided by non-human sensing devices:

- Here is a computer-based chemical analysis of a MASINT soil sample collected near Leninsk in Russia.
- Here is an ELINT image showing the probable locations of radar installations within a hundred-mile radius of Esfahan in iran.
- Here is a computer-based recording of a series of entries taken from a known Islamic Jihadist website.

We are limited only by our own imaginations is seeing how many kinds of tangible evidence we might encounter in intelligence analysis. You may easily think of better examples than those we have just listed. But we hope we have at least shown that we have some useful ways of categorizing tangible evidence. Our next task is to identify the essential credibility attributes of these tangible forms of evidence. Following this, we will then consider attributes of the credibility of human sources who provide us with tangible evidence items.

3.3.3 Testimonial Evidence

As indicated in Figure 38, there are several types of testimonial evidence. If the source does not hedge or equivocate about what he/she observed (i.e., the source reports that he/she is certain that the event did occur), then we have unequivocal testimonial evidence*.* If, however, the source hedges or equivocates in any way (e.g., "I'm fairly sure that E occurred") then we have equivocal testimonial evidence. The first question we would ask this source of unequivocal testimonial evidence is, *"How did you obtain information about what you have just reported?"* It seems that this source has four possible answers to this question.

A first possible answer is, "I made a *direct observation* myself*.*" In this case we have unequivocal testimonial evidence based upon direct observation*.*

A second possible answer is, *"I did not observe this event myself but heard about its occurrence (or nonoccurrence) from another person."* Here we have a case of secondhand or hearsay evidence, called unequivocal testimonial evidence obtained at second hand.

A third possible answer is *"I did not observe this event myself but I learned about it from an intelligence report."* The truth is that lots of things might be done to an item of incoming intelligence information between the time it is first received, and the time it reaches the desk or the computer of intelligence analysts. Any number of persons or devices may have had access to this information item and may have done various things to it. One result is that the intelligence analyst may not have received either an authentic or a complete account of the actual information received from the source. In law this fact is recognized and so there are elaborate procedures for dealing with what are called chains of custody. Well-qualified persons are designated as evidence custodians who make careful records of every person who had access to an evidence item from the time it was received, what they did with this item, how long they held the item, and who next received the item before it was finally introduced at trial. Although we are not privy to the procedures for dealing with chains of custody of information received in intelligence analysis, and whether there are any persons who act as evidence custodians, we have written a paper for intelligence agencies on such matters that also illustrates how Disciple-LTA (Tecuci et al., 2005a; 2007b; 2008a), a forerunner of the Cogent system, can assist analysts to capture doubts associated with chains of custody of intelligence information (Schum et al., 2009).

A detailed discussion of the analysis of the chains of custody is provided in Chapter 4.5.

A forth answer is possible, *"I did not observe event E myself, nor did I hear about it from another source. But I did observe events C and D and inferred from them that event E definitely occurred."* This is called testimonial evidence based on opinion and it requires some very difficult questions. The first concerns the source's credibility as far as his/her observation of event C and D; the second involves our examination of whether we ourselves would infer E based on events C and D. This matter involves our assessment of the source's inferential ability It might well be the case that we do not question this source's credibility in observing events C and D, but we question the conclusion that event E occurred the source has drawn from his observations. We would also question the certainty with which the source has reported an opinion that E occurred. Despite the source's conclusion that "event E definitely occurred", and because of many sources of uncertainty, we should consider that testimonial evidence based on opinion is a type of equivocal testimonial evidence (see Figure 38).

There are two other types of equivocal testimonial evidence. The first we call completely equivocal testimonial evidence. Asked whether event E occurred or did not, our source says, "I don't know", or "I can't remember." This is an interesting response of the sort we so frequently observe during congressional hearings. There are two possible explanations for this complete equivocation. The first is that the source is honestly impeaching or undermining his own credibility; he does not know or cannot remember. A frequent addition to this equivocation might be the further statement, "I'm not a good source here, perhaps you ought to ask X (another possible source)." Unfortunately, there is another possible explanation for this complete equivocation; the source does know or can remember, but refuses to tell us whether E occurred or not. If we had evidence that this source did know or did remember whether or not event E occurred, this would be evidence that our source is a double and has more than one employer.

But there is another way a source of HUMINT can equivocate; the source can provide probabilistically equivocal testimonial evidence in various ways. One way is numerical, as in the following example. Asked whether event E occurred or did not, the source might say, "I'm 60 percent sure that event E happened and 40 percent sure that it didn't happen." We could look upon this particular probabilistic equivocation as an assessment by the source of his or her own observational sensitivity. However, if we had evidence pointing to his underassessment of how sure he was that event occurred, we might be inclined to view this evidence as bearing on the source's veracity. Another way probabilistic equivocation can be expressed is in words rather than in numbers. Asked whether event E occurred, the source might say such things as, "I'm fairly sure that E occurred;" "It is quite probable that E occurred;" or "It is very unlikely that E occurred." Here again we would wish to determine whether the source's stated degree of equivocation was legitimate or not.

Example 8. Consider the evidence item E5 Ralph in Table 2. Here Ralph hedges a bit by saying that the lock on the hazardous materials storage area appears to have been forced. He cannot say for sure that the lock had been forced, so he hedges in what he tells us.

3.3.4 Mixed Evidence

There are also other situations in which individual items can reveal various mixtures of the types of evidence. One example involves a tangible document containing a testimonial assertion based on other alleged tangible evidence. As we noted, these forms of evidence are not mutually exclusive; they can occur together in a single item of evidence. We might say that mixtures of them get crowded into the same item of evidence.

Example 9. Here is an obvious example of a mixture of two or more items of tangible evidence; it is called a passport. A passport is a tangible document alleging the existence of other tangible documents recording the place of birth and country of origin of the holder of the passport. In other words, a passport sets up a paper trail certifying the identity of the holder of the passport. In addition to the authenticity of the passport itself, we are also interested in the authenticity of all the other tangible documents on which this passport is based.

3.3.5 Missing Evidence

To say that evidence is missing entails that we must have had some basis for expecting we could obtain it. There are some important sources of uncertainty as far as missing evidence is concerned. In certain situations, missing evidence can itself be evidence. To begin with, consider some form of tangible evidence, such as a document, that we have been unable to obtain. There are several reasons for our inability to find it, some of which are more important than others. First, it is possible that this tangible item never existed in the first place; our expectation that it existed was wrong. Second, the tangible item exists but we have simply been looking in the wrong places for it. Third, the tangible item existed at one time but has been destroyed or misplaced. Fourth, the tangible item exists but someone is keeping it from us. This fourth consideration has some very important inferential implications including denial and possibly deception. An adverse inference can be drawn from someone's failure to produce evidence. The failure to produce requested evidence may mean that producing it would not be in the best interests of the person(s) from whom the item was requested. If these interests coincide with those of an adversary, we could conclude that this failure to produce evidence is part of an attempt to deceive us since, if we did obtain this evidence, it would be in our best interests and not the best interests of the adversary.

Now consider missing testimonial evidence. Suppose we expect that a HUMINT asset $\mathcal{A}$ could tell us about some event of importance to us. There are several interesting possibilities here. First, $\mathcal{A}$

might never respond to our inquiry; put another way, $\mathcal{A}$ responds to our inquiry with silence. There are different rules that apply in intelligence analysis than those applying in our courts of law. In law, a defendant or a witness can claim protection under the Fifth Amendment to our Constitution. He cannot be compelled to testify, and no adverse inference is allowed by his failure to do so. But there is no such privilege in intelligence analysis. We would be entitled to draw an adverse inference about $\mathcal{A}$'s silence in response to our request for information to which we believe $\mathcal{A}$ has access. Another possibility is that $\mathcal{A}$ acts to impeach his own competence; $\mathcal{A}$ tells us that he never has made any observation of events such as those in which we are presently interested. This may sound like the complete equivocation we discussed above. The difference in this case is that $\mathcal{A}$ gives a particular reason why he does not know whether this event occurred or not; $\mathcal{A}$ says he was never in a position to observe this event or had no access to the information requested of him. If, on evidence, we learned that $\mathcal{A}$ did make an observation or did have access to information about the requested event, we would certainly be entitled to draw an adverse inference concerning $\mathcal{A}$'s behavior and the inferential consequences of his refusing to reveal the information we are seeking from him.

In summary, there are very important uncertainties associated with missing evidence, either tangible or testimonial in nature. But there is one final matter to consider about which most analysts will already know. *We should not confuse negative evidence with missing evidence*. To adopt a common phrase, *"evidence of absence (negative evidence) is not the same as absence of evidence (missing evidence)."* Entirely different conclusions can be drawn from evidence that an event did not occur than can be drawn from our failure to find evidence. We are obliged to ask different questions in these two situations. Missing evidence may be either tangible or testimonial in nature.

3.3.6 Authoritative Records

There is one final category of evidence about which we would never be obliged to assess its credibility. In intelligence analyses, and in many analyses in other contexts, we routinely need information whose credibility we would never be expected to defend. In fact, in certain instances we could never establish the credibility of this information. Tabled information of various sorts such as tide tables, celestial tables, tables of physical or mathematical results (such as probabilities associated with statistical calculations), and many other tables of information we would accept as being believable provided that we used these tables correctly. In some instances, of course, tabled information might contain errors. For example, tables showing the ranges or explosive power of a weapons system under various conditions might have incorrect entries that are discovered and corrected. Many other items of information are accepted facts whose credibility is assumed. For example, an analyst would not be obliged to prove that temperatures in Iraq can be around 120° Fahrenheit in summer months, or that the population of Baghdad is greater than that of Basra.

3.3.7 Postive and Ngative Evidence

Positive evidence, either tangible or testimonial, records the occurrence of some event(s) of interest. Negative evidence records the nonoccurrence of event(s) of interest. As we know, discovering that some event did not happen can be just as valuable as discovering that some event did happen. We have often wondered whether in intelligence analysis, or in other contexts, negative evidence is dismissed as being uninteresting. To the extent to which this may occur is troublesome since this may mean that negative evidence having great weight or force is available but is being overlooked. Many of the Sherlock Holmes stories illustrate how valuable negative evidence can be.

3.4 Recurrent Substance-Blind Combinations of Evidence

We have considered a categorization of individual items of evidence but have also mentioned situations in which individual items can reveal various mixtures of the types of evidence shown in Figure 38 (p. 97). We now consider combinations of two or more individual items of evidence. These combinations are also recurrent and do not involve the substance or content of the evidence.

There are three main classes of evidence combinations: harmonious, dissonant, and redundant, all of which we may encounter in a mass of evidence being considered in an intelligence analysis, or an analysis in any other context.

3.4.1 Harmonious Evidence

Two or more items of evidence are harmonious if they are directionally consistent in the sense that they all point toward, or favor, the same hypothesis or possible conclusion.

There are two basic forms of harmonious evidence. The first is called corroborative evidence. In this combination of evidence we first have two or more sources telling us that the same event has occurred. Suppose both of these sources report that event E has occurred. Directional consistency is apparent here since E is consistent with itself. The sources of corroborative evidence may be any combination of the “INTs” we have mentioned. For example, we may have both IMINT and HUMINT telling us that a certain event has occurred at a location at a certain time. Or, we may have IMINT and COMINT both saying that a certain event occurred. This form of corroboration often, but not always, allows us to have greater confidence that the event in question did occur. In such cases we would say that one source has verified what the other source has told us. The exception involves instances in which we have other evidence suggesting that two or more HUMINT sources collaborated in deciding what to tell us, or that one source influenced or coerced another source to report the same event. As we know, HUMINT sources are frequently not independent; they can interact in ways designed to deceive us.

But there is another way evidence can be corroborative in nature involving items of directly relevant and ancillary evidence. Suppose HUMINT asset $\mathcal{A}$ reports an event we take to be directly relevant in an analysis. Suppose that in assessing $\mathcal{A}$'s credibility we have ancillary evidence that we believe supports an attribute of $\mathcal{A}$'s credibility. Such evidence would be corroborative in the sense that we gain further confidence that the event $\mathcal{A}$ reports did in fact occur. An example would involve information about $\mathcal{A}$'s track record in his previous reports (if he made any). Such ancillary evidence would support $\mathcal{A}$'s veracity or observational sensitivity as far as his present report is concerned. But ancillary evidence can bear upon a human source's competence as well. In assessing asset $\mathcal{A}$'s competence we may have evidence from another asset $\mathcal{B}$ who says that $\mathcal{A}$ could in fact have made the observation $\mathcal{A}$ says he made. We could also verify $\mathcal{A}$'s competence by IMINT showing that $\mathcal{A}$ was at the place where he says he made his observation.

But there is another combination of harmonious evidence that differs from corroborative evidence of the same event; it is called convergent evidence. This combination of evidence involves two or more evidence items that concern *different events which point toward or favor the same hypothesis*. Convergent evidence can involve any of the "INTs." Suppose we have the following situation. We have IMINT evidence that event E occurred, and we have MASINT evidence that event F occurred. But we believe that both of the events E and F would point to or favor the same hypothesis H. In other words, these two events are directionally consistent; they both point us in the same inferential direction. But convergent evidence can have an additional and most important property that we will now explain.

Convergent evidence can exhibit what is called evidential synergism. In many situations, two or more evidence items, considered jointly, have greater inferential force or weight than they would have if considered separately or independently. Another equivalent way to characterize evidential synergism is to say that one item of evidence can have greater force if we consider it in light of other evidence we have. Suppose again that we have evidence about events E and F that converge in favoring hypothesis H. But when taken together, or considered jointly, these two events have additional force favoring hypothesis H. Additionally, we might observe that evidence about event F seems to have more force or weight when we consider it in light of evidence about event E. As we have mentioned before, one tragic example of our failure to exploit evidential synergism involves events that occurred before September 11, 2001. The FBI had evidence of persons from the Middle East who arrived at flying schools here in the U.S.A. paying in cash for their flying lessons. But these students wished only to learn how to steer and navigate heavy aircraft, and not how to make takeoffs and landings in these aircraft. At the same time, our intelligence services had evidence that new attacks would be made on the World Trade

Center in New York, this time using airliners. Unfortunately, for various reasons, these items were never considered jointly and hypotheses suggested by these joint events were never considered.

What it comes to is that important evidential synergisms will never be recognized unless intelligence evidence is shared among all agencies involved in its collection and analysis. We have written more on the probabilistic underpinnings of evidential synergism (Schum, 1994/2001, pp. 401-409; Anderson et al., 2005, pp.46-50).

Recall our saying the harmonious evidence is directionally consistent because all of it points to the same conclusion. Here are some examples and questions.

In the cesium-137 scenario we have examples of both forms of evidential harmony. The first involves what we called corroborative evidence, in which two or more sources report the same event. Have a look at evidence E1 Washington Gazette and E4 Not checked-out in Table 2. Here we have both Willard and Ralph telling us that the cesium-137 canister was missing from the STEMQ company warehouse. Ralph's report corroborates Willard's initial report.

Convergent evidence involves evidence about different events, all of which point to the same conclusion. Evidential synergism was illustrated by the following example.

Example 10. Person Y has been under surveillance in connection with terrorist activities. We suspect that Y will attempt to leave the country in a short while. Three days ago we received information that Y sold his car. Today, we received information that he closed his account at his bank. Either item of evidence does not tell us much. He could be planning to buy a new car. He could also be dissatisfied with his bank. But, taken together, these two items of evidence are suggestive that Y is planning to leave the country.

3.4.2 Dissonant Evidence

Dissonant evidence involves combinations of two or more items that are directionally inconsistent; they can point us in different inferential directions or toward different hypotheses.

There are two basic forms of evidential dissonance; the first involves contradictory evidence. Contradictory evidence always involves events that are mutually exclusive (they cannot have occurred jointly). From one source we learn that event E occurred; but from another source we learn that this same event did not occur. The dissonance seems obvious in this case since event E cannot have both occurred and not have occurred. Contradictory evidence can involve any sources of intelligence evidence and any number of sources. For example, we may have some five sources, three telling us that event E occurred and two telling us that event E did not occur. We must first be a bit careful in discussing the directional inconsistency of contradictory evidence. Suppose we are considering whether hypothesis H is true; an obvious alternative is the hypothesis ¬H (H is not true). We further believe that event E, if it occurred, would favor hypothesis H. With some views of probability, which we discuss in Chapter 5, this means that hypothesis ¬H would be favored by event E not occurring. However, with another view of probability we will consider, we may not always believe it necessary to say that the nonoccurrence of E favors the nonoccurrence of H. On other occasions we may not even be sure what the nonoccurrence of E is telling us.

In any case, evidential contradictions are always resolved on credibility grounds. There is quite an interesting history concerning how we have come to rely on this form of resolution. As an example, suppose we have three HUMINT sources who tell us that event E occurred, and one HUMINT source who tells us that event E did not occur. In the not so distant past, it was believed that we should always resolve the contradiction by counting heads; i.e. majority rules. So, on this basis we would side with the three sources who tell us that event E did occur. This reliance on the number of witnesses on either side of a contradiction has a biblical origin. As the Bible records (Deuteronomy, 19:15): "…for any iniquity, at the mouth of two witnesses, or at the mouth of three witnesses shall the matter be decided." This numerical strategy lasted at least till the time of Napoleon. As Wigmore records (Wigmore, 1940, p.256), Napoleon was disturbed by the strategy, saying, "Thus one honorable man by his testimony could not prove a single rascal guilty, though two rascals by their testimony could always prove an honorable man guilty." The trouble here is that counting heads assumes that all of the four sources involved in this episode of contradictory evidence have equal credibility. This may be a very bad assumption since, on further evidence about these four sources, we may well believe that the one source telling us that E did not occur has greater credibility than does the aggregate credibility of the three sources who tell us that event E did occur. As Wigmore also observed, our courts do not accept a majority rule interpretation. What matters is the aggregate credibility of the witnesses on either side. This happens to be entirely consistent with what a Bayesian analysis tells us (Schum, 1994/2001a, pp. 409-412). *So, what matters in resolving evidential contradictions is the aggregate credibility of the sources on either side of this contradiction.*

In some accounts we have read, dissonant evidence is described as being necessarily contradictory in nature; but this is quite erroneous. There is another quite different form of dissonant evidence called <u>divergent evidence</u>. Contradictory evidence involves whether one event occurred or did not occur. But divergent evidence involves entirely different events. The directional inconsistency here means that these events point us toward different hypotheses. In one case, suppose believable evidence about event E would favor hypothesis H, but believable evidence about event F would favor hypothesis ¬H. In a more general case, suppose an analyst is considering four hypotheses {H_1, H_2, H_3, and H_4}. One body of evidence is consistent in pointing most strongly to hypothesis H_1, while another body of evidence is consistent in pointing most strongly toward H_3.

Resolving evidential divergences, or evidential conflicts as it is sometimes called, is a more difficult matter than resolving evidential contradictions. Credibility assessment does play an important role in both cases, but there is an additional difficulty with divergent or conflicting evidence. Suppose we return to the simple situation in which we have believable evidence items regarding events E and F. Two analysts agree that evidence of E favors hypothesis H but the evidence of F favors hypothesis ¬H. A third analyst observes the two analysts who have just agreed that these two evidence items are divergent or conflicting. This third analyst says, "I don't agree with your assessment of this evidence. The trouble is that you are considering these two evidence items separately. If you consider them together the apparent conflict disappears. The reason is that the occurrence of event E would effectively swamp the occurrence of event F and so there is no conflict here. These two items of evidence taken together makes their dissonance disappear." So, dissonance involving divergent evidence always calls for an analyst's judgments on the directionality of evidence about different events. It is often the case that if we knew more about the situation in which different events have occurred we might be able to explain away divergences or conflicts that seem to appear.

We discussed two forms of dissonant, or directionally inconsistent, evidence that point us toward different hypotheses or possible conclusions. Such evidence can be either contradictory or divergent.

3.4.3 Redundant Evidence

Redundant evidence involves combinations of two or more items that either says the same thing over again or does not add anything to what we already have.

This final recurrent and substance-blind combination of evidence is in effect the opposite of the possible evidential synergism mentioned above for convergent evidence. We often encounter two or more items of evidence in which the first item acts to reduce the force of subsequent items of evidence. Stated another way, the first item acts to make subsequent items redundant to some degree. There are two ways this can happen as we will see.

The first form of evidential redundance involves the corroborative evidence we discussed above in Section 3.4.1. In this case we have repeated evidence of the same events. Although having corroborative evidence does add to our confidence that an event of interest did occur, each additional item adds less and less to our confidence. At some point we will surely say, "We already believe that this event occurred, we don't need any further evidence about this event." We refer to this situation as corroborative redundance. The credibility of our sources plays a crucial role in determining how redundant successive reports of the same event will be. To illustrate, suppose event E favors hypothesis H and we have successive items of evidence E_1^*, E_2^*, and E_3^*, all telling us that event E occurred. First, suppose we believe that our first source of evidence, the one who provided E_1^*, is perfectly credible. In short, we know for sure that event E occurred. In this case, having evidence E_2^* and E_3^* would tell us nothing we do not already know, so they are perfectly redundant. But now suppose that the first source is not perfectly credible. In this case E_2^* can add to the verification that event E occurred, depending on how credible we believe the second source to be. To the extent that the second source is not perfectly credible, the third source can add some additional verification.

The second form of redundancy involves different events in which evidence about one event, if credible, takes something away from the inferential force of evidence about another event. We have called this cumulative redundance. The word "cumulative" is an expression used in law to refer to evidence that does not add anything to what we already know. An example of cumulative redundance is the following one.

Example 11. Suppose asset $\mathcal{A}$ tells us that it was Omar, the terrorist, who he saw two days ago planting the shaped explosive device that killed two American soldiers outside of Tikrit. Then asset $\mathcal{B}$ tells us that he saw someone who looked like Omar planting this device two days ago at this same location outside of Tikrit. We have two different events here; the first source says it was Omar, but the second source only says it was someone who looked like Omar. Suppose we believe that the first source is perfectly credible. Then the report from the second source is completely redundant. If Omar was there, it follows necessarily that someone who *looked like* Omar was there. The report of the second source only springs to life if the first source is not credible.

The importance of considering these two forms of evidential redundancy cannot be overstated. In the case of corroborative redundance we risk double counting evidence about the same event and ascribing additional weight to the evidence which it does not have. In the case of cumulative redundance we risk getting more inferential mileage out of the evidence than can be justified.

We mentioned two forms or redundant evidence which we called corroborative and cumulative redundance. Corroborative redundance involves repeated evidence of the same event; cumulative redundance involves evidence about different events. In either case we have instances in which some evidence can reduce the inferential force of other evidence. Here are some examples and questions concerning redundance:

Example 12. Look again at Willard's and Ralph's reports in evidence items E1 Washington Gazette and E4 Not checked-out (see Table 2). If we believed that Willard was completely believable in his report that the cesium-137 canister was missing, then Ralph's report would tell us nothing we do not already know for sure. This would make Ralph's report completely redundant. Ralph's report has value to the extent that Willard is not completely believable. But we could have asked other persons if they believed the cesium-137 canister to be missing. Here comes Joe and then

Frank who each tell us that the canister was missing. Each new report of the same event tells us less and less depending on the credibility of earlier reports of the same event. This is what corroborative redundance is all about.

3.4.4 Why Considering Evidence Combinations Is Important

There may be very few, if any, situations in which the conclusion of an intelligence analysis problem is based on just a single item of evidence. Intelligence analysts usually draw conclusions based on masses of evidence of different kinds and coming from a variety of different sources. What is obvious is that careful assessment of the joint impact or force of masses of evidence is crucial in drawing conclusions about what is happening or what will happen in a situation of interest. Determining the joint impact or force of a mass of evidence is no easy matter. The basic reason is that individual evidence items can have a variety of different effects on each other. Another way of saying this is to say that our items of evidence can interact in different ways; this is what makes assessments of joint impact or force so complex. It also means that we must consider combinations of evidence to see how they might interact. We have presented several

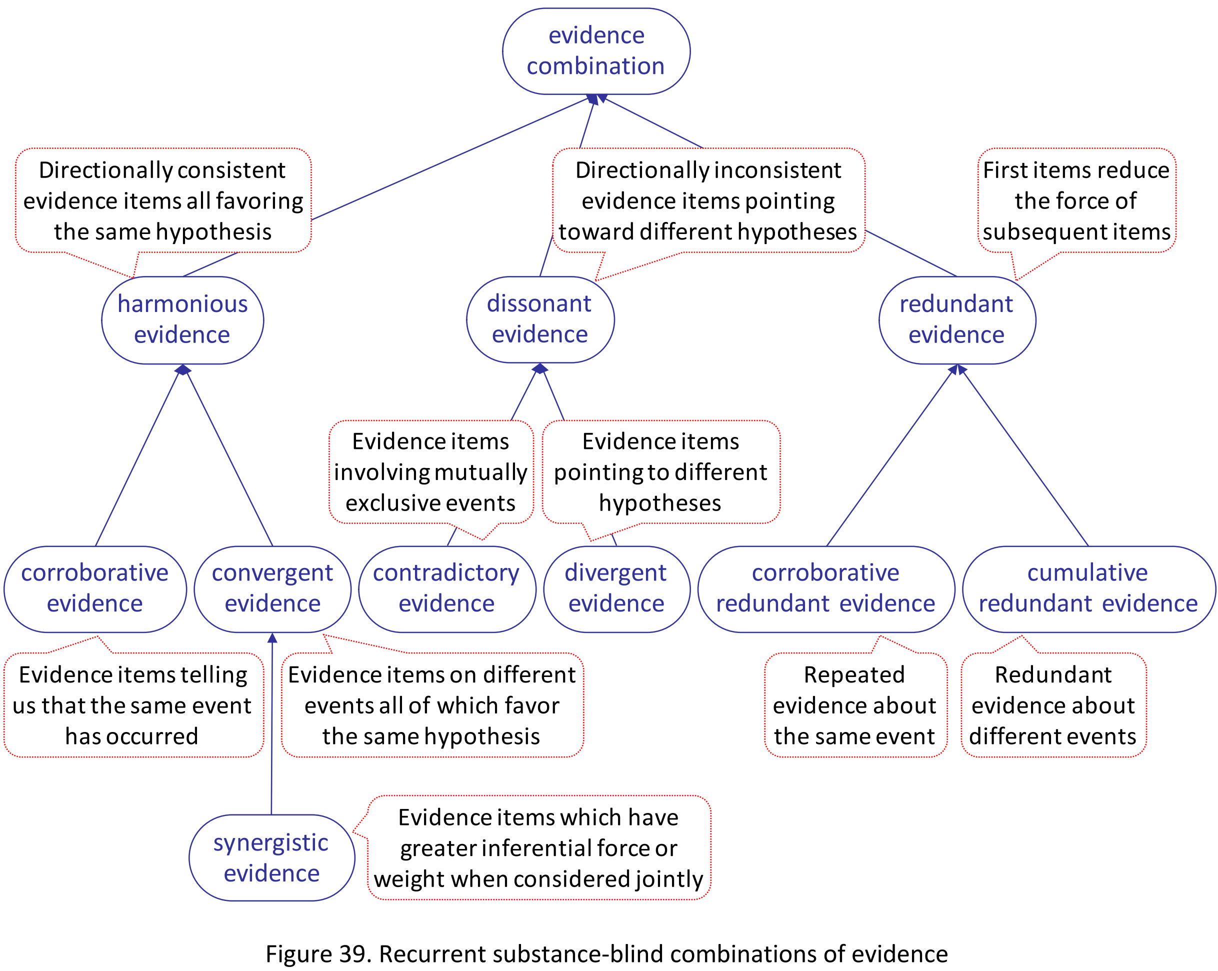
evidence combination
Directionally consistent evidence items all favoring the same hypothesis
Directionally inconsistent evidence items pointing toward different hypotheses
First items reduce the force of subsequent items
harmonious evidence
dissonant evidence
redundant evidence
Evidence items involving mutually exclusive events
Evidence items pointing to different hypotheses
corroborative evidence
convergent evidence
contradictory evidence
divergent evidence
corroborative redundant evidence
cumulative redundant evidence
Evidence items telling us that the same event has occurred
Evidence items on different events all of which favor the same hypothesis
Repeated evidence about the same event
Redundant evidence about different events
synergistic evidence
Evidence items which have greater inferential force or weight when considered jointly

Figure 39. Recurrent substance-blind combinations of evidence

basic patterns of interactions among evidence items when we consider possible combinations of evidence items. They are summarized in Figure 39.

One reason for carefully considering these combinations of evidence is that they are often confused or incorrectly identified, leading to mistakes in how the evidence is described in an analysis. But perhaps the most important reason is that *there are very important sources of uncertainty lurking in these evidential combinations*. But we have one more task to perform regarding uncertainty and evidence. There are five characteristics of evidence that make conclusions drawn from it necessarily probabilistic in nature. They are described in detail later in Section 3.5.

3.5 Major Sources of Uncertainty in Masses of Evidence

There are five major reasons why conclusions reached in intelligence analysis based on evidence will be necessarily probabilistic in nature: our evidence is always incomplete, usually inconclusive, frequently ambiguous, commonly dissonant, and with various degrees of credibility. Any one of these reasons can lead to uncertain conclusions, but an analyst drawing conclusions based on masses of different forms and combinations of evidence will likely encounter all five of these reasons at the same time. Intelligence analysts who report conclusions with varying degrees of uncertainty are often unjustly criticized for doing so. Persons providing such criticisms are doubtless unaware of many or most of the reasons why it is necessary to hedge conclusions in probabilistic terms. *One main reason we have for providing a careful account of these five reasons is that no single view of probability we know about captures well all five of these sources of uncertainty.* Each view of probability we will mention provides useful insights about some of these sources of uncertainty, but no single view says all there is to be said.

3.5.1 Incompleteness

We may have all heard someone say, “I am going to wait until I have all the evidence before I draw a conclusion or make a decision.” This person faces an infinitely long wait because there is no situation in which we can say we have all the possible evidence. The first way of showing that this statement is true is to consider the distinction we have made between directly relevant and ancillary evidence. In doing so, we also referred to ancillary evidence as being meta-evidence, or evidence about evidence. The trouble here is that we face an infinite regress in which we have evidence, ancillary evidence about this evidence, ancillary evidence about this ancillary evidence, and so on, ad infinitum. Suppose we have an HUMINT asset who provides us with some interesting evidence; call this asset our primary source. But then we have a secondary source who provides us with ancillary evidence about the credibility of our primary source. But then we have a tertiary source who provides ancillary evidence about the credibility of our secondary source. This process could go on and on indefinitely. This fact was noted years ago by the CIA’s

James J. Angleton, who encountered situations in which chains of HUMINT sources provided contradictory and divergent evidence about each other's credibility. He described this situation as similar to being in a "wilderness of mirrors" (Martin, 1980).

There are many other situations in which we could have endless chains of evidence, and evidence about evidence, even in empirical statistical situations. It is often said that the conclusions reached by statisticians are always misleading to some extent and that they must choose ways of reporting conclusions that are minimally misleading. Here is an intelligence analyst who reports the results of a statistical analysis of the capabilities of a weapon system of some sort. Her conclusions are challenged by another analyst. In turn, the comments made by this challenger are then challenged by a third analyst. This process could go on forever until someone decides to call a halt to this evidence about evidence situation. This fact is noted in our procedures for trials at law that limit the extent to which we can follow this chain. If this were not the case, a trial might go on for years without any verdict ever being reached.

But there is another even more important reason for our evidence always being incomplete. In how many intelligence analyses could it ever be said that every one of the relevant questions that could have been asked was in fact answered by the evidence that was gathered and analyzed? There probably has never been an intelligence analysis in which there were no lingering unanswered questions at the time a conclusion was required. In the absence of clairvoyance, there may be considerable uncertainty about what questions should be asked in an intelligence analysis. In Section 3.3.5 we discussed missing evidence as a possible category of evidence. In such cases we attempted to answer certain questions but were unable to do so. But there will be many questions lingering that we have never even attempted to answer as well as many questions that we may not even recognize as being relevant. In Section 5.6 we will discuss a view of probability, called the <u>Baconian</u> View, which uniquely places special emphasis on the extent to which our evidence in an analysis is complete in its coverage of all the questions we recognize as being relevant to the conclusions we must reach. This view requires us to consider the force of evidence we do have, but it also says that this force depends on questions that are unanswered by the present evidence we are considering. This issue of completeness is bound to be of interest to intelligence analysts engaged in <u>current intelligence</u> in which conclusions are often required in a very short time. An issue here concerns the extent to which any analyst could have the time to cover all the questions that might occur to this person as being relevant to the conclusion being requested.

Example 13. On various matters concerning events in the past, it might be argued that we have complete evidence. For example, we believe we have evidence that allows us <u>now</u> to conclude, <u>beyond all shadow of doubt</u>, that the twin towers of the World Trade Center in New York City were destroyed on September 11, 2001, and that the then President of the United States, John

F. Kennedy, was assassinated in Dallas, Texas, on November 22, 1963. We can go to New York and see for ourselves where the World Trade Center used to be and then view television images of the two aircraft that slammed into the buildings, and the horror of the subsequent collapse of the buildings. Or, we can view the monument to John F. Kennedy in Arlington National Cemetery in Virginia that records the place and time of his death. We can also watch the so-called Zapruder film that shows President Kennedy being struck in the head and then collapsing into his wife Jacqueline's arms after he was struck. In each of these cases we have what can be regarding as conclusive evidence (more on conclusive evidence in a moment).

There are two points to be made about the two examples we have just provided. First, it does appear that the occurrence of these two events is no longer a topic of analytic interest, since we are certain that they occurred. But there are very many other questions concerning these two events that are now of great interest and will continue to be of interest in future analyses. Concerning the tragedy in New York, we are asking such questions as, "Who were all the persons involved in the planning of this action?" and "Why did these persons take this action at this particular time?" As far as the Kennedy assassination is concerned, it is still being asked, "Did Lee Harvey Oswald act by himself or was he part of a conspiracy to kill President Kennedy?" The evidence on these questions will never be complete. The second point concerns the stability over time of the conclusions we have now reached with certainty. Five hundred years from now, will these two conclusions about the destruction of the World Trade Center and the assassination of President Kennedy still be regarded as certain (if they are, indeed, matters of concern at all)? Perhaps the evidence we now regard as conclusive will have vanished long before the year 2514.

Example 14. Intelligence analysts will not normally be concerned about what will be inferred five hundred years from now about events of interest today. In many cases they are asked to reach conclusions concerning past events based on evidence that is far from complete. For example, new evidence is coming to light as we write these words on whether Lee Harvey Oswald acted alone in the assassination of President Kennedy. New accounts of the identities and motivations of the nineteen terrorists who destroyed the World Trade Center appear regularly as well. Even more evidence about these events will certainly emerge in the future. Now, consider an intelligence analysis that involves the prediction of some future event, such as whether countries such as Iran, Iraq, Saudi Arabia, Syria, Egypt, and Turkey will engage in all-out war if we remove all of our forces from countries in the Middle East. Unless we have a person who is certifiably clairvoyant and can see into the future we will never know for certain. The only thing certain is that we can only draw inferences about events such as these based on evidence about events in the past and the ever-fleeing present. This evidence will obviously be incomplete, if only because we may not be asking the right questions. In addition, and of the greatest importance, there will always be questions unanswered by the evidence we do have. These is only one view of probabilistic reasoning, the Baconian view, that asks how complete is our coverage of evidence and how many questions remain unanswered by the evidence we have considered.

We leave the issue of incompleteness with two thoughts we ask you to keep in mind. With the exception of situations like the ones given in Example 13 concerning past events and conclusive evidence, in any other case we have two reasons why our evidence is never complete:
We will always have some unanswered questions;
We will always encounter the need for meta-evidence or evidence about evidence, there being no end to this need.

3.5.2 Inconclusiveness

The evidence encountered in intelligence analysis is commonly inconclusive in nature. This means that evidence is consistent with more than one possibility, hypothesis, or explanation. Conclusions reached from such evidence can only be probabilistic in nature. Another term we might use here is to say that intelligence analysis usually involves circumstantial evidence. Recall that circumstantial evidence, even if credible, supplies only some but not complete grounds for a conclusion. Conclusive evidence, that which is consistent with only one possibility or hypothesis, is usually in very short supply. Conclusive evidence would supply complete grounds for, or make necessary or certain, some hypothesis or possible conclusion.

There is an expression used by intelligence analysts with reference to conclusive evidence on some major hypothesis that comes from a completely credible source; the term nugget is used with reference to such evidence. Here is an example of such a nugget using our dirty bomb (cesium-137) example. Suppose we have a trusted source who reports the following events. This source tells us, “Persons associated with the North American Jihadist Organization (NAJO) in Silver Spring, MD, did acquire the cesium-137 that was stolen from the STEMQ warehouse in Baltimore. They are now constructing a dirty bomb in the garage of a residence at 221 Colesville Rd. in Silver Spring, MD, which they intend to set off on the grounds of the capitol building in Washington, DC, next Thursday at 12:00PM.” If we had such a nugget, we could easily prevent this disaster from occurring. Barring the acquisition of such a nugget we must, as the phrase goes, “mine lots of lower grade ore” in the form of inconclusive and circumstantial evidence.

3.5.3 Ambiguity

Evidence is ambiguous to the extent to which we cannot determine what it is telling us. We might also describe ambiguous evidence as being imprecisely stated. In addition, ambiguous evidence goes beyond being merely inconclusive. We may have precisely stated, non-ambiguous, evidence that is still circumstantial or consistent with more than one hypothesis or possible conclusion. As an example of ambiguity, we will discuss a conversation we have intercepted between two persons of interest who are involved in known terrorist activities. In this conversation they make use of words used to describe persons and their activities that are designed to conceal their intentions or to mislead others who may be listening. It is well known, for example, that terrorist organizations in the Middle East spend a great amount of time training their operatives to be skillful in disguising their identities, capabilities, and intentions. Even our best efforts to disambiguate evidence we receive are not always successful. A person listening to such a conversation may be perfectly fluent in the language and dialect being used by the two persons overheard, but still be unable to tell us exactly what these two persons were saying.

Example 15. In many cases we may encounter information that would seem relevant if we could decide what it is telling us; in other words, it is ambiguous, imprecise, or vague. We may, for example, be provided with a document concerning apparently current plans of a certain terrorist organization. This document, on careful examination, is reckoned to be authentic. The document says, "Our destiny is now approaching. We will meet at the usual place at the agreed upon time and proceed to the Crusader site. If Allah is willing, we will kill many of them." This message is certainly vague with respect to the time and place of what seems to be an intended terrorist operation. With any luck, we might have other evidence about this group that might allow us to remove at least some of the ambiguity apparent in this document. The basic trouble with ambiguous evidence is that it certainly generates uncertainty but we are often at a loss about how to remove at least some of this uncertainty.

Example 16. In many cases involving intelligence reports you receive, you cannot tell exactly what information is being conveyed in these reports. Very good examples are provided by the equivocal testimonial evidence. Here is a source of HUMINT who hedges when asked if event E occurred. The source says, “I am not sure, but I believe it is about 60% likely that event E happened.” The trouble is that “60%” for this source might not be the same as what you believe “60%” means. Evidence is frequently encountered that is imprecise in some way with the result that you cannot tell exactly what event is being reported. If you can’t be sure what your evidence says you can hardly avoid hedging your conclusions.

Example 17. In so many instances of HUMINT reports, the sources of it provide ambiguous accounts of observed events. Here is a source who reports seeing, “A large number of Taliban fighters assembling, a short time ago, near location X in Afghanistan.” This is an imprecise or fuzzy account of this situation. Hearing this account, we do not know exactly how many, where, and when this assemblage of Taliban fighters occurred. Was the exact number 20, 50, 100, 500, or larger? Was the time 12 hours ago, 3 hours, 30 minutes, or shorter? Was the location 10 km, 5 km, 2 km, or only 500 m from location X? Why do such ambiguous reports so often occur? The reason is quite simple. Human observers cannot make precise judgments under many conditions and are trying to do their best in honestly reporting what they observed. Suppose instead, this source had given the following report: “At exactly two hours and 17 minutes ago, I observed exactly 257 Taliban fighters assembling exactly 0.76 km from location X.” We might be quite suspicious of a source who provided such a precise account, especially when we hear about the conditions under which the source said he made this observation.

Example 18. Often analysts are criticized for providing ambiguous conclusions in their analyses. For example, here is an intelligence analyst who says her analysis shows that it is quite probable that the Russians are increasing the number of their air defense forces in Damascus in Syria. She is criticized for not providing a precise probability to go along with her conclusion. She says, "OK you want me to give you a precise numerical probability here and so I will say that the probability is exactly 0.73 that the Russians are increasing the number of their air defense forces in Damascus in Syria." After she provides this number, she tells a colleague: "I have just responded to a stupid request. I was asked to give a precise probability to the Russians increasing the number of their air defense forces in Damascus in Syria. I said this probability was 0.73. But I certainly did not base this number on 100 past occasions on which the Russians increased the number of their air defense forces 73 times. No one has such a record. The number I gave is just a good guess about what, for me, 'quite probable' means. But this number might actually be any number between 0.65 and 0.85."

3.5.4 Dissonance

We devoted Section 3.4.2 to a discussion of dissonant evidence as being a recurrent combination of intelligence evidence. We described dissonant evidence as being directionally inconsistent in the sense that it points us toward more than one hypothesis or possible conclusion. We described the two forms of dissonance as involving contradictory and divergent evidence. Suppose we have an analysis in which it is stated that *all* of the evidence analyzed was harmonious or directionally consistent in favoring the conclusion reported in the analysis. The first question someone should ask is, "Are you sure you did not gather and analyze just the evidence you believe would favor the conclusion you reported?" In some situations this is called "cherry picking." At least some pattern of dissonance may be expected in any intelligence analysis, especially one that involves any degree of complexity. To be sure, some of the evidence may be harmonious in pointing in one direction, but on balance we will have other harmonious evidence pointing in another direction. In short, some dissonance will be expected in every intelligence analysis having any degree of complexity. Failure to report this dissonance will arouse justified suspicions of the intentions and the competence of the persons performing the analysis.

3.5.5 Imperfect Credibility

Sources of evidence of any kind (i.e., all of the possible "INTs") have any possible degree of credibility. No mechanical or electronic sensor is perfectly accurate; and on occasion sensor records of various sorts can even be faked. As we all know, no human observer is always perfectly credible. In some cases, a person may not even be a competent source of information. Here are some examples of these distinctions. Radars and other sensory images are not always sensitive enough to discriminate between events of interest to us. On occasion we can obtain estimates of a sensor's hit rate and false-positive rate. It is certainly not unheard of in the intelligence business for documents, photographs, and other similar items of evidence to be faked or forged. At least some items of evidence you may encounter will not be authentic; they are not what they seem to be. Human sources are not always truthful, objective, or accurate as observers. Further, a person may be truthful, objective, or accurate about some matters but not about others. Just because a human source has given us what we take to be believable evidence in the past is no guarantee that he will continue to do so in future. Failure to recognize this simple fact has produced more than one intelligence-related catastrophe.

An experienced analyst reading the account of uncertainty in intelligence evidence just provided may easily recognize all five of these characteristics, and will have encountered most of them at one time or another. Any critic of intelligence analysts for expressing their uncertainty in reporting conclusions should be made aware of the five reasons we have mentioned that require intelligence analysts to acknowledge the extent of their uncertainty, which will always be present. But there are additional reasons why uncertainty is evident in all complex activities such as intelligence analysis.

We suspect that one reason for the current interest in uncertainty in intelligence analysis is that analysts have encountered sources of uncertainty that are not captured by conventional views of probability. This is the main reason we have for presenting alternative systems of probabilistic reasoning in the next section.

3.6 Review Questions

37. Evidence, especially testimonial evidence (e.g., HUMINT), often relates the occurrence of several events. For example, here is an item of evidence coming from a human source where the source tells us several things. The source says: "I observed person P in company with a known Al-Qaeda operative in Vienna, Austria on August 21, 2002. During their conversation, I observed the Al-Qaeda operative taking a stack of $100 bills and a document from Person P. The document looked like a flight manual." This report, carefully parsed, contains several events. Can you identify them?

38. Here’s an interesting and important question for you to consider. If you have a weak link in an argument from evidence to some major hypothesis, is it worse to have this weak link at the top or at the bottom of your argument?

39. The third credential of evidence, its force or weight, depends upon our beliefs about the other two credentials: relevance and credibility. Give some examples of this relationship.

You can do this in words and without any equations.

40. We have emphasized the fact that evidence about some event is not the same as knowing that this event actually occurred. Suppose we have some evidence E* that event E occurred. We gave an example involving HUMINT evidence E* from a source named Mouse that event E occurred, where E is the event that Amad M. attended an al Qaeda weapons training class near Madyan in Northwest Pakistan in October, 2013. Here we had the task of inferring E based on evidence E*. Can you think of another example in which we infer an event E, based on evidence E*?

41. Indicate and justify what type of evidence is each of the following items:
    (a) A spent shell casing.
    (b) Human source X reports to us that military coup is to be expected in Country A within the next two weeks.
    (c) A captured document.
    (d) You take your car for an oil change, expecting the bill to be about $25. Instead, the bill is $350. You ask the mechanic why an oil change costs so much. The mechanic tells you that you needed a new fuel pump and a new water pump, which he changed in the interests of your safety. You ask the mechanic to let you see these two pumps which you believed were working perfectly. The mechanic tells you how sorry he is that these two items have gone missing.
    (e) Human source Y reports to us that the morale among combat troops in Country B is at an all-time low.
    (f) A sensor image (radar, IR, photo) of some ground installations in a certain territory.
    (g) A table showing the reliability of a certain system after various numbers of hours of operation.

42. In discussing the relevance of evidence we noted that this credential of evidence answers the question: So what? How is this evidence linked to hypotheses we are trying to prove or disprove? Consider evidence E* and event E (that M did rent the car from Quick car rental company on 24 November) in the answer to Question 12. From E we infer F that M drove the rented car the next day, 25 November.) Then consider the hypothesis "H: Person P was acting on behalf of Al Qaeda in this car bombing incident." How would you defend the relevance of this evidence E* on hypothesis H?

43. Consider our answer to Question 39 in which we proposed a chain of reasoning between evidence E* and hypothesis H. The links we considered in this relevance argument consisted of the events E, A, B, C, and H. All these links are sources of doubt or uncertainty about these links. In other words, any of these events might not be true. Provide some reasons why these events might not be true.

44. There is no such thing as a perfect argument or one that is absolutely correct and complete. For example, someone can always find one or more missing links in an argument that should be considered. Can you find any missing links in our argument shown in the answer to Question 40?

45. Give some examples from your own experience when you have heard people providing information about which they hedge or equivocate.

46. Why are credibility questions different for different forms of evidence and its sources?

47. The third credential of evidence, its force or weight, depends upon our beliefs about the other two credentials: relevance and believability. Give some examples of this relationship. You can do this in words and without any equations.

48. In discussing the relevance of evidence we noted that this credential of evidence answers the question: So what? How is this evidence linked to hypotheses we are trying to prove or disprove? Consider evidence E* and event E (that M did rent the car from Quick car rental company on 24 November) in the answer to Question 12. From E we infer F that M drove the rented car the next day, 25 November.) Then consider the hypothesis "H: Person P was acting on behalf of Al Qaeda in this car bombing incident." How would you defend the relevance of this evidence E* on hypothesis H?

49. Consider our answer to Question 39 in which we proposed a chain of reasoning between evidence E* and hypothesis H. The links we considered in this relevance argument consisted of the events E, A, B, C, and H. All these links are sources of doubt or uncertainty about these links. In other words, any of these events might not be true. Provide some reasons why these events might not be true.

50. There is no such thing as a perfect argument or one that is absolutely correct and complete. For example, someone can always find one or more missing links in an argument that should be considered. Can you find any missing links in our argument shown in the answer to Question 40?

51. Give some examples from your own experience when you have heard people providing information about which they hedge or equivocate.

52. Can you provide other examples of mixtures of evidence from your own experience?

53. What inferences might we draw from Omar al-Massari's refusal to provide us with his laptop computer?

54. Can you make up some examples of evidence that corroborates other evidence in our dirty

bomb scenario? Ask yourself what items of evidence we now have that you would like to see corroborated.

55. You have two sources reporting the current location of a certain tank column. One source says it is five miles away to the north; the other says it is three miles away to the east. How would you characterize these two items of evidence?

56. Two human sources each report observing person X in company yesterday with a known distributor of narcotics. Is this evidence corroborative or convergent?

57. You have an aerial photograph you believe shows three surface-to-surface missiles of a certain sort at map coordinates (x, y); this photograph was taken one week ago. You also have HUMINT from a source who reports observing three missiles of this sort one week ago at nearly the same coordinates. Is this redundant evidence?

58. One source tells us that T, a known terrorist, was observed at location X at 10:00PM last Friday. Then another source tells us that T was at location Y (100 miles away from X) at this same time. What kind of evidence is this?

59. From one source we receive information that Country A is moving military forces in the direction of Country B; we believe this favors hostile action between countries A and B. But another source tells us of recent secret negotiations between representatives of countries A and B that were successful in resolving major differences between countries A and B. This evidence you believe favors the possibility that there will be no hostile action involving A and B. What kind of evidence is this?

60. In assessing these three forms of evidence combinations (i.e., harmonious, dissonant, and redundant), show why it is so necessary to carefully consider the credibility of sources of the evidence being considered.

61. We naturally encounter instances of harmonious, dissonant, and redundant combinations of evidence in our daily lives. Provide some examples.

62. Here is a problem that involves resolving dissonant patterns of contradictory evidence. Suppose you have six persons who say that event E occurred, and only two persons who say that E did not occur. Can you resolve this contradiction by simply counting heads and siding with the majority?

63. We suspect that person P is a double agent and is presently passing classified information to a certain potential adversary. We thought he worked only for us but now have some grounds for a belief that he is also employed by this adversary. What would constitute absolutely complete evidence that P is in fact a double agent?

64. Give an example of corroborative redundance.

65. One of our military transport aircraft made a stop at a civilian airfield in Country C, with whom we have had friendly relations over the years. Two days ago this aircraft was destroyed on the ground by an explosive device. There are identified groups in Country C that do not favor C's continued friendly associations with us; one of these groups is called the "Purples." Person Q, believed to be associated with the "Purples" was observed, by a usually believable source, in an aircraft parking area of this airfield just one hour before our aircraft was destroyed. Why is this just inconclusive evidence that the "Purples" were involved in this incident?

66. We believe the person shown in the photograph is P. However, the figure is blurred. Why is this evidence ambiguous?

67. A radar image shows the possibility of <u>one or more</u> aircraft at a certain location. Why is this evidence ambiguous?

68. An observer tells us that he saw a tall man with very dark hair driving away in an old car from the sight of a terrorist incident shortly after it occurred. Why is this evidence ambiguous?

69. From your own experience, can you recall other items of ambiguous evidence you have received?

70. You also must consider this person's competence. Provide an example of a source who is not credible because the sourse is not competent.

71. How does intelligence analysis differ from evidential analyses in law trials as far as the completeness of evidence is concerned?

72. Can you think of instances in which you might say you have conclusive evidence when this is actually not correct?

73. What are some causes of evidence to be ambiguous, and how does ambiguity differ from inconclusiveness?

74. Show how evidential dissonance and selectivity are related in ways that can be inferentially hazardous.

75. It can be argued that of all the inferential issues involved in intelligence analysis, the most important and interesting ones involve the <u>credibility</u> of evidence and its sources. Give some reasons why this is so.

76. We suspect that person P is a double agent and is presently passing classified information to

a certain potential adversary. We thought he worked only for us but now have some grounds for a belief that he is also employed by this adversary. What would constitute absolutely complete evidence that P is in fact a double agent?

77. One of our military transport aircraft made a stop at a civilian airfield in Country C, with whom we have had friendly relations over the years. Two days ago this aircraft was destroyed on the ground by an explosive device. There are identified groups in Country C that do not favor C's continued friendly associations with us; one of these groups is called the "Purples." Person Q, believed to be associated with the "Purples" was observed, by a usually believable source, in an aircraft parking area of this airfield just one hour before our aircraft was destroyed. Why is this just inconclusive evidence that the "Purples" were involved in this incident?

78. We believe the person shown in the photograph is P. However, the figure is blurred. Why is this evidence ambiguous?

79. A radar image shows the possibility of <u>one or more</u> aircraft at a certain location. Why is this evidence ambiguous?

80. An observer tells us that he saw a tall man with very dark hair driving away in an old car from the sight of a terrorist incident shortly after it occurred. Why is this evidence ambiguous?

81. From your own experience, can you recall other items of ambiguous evidence you have received?

82. You also must consider this person's competence. Provide an example of a source who is not credible because the source is not competent.

83. How does intelligence analysis differ from evidential analyses in law trials as far as the completeness of evidence is concerned?

84. Can you think of instances in which you might say you have conclusive evidence when this is actually not correct?

85. What are some causes of evidence to be ambiguous, and how does ambiguity differ from inconclusiveness?

86. Show how evidential dissonance and selectivity are related in ways that can be inferentially hazardous.

87. It can be argued that of all the inferential issues involved in intelligence analysis, the most important and interesting ones involve the <u>credibility</u> of evidence and its sources. Give some reasons why this is so.

88. The competence, veracity, objectivity, and observational sensitivity need to be considered for all of the human sources in our dirty bomb example. Pick a human source in this example and state what kinds of questions you would ask about the competence, veracity, objectivity, and observational sensitivity of the person you have chosen.

89. Identify the "dots," details, or "trifles" [as Sherlock Holmes called them] in the following intelligence report:

    FBI Report 1: [1 April, this Year. Abdul R is the owner of a Gourmet Foods shop in City A, in Virginia. [Phone number 703-abc-defg]. BB Union National Bank lists Gourmet Foods as holding account number 10701xxxxxx Six checks totaling $35,000 have been deposited in this account in the past four months and are recorded as having been drawn on accounts at the Pyramid Bank of Cairo, Egypt and the Central Bank of Dubai, United Arab Emirates. Both of these banks have just been listed as possible conduits in money laundering schemes.

90. Identify the "dots," details, or "trifles" in the following intelligence report:

    FBI Report 2: [12 May, this year [From MI-5, UK]. Riyad Yasser, a UK citizen, was arrested on 1 May, this year following an accident on the M4 Motorway near the Heston Service Area outside of London. Yasser has been an air traffic controller at Heathrow Airport for the past six years. Two kilos of Semtex were found in the trunk of his car. A videocassette of a sermon given by Omar Mahmoud Othman, formerly a Salafi jihad preacher at the Baker St. Mosque in London, was found in Yasser's apartment at # 44, Northumberland Circle, East Bedfont, London. Also found in Yasser's apartment was a note containing several addresses in Canada, the USA, and in Nassau in the Bahamas. The addresses are: 7xx St. Clare St., Montreal; 4xx 11th Street, Miami Beach, FL; 17xx Ferry Ave., Camden, NJ, and **xx** Apple St. in Nassau, The Bahamas.

91. What are some causes of evidence to be ambiguous, and how does ambiguity differ from inconclusiveness?

4 RELEVANCE AND CREDIBILITY

4.1 Relevance of Evidence and Arguments

Accurate assessments of the relevance of evidence and arguments in an argumentation structure are critical for accurate hypothesis assessment. Relevance depends on:

- how **recent** is the evidence,
- how **unambiguous** is the evidence (or the argument),
- how **conclusive** is the link between the evidence (or the argument) and the hypothesis,
- how **complete** is the evidence or the argument.

In the following we provide more details about each of these relevance credentials.

4.1.1 Recentness

The information in the evidence may be dated as it relates to the hypothesis, and its relevance, therefore, should be less than certain. The main criterion for judging whether the evidence is dated is how likely it is that the information has changed and thus is no longer accurate. The implied assumption for evidence that is dated is that the evidence is still valid.

If country X has a longstanding hate relationship with country Y that goes back decades, and there is evidence that the Foreign Minister of country X three months ago said there will never be peace between the two countries, the relevance of this information to the hypothesis that prospects for peace are poor remains high, even if the information is three-months old.

If country X has reneged on a promise to stop selling arms to another country three times in the past six months, and evidence from a month ago reports that country X will stop selling arms to that country, the relevance of that evidence is not high.

Each of the examples below shows a hypothesis and one or several items of evidence (that could also be interpreted as sub-hypothees), with a justification of its relevance to the hypothesis.

Example 19. *Hypothesis:* Bogustan is not yet producing chemical warfare agents at the Tanan chemical plant as of mid-November.

Evidence: A collection drone operating near a plant under construction at Tanan did not detect any chemical warfare agents in May.

> A relevance of lacking support would be defensible given the information is rather old; construction could have been completed and production could have started up in the intervening six months before November.

Evidence: A collection drone operating near Tanan did not detect any chemical warfare agents in mid-October.

> A relevance of barely likely would be defensible given the information is more recent; nonetheless, production could have started up in the weeks before mid-November.

Evidence: A collection drone operating near Tanan did not detect any chemical warfare agents in early November.

> A relevance of likely would be defensible given the information is more recent; nonetheless, production could have started up in the days before mid-November.

4.1.2 Ambiguity

The evidence or sub-hypothesis may be ambiguous, and its relevance therefore should be less than certain. Information is ambiguous when it can be interpreted in multiple ways or mean different things. Because of this ambiguity, the evidence or sub-hypothesis, depending on how it is interpreted, can support various hypotheses. The extent of the ambiguity can vary--the more ambiguous the information, the less its relevance.

The ambiguity can involve the description of the event, the individuals or organizations involved in the event, and/or the time frame for the event. References to specific entities are ambiguous when the purpose of the entity is not clear and is assumed based on other evidence.

For example, evidence of contact between a “Department 10”—a very specific reference—and a specific plant that is being used to support a hypothesis that the plant is involved in chemical weapons production is ambiguous if there is uncertainty regarding what Department 10 is and what is its exact connection to chemical weapons production.

Example 20. *Hypothesis:* Shamland’s president believes that Shamland needs nuclear weapons to counter an existential threat from Aggressia.

Evidence: The head of Shamland's national security council told the Defense Minister that the president believed that Aggressia's testing of a nuclear weapon last year would drive Shamland to produce special weapons to offset Aggressia's new superiority.

> A defensible assessment of relevance in this case is barely likely given the ambiguous reference to “special weapons”—almost certainly a reference to weapons of mass destruction that includes nuclear weapons but biological and chemical weapons as well.

Evidence: The head of Shamland's national security council on 25 February told the Defense Minister that the president believed that Aggressia's testing of a nuclear weapon last year would drive Shamland to pursue the nuclear option to offset Aggressia's new superiority.

> A defensible assessment of relevance in this case is very likely given the apparent reference

to nuclear weapons. The reference ("nuclear option"), however, also could include "dirty bombs" that spread radiation (but lack the destructive power of a nuclear weapon) so the relevance cannot be certain.

Evidence: The head of Shamland's national security council on 25 February told the Defense Minister that the president believed that Aggressia's testing of a nuclear weapon last year would drive Shamland to produce nuclear weapons to offset Aggressia's newly acquired nuclear weapons capability.

A defensible assessment of relevance in this case is certain given the specific reference to nuclear weapons.

Example 21. *Hypothesis:* Equipment is operating at the Tanan plant in late November.

Evidence: According to an intercepted communication, an unidentified Ministry of Chemicals official on 23 November told an official at a chemical research institute that "everything at Tanaka is working well."

A defensible assessment of relevance in this case is barely likely given the ambiguous reference to "everything" and "working well." "Everything" in this example conceivably could be a reference to personnel issues.

Evidence: According to an intercepted communication, an unidentified Ministry of Chemicals official on 23 November told an official at a chemical research institute that "equipment at Tanaka is working well."

The relevance of this evidence would increase to likely or more than likely but the ambiguous reference "working well" remains. "Working well" could be a reference to "remains in working order": The equipment is capable of operation but is not operating.

Evidence can be ambiguous when individuals are referred to not by specific name but by a characterization that relates to their relative position in the entity's hierarchy. For example, references to "senior officials" or "high-ranking officials" could be used accurately to describe cabinet-level officials as well as office or department heads who may occupy positions several rungs lower in the entity's bureaucracy. Organizations can similarly be described not by name but by area of responsibility. Uncertainty in assessing the relevance of evidence using such references occurs because multiple organizations may have similar responsibilities.

Example 22. Hypothesis: Bogustan agreed to covertly sell chemical precursors to Aggressia.

Evidence: An official in Aggressia told an Aggressian diplomat stationed in Bogustan that the highest levels of Bogustan's government approved a deal to covertly sell Aggressia chemical precursors.

The relevance of the evidence in this example is very likely: The "highest levels" of the government strongly suggests a state sponsored deal and not unauthorized activity. If only "senior officials" were cited in the evidence, the relevance of the evidence to the hypothesis would be only barely likely, because of the possibility that the corrupt senior officials were acting without the approval of Bogustan's president.

In some cases, the evidence is ambiguous because the evidence includes references to codewords that are being used to conceal the true purpose of what is being discussed and/or the identity of who is involved in the discussion. In these cases, there may be significant uncertainty regarding *the who and the what* in the evidence. The relevance of this evidence can vary from lacking support, if little or nothing is known about the codeword, to almost certain. During World War II, the Japanese used the codeword "AF" to describe the objective of an upcoming operation. US intelligence suspected that AF referred to the Midway Island but were not sure until the US sought additional information that confirmed that AF was Midway.

4.1.3 Conclusiveness

The evidence or sub-hypothesis may be inconclusive, and therefore its relevance should be less than certain. When the evidence or sub-hypotheses is inconclusive, the evidence or sub-hypothesis can support other scenarios despite the precision of the language describing the evidence or sub-hypothesis. Evidence that all nuclear power plants have been placed on a higher security level supports a scenario that an attack of some sort has occurred against nuclear power plant A. The evidence, however, also supports the scenario that an attack occurred against some other nuclear power plant or that official concern about a possible attack against nuclear power plants has increased but an attack has not yet occurred. Any one of those three scenarios could drive officials to impose tighter security at all plants. Without additional indicators, the relevance of the evidence of "higher security level" would be lacking support for each of the three equally plausible hypotheses. A situation in which there were just two equally plausible hypotheses would allow a relevance assessment of barely likely. In conjunction with other indicators, however, the combined relevance would exceed lacking support and barely likely, respectively, in both situations (two or three plausible hypotheses) depending on the strength of the other indicators.

How big an assumption has to be made to establish a relevance of certain is the main criteria for assessing the relevance, or, in other words, how big an analytic leap is necessary to accept the hypothesis given the evidence or sub-hypothesis—the larger the analytic leap, the less the evidence's relevance.

Example 23. *Hypothesis:* Managers were concerned about the possibility of an explosion at the Tanan plant in Bogustan prior to 25 November.

Evidence: Plant managers expressed concerns about safety procedures not being followed in an area of the plant where hazardous chemicals were being stored.

A defensible assessment of relevance in this case is barely likely and maybe lacking support. The evidence also supports the hypotheses/scenarios where managers are concerned about employees being injured as well as the hypothesis that managers are concerned about being cited for safety code violations. Knowledge of how much plant managers in Bogustan are concerned about the welfare of their workers and how strictly government agencies enforce safety regulations would inform the relevance assessment.

Evidence: Officials expressed concerns about safety procedures not being followed in an area of the plant where unstable chemicals were being stored.

A defensible assessment of relevance in this case is likely because unstable chemicals could be a source for a fire or explosion. The evidence still is not conclusive, but a smaller analytic leap needs to be made to get to a possible explosion, especially if Bogustan does not care very much about the welfare of factory workers and safety regulations are almost non-existent.

Example 24. *Hypothesis:* People and equipment on 25 November were sent to the Tanan chemical plant to fight an ongoing fire.

Evidence: A Defense Ministry official on 25 November said that fire-fighting equipment at Tanaka was not adequate and needed to be immediately increased.

A defensible assessment of relevance in this case is barely likely and possibly lacking support. The evidence supports a scenario where a fire is ongoing but it just as easily supports a scenario where officials are simply concerned about a possible fire and want to prepare for that contingency.

Evidence: A Defense Ministry official on 25 November said that the fire-fighting equipment at Tanaka was not adequate and needed to be immediately increased or the situation will surely deteriorate.

This evidence is still inconclusive, but its assessed relevance should be higher than in the previous example. The reference in the evidence to the possibility of the situation deteriorating more strongly supports a scenario where a fire is ongoing.

Example 25. *Hypothesis:* Uranium hexafluoride feedstock has been delivered to the Destructville plant.

Evidence: Several large flatbed trucks and a rail train were observed at the plant unloading large, multi-ton cylinders that are consistent in size and shape with cylinders used to transport uranium hexafluoride.

A defensible assessment of relevance in this case is likely. The evidence is inconclusive because we do not know whether the cylinders are filled or contain something other than uranium hexafluoride.

Evidence: Several large flat-bed trucks and a rail train were observed at the plant unloading large, multi-ton cylinders that are consistent in size and shape with cylinders used to transport uranium hexafluoride. A “UF6” label with a date was visible on the side of the cylinders.

A defensible assessment of relevance in this case is more than likely. We still do not know whether the cylinders are filled, but if they are, the likelihood that the cylinders contain something other than uranium hexafluoride is less.

4.1.4 Completeness

The evidence or sub-hypothesis may contain major information gaps or address only a portion of the geographic area and time period covered by the hypothesis, and thus is incomplete. The assessed relevance of evidence that is incomplete depends on how complete the information is. The more complete the information, the higher the relevance.

Major information gaps can occur for various reasons: The individuals discussing or reporting the information did not hear or have access to all of the information, or the information being discussed was deliberately left incomplete. In some cases, elements noted in the information provide a framework for positing a scenario but not much more.

Example 26. *Hypothesis:* People near the Tanan chemical plant are being evacuated.

Evidence: An individual at the plant reported that he overheard a conversation in which the plant manager said a situation involving buses and nearby residents who were in danger was urgent.

The information plausibly supports the hypothesis but the evidence does not warrant a relevance above barely likely. Other scenarios are possible: A bus accident involving multiple buses in which an oil or gas pipeline/facility/refueling station was damaged and was burning is one such alternative scenario.

The relevance of information from technical collection systems often depends on the extent of the coverage of the technical system. The extent of time and/or spatial coverage depends on the time or area that is the basis for the hypothesis. This can vary from one observation post or one small geographic area to an entire country.

The relevance of evidence stating that a review of optical imagery of facilities capable of producing chemical precursors showed no expansion to a hypothesis that a country has not expanded its precursor production capacity would depend on the extent to which the entire country was examined. The relevance would decrease when new plants could have been built and plants not known to produce chemical precursors could have been expanded to produce chemical precursors.

The relevance of evidence that observed guards walking the perimeter at three different times during a day to a hypothesis that guards regularly conduct around-the-clock patrols at the plant would have to be lowered to take into account the incompleteness of the coverage. Evidence that observed guards conducting perimeter patrols on different days at different times would have greater relevance to the hypothesis but would still not have a relevance of certain.

Inconclusive evidence typically supports multiple, different hypotheses whereas incomplete evidence typically can support only one hypothesis (or variations of the hypothesis that are different only in the degree to which they posit a specific scenario).

Example 27. *Hypothesis:* Bogustan has not deployed any SAM batteries to the border.

Evidence: Optical imagery of the border area did not observe any SAM batteries. (About 2/3 of the border area was imaged).

> If the unimaged portion of border area was suitable for the deployment of SAMs, the relevance of this evidence would be likely, all other factors being equal. Alternatively, the hypothesis could be changed to reference 2/3 of the border area, in which case, the relevance of the evidence would be certain. (The credibility of the evidence would address the issues of denial and deception and simple error-SAM batteries were present but were missed).

4.1.5 Mixed Relevance Credentials

All four components of relevance may need to be considered together when assessing relevance. Information can be irrelevant for different reasons at the same time.

Example 28. *Hypothesis*: Bogustan in November believes that it needs to produce chemical-warfare agents to counter an existential threat from Halifaza.

Evidence: The head of Bosgustan's national security council in January told the Defense Minister that the president was considering issuing an order to produce new weapons to offset Halifaza's military superiority.

> Ths information is dated (almost a year old); ambiguous (the reference to “new weapons” could mean almost anything); and inconclusive (the president had not made a decision). The relevance of this evidence to the hypothesis is lacking support there are too many uncertainties related to different aspects of the evidence to assess a relevance greater than 50 percent.

4.2 Credibility of Testimonial Evidence

We present an approach to assessing the credibility of a source which relies on the rich legacy of experience gathered over the past five hundred years in the Anglo-American adversarial judicial system concerning questions to ask of witnesses who appear in trials at law.

The credibility of testimonial evidence depends on the credibility of its source which, in turn, depends on the following credentials: competence, veracity, objectivity, and observational sensitivity or accuracy.

4.2.1 Competence

Access

The first question to ask related to competence is whether this source actually made the observation he claims to have made or had access to the information that he reports. In several accounts of intelligence analysis we have observed a glaring non sequitur. These accounts all say, “We can believe what this source has reported to us because he had good access to the information he reports.” This source may have had all the access in the world, but still not be credible in his report about what he observed. *The problem here is that HUMINT asset competence does not entail the asset’s credibility; competence is just one of the credentials of credibility*.

Understandability

The second competence question concerns whether this source understood what was being observed well enough to provide us with an intelligible account of what was observed. Thus, besides access, competence also involves understandability.

There are other situations, important in intelligence analysis, in which human competence is an issue. As you know, we rely on persons to inform us about the meaning of various forms of tangible and testimonial evidence. In addition, we rely upon many persons who process, edit, and transmit intelligence information of many kinds. Here are some examples:

- Analysts who interpret any form of information obtained from sensors;
- Persons who translate documents written in foreign languages;
- Persons who edit, transcribe, or summarize intelligence information;
- Persons who process raw sensor records.

In all of these situations, the competence at issue concerns the *skill* these persons demonstrate in performing their tasks. We might be as misled by a photo that is misinterpreted as we would be if a document written in a foreign language suffered from translation errors.

Table 5 presents 3 questions or tests concerning competence of human sources.

Table 5. Questions concerning the competence of human sources.

1. Access: Did this asset actually make the observation being claimed or have access to the information reported?
2. Understandability: Does this asset have an understanding of what was observed or have any knowledge or expertise regarding this observation?
3. Observational capability: Is this asset generally a capable observer?

4.2.2 Veracity, Sincerity, or Truthfulness

From experience we learn that people do not always believe what they are telling us. We would not say that a person was being untruthful if this person believed what he/she just reported. So the first question we should ask about the source telling us that event E occurred is: *Does this source believe that event E occurred?* Assessing veracity, sincerity, or truthfulness has been a problem for centuries and some mistakes have been made in explaining what veracity means. In many old and in some newer works it is said that a source is being truthful only if the event reported actually occurred. There is great trouble here since this explanation confounds veracity with the two other credibility attributes we need to consider. As an example, our source tells us that event E occurred and we later find out that E did not occur. Was this source being untruthful? Not necessarily, since this source might simply have been mistaken. So, the veracity of sources of HUMINT who report on what they observe depends on our assessment of whether these sources actually believe what they are reporting to us.

Table 6 presents 10 questions or tests concerning veracity of human sources. Take, for instance, the third veracity question: "Exploitation Potential: Is this source subject to any significant exploitation by other persons or organizations to provide us this information?" If there is evidence that this source is subject to any significant exploitation by other persons or organizations to provide us this information, then we cannot believe this source. Consider, as an example, a source whose family is detained by Al Qaeda and who has received threats from it to provide us this information.

Now let us consider the fourth veracity question:

> Any Existing Contradictory or Divergent Evidence: Is there any existing evidence that contradicts or conflicts with what the source has reported to us?

The following is an example of contradictory evidence where a source is inconsistent, telling us different things at different times.

Example 29. Here is a HUMINT source called "Rosebud." Two weeks ago, Rosebud told us she observed Amir D. running away from a car, just before a bomb in this car exploded in a crowded market in Baghdad on 14 May, 2013. Now today, she tells us that it was Omar T. who was running away from this car at this time and place. We have two options here, the first concerns Rosebud's veracity. We might say that her inconsistency shows that she is not keeping her stories straight, and so we should believe she made up a story and is not being truthful. Alternatively, we might believe that Rosebud has simply forgot who she observed running away from this car just before it exploded.

Table 6. Questions concerning the veracity of human sources.

1. *Goals of this Source:* Does what this source tells us support any of his or her goals?
2. *Present Influences on this Source:* Could this source have been influenced in any way to provide us with this report?
3. *Exploitation Potential:* Is this source subject to any significant exploitation by other persons or organizations to provide us this information?
4. *Any Existing Contradictory or Divergent Evidence*: Is there any existing evidence that contradicts or conflicts with what the source has reported to us?
5. *Any Existing Corroborative or Confirming Evidence*: Is there any other evidence that corroborates or confirms this source's report?
6. *Veracity Concerning Collateral Details*: Are there any contradictions or conflicts in the collateral details provided by this source that reflect the possibility of this source's dishonesty?
7. *Source's Character:* What evidence do we have about this source's character and honesty that bears upon this source's veracity?
8. *Reporting Record:* What does the record show about the truthfulness of this source's previous reports to us?
9. *Source Expectations about Us:* Is there any evidence that this source may be reporting events this source believes we will wish to hear or see?
10. *Interview Behavior:* If this source reported these events to us, what was this source's demeanor and bearing while giving us this report?

Here is another case of inconsistency.

Example 30. Another source named "Dingbat" tells us that it was Umar who was running away from the car just before a bomb in this car exploded in a crowded market in Baghdad on 14 May, 2013. We have relied on Dingbat to keep us informed about Umar in the past. We recall that Dingbat had previously told us that he and Umar were in Karbala all day on 14 May, 2013. Here again in this case, Dingbat's inconsistency may indicate his lack of veracity; but it may also mean that he was mistaken about the dates of these two events.

In assessing the veracity of a source, we should attempt to answer each of the above questions based on evidence. If we have evidence that the answer to each of the above questions supports the veracity of the source, then we can have high confidence in this source's veracity. If, however, many of the questions in Table 6 are not answered by our evidence, then our confidence in our assessment has to be lower. A very cautious approach would be to conclude that the source is not truthful if any of these tests is not passed. This approach is justified by the fact that we would do such an analysis for critical evidence for which we would need to have a high degree of confidence that it is true.

4.2.3 Objectivity

From common experience we observe that persons, including ourselves, often believe that some event has occurred because we either expect it to occur or want it to occur, regardless of what our senses are telling us. In such instances, we would say that this source lacks objectivity. An objective observer is one who bases a belief on the sensory evidence he/she received rather than on what this person expected or desired to observe. Suppose we believe that the source telling us that event E occurred is being truthful; he does believe that event E occurred. But now the question is: *Was this belief based on the sensory evidence this source received, or was it based on what this source expected or wished to observe?* One additional important matter concerns the role of memory. The reason is that our beliefs are elastic in nature; they change over time and often in response to new information we receive. If a HUMINT source made the observation some time ago, we might well question whether this source had the same belief at the time of his/her observation that this person now has while reporting to us. Is this person now telling us what he/she expected or wished to occur instead of basing this report on this person's recollection of what his/her senses recorded?

Table 7 presents 6 questions or tests concerning the objectivity of a human source. The following is a situation in which we have a HUMINT source giving us evidence about the credibility of another HUMINT source. In particular, it provides evidence relevant to answering the first two questions in Table 7.

Example 31. Suppose we are concerned about the credibility of a source called "Mable." Mable says he observed Yaqub M. placing an IED on a road leading from Baghdad to Samarra two days ago. We have another source, "Foxtrot" who knows "Mable" and who has told us useful things about Mable in the past. So, we tell Foxtrot about Mable's telling us about observing Yaqub M, placing the IED on the road between Baghdad and Samarra two days ago. So, we ask Foxtrot if Mable either expected Yaqub, or wished Yaqub, to be the person he saw placing the IED. Foxtrot says that Mable barely knows Yaqub and has no grounds for expecting or wishing that Yaqub was the person he saw. Foxtrot adds, "You can be confident that Mable told you what he did see and not what he expected or what he wanted to see."

Here is an example about memory and possible changes in a source's HUMINT testimony, related to the fourth question in Table 7.

Example 32. Here is Rosebud again who first tells us that it was Amir D. running away from a car, just before a bomb in this car exploded in a crowded market in Baghdad on 14 May, 2013, and then two weeks later now tells us that it was Omar T. and not Amir D. This may simply be the result of Rosebud’s changing her mind about who she saw, and not the result of her failing to keep her story straight.

Table 7. Questions concerning the objectivity of human sources.

1. *Source's Observational Expectations:* As far as this present report is concerned, do we have any evidence concerning what this source may have expected to observe?
2. *Source's Observational Desires:* As far as this present report is concerned, do we have any evidence bearing on what this source may have wished to observe?
3. *Belief-Formation Objectives*: As far as this present report is concerned, is there any evidence that this source may have believed it risky to form certain beliefs about what was being observed?
4. *Memory Effects on Beliefs:* Suppose this source is reporting about events he/she observed some time ago. How certain can we be that this source's present beliefs are the same as this source's beliefs were at the time of the source's observation? This not only involves how good the source's memory is, but also involves possible reasons why this source may have changed a belief.
5. *Any Existing Contradictory or Divergent Evidence:* It is entirely possible that any existing contradictory or divergent evidence may bear on this source's objectivity rather than on this source's veracity. Such evidence may point to this source's only lacking objectivity and not veracity.
6. *Any Existing Corroborative Evidence:* We may have evidence that may bear corroboratively on a source's objectivity rather than on this source's veracity.

4.2.4 Observational Sensitivity or Accuracy

Suppose we believe our source to be truthful and objective in their report that event E occurred. This source does believe that E occurred and this source did base this belief on sensory evidence obtained during a relevant observation. But now the question is: *How good was the sensory evidence this source received under the conditions in which this observation was made?* As we know, none of our senses are perfectly accurate or sensitive, particularly under a variety of ambient conditions such as reduced visibility and high noise levels. The physical condition of the source is also relevant here. We would question the adequacy of the sensory evidence this source obtained if he/she had some sensory defect or was intoxicated at the time the observation was made. Common experience tells us that human senses are not infallible and that we are all prone to make mistakes in our observations.

Table 8 presents 6 questions or tests concerning the observational sensitivity of a human source.

The following is a situation providing evidence on the allocation of attention.

Example 33. Here comes Mable again who tells us that he observed Yaqub M. placing an IED on a road leading from Baghdad to Samarra two days ago. We ask Mable how sure he is that the person he saw was really Yaqub M. Mable says, “Well I am pretty sure it was Yaqub M. who I saw as I was driving by, but I actually got only a brief look at him.”

Table 8. Questions concerning the observational sensitivity of human sources.

1. *Relevant Sensory/Physical Capacity:* What evidence exists concerning the source's physical and sensory capacities at the time this source made the observations forming the basis for this report?
2. *Allocation of Attention:* What do we know about the allocation of attention of this source on the reported event?
3. *Observational Conditions:* What do we know about the ambient conditions existing during the time this source made the observation forming the basis for this report? Did any conditions exist that could have influenced the accuracy of these observations?
4. *Past Accuracy Record:* What does the record show about this source's observational accuracy in past reports this source has provided?
5. *Any Existing Contradictory or Divergent Evidence:* It is entirely possible that any existing contradictory or divergent evidence may bear on this source's observational accuracy rather than on this source's veracity or objectivity.
6. *Inaccuracy Concerning Collateral Details:* Are there any less important details in this source's report that we suspect are inaccurate?

The following is a situation involving the observational conditions.

Example 34. Here comes our source “Dingbat” who tells us that it was Umar who was running away from the car just before a bomb in this car exploded in a crowded market in Baghdad on 14 May, 2013. We ask Dingbat how sure he is that it was Umar who was running away from the car. Dingbat says, “I am pretty sure it was Umar but I can’t be sure since it was a foggy day and I was about half a block away from the car.”

4.3 Credibility of Tangible Evidence

Before we describe tangible credibility attributes in some detail, we will first argue that we should preserve in Cogent the distinction between real and demonstrative tangible evidence. Again, real evidence is a thing itself and demonstrative evidence is just a representation of a thing. One reason for preserving this distinction is that attributes of the credibility of tangible evidence, and the questions they raise, depend upon whether the tangible item is real or demonstrative in nature. One problem that we will return to later is that we may have a bit of difficulty on occasion deciding whether an item of tangible evidence is real or demonstrative in nature. Discussions of tangible evidence reveal three major attributes of the credibility of these forms of evidence: authenticity, reliability, and accuracy. We will show which ones of these attributes arise in real and in demonstrative tangible evidence and the questions they raise for the evaluators of such evidence. Exactly who evaluates the credibility of tangible evidence in intelligence agencies is an interesting problem all on its own.

The major problem is that certain forms of tangible evidence will not require consideration of all of these three credibility attributes. For example, real tangible evidence may require only authenticity considerations. Demonstrative tangible evidence, on the other hand, may require consideration of all three attributes: authenticity, reliability, and accuracy. What is comes to is that authenticity is always a requirement for any form of tangible evidence, but reliability and accuracy are attributes mainly of demonstrative tangible evidence. Further, what questions are raised concerning reliability and accuracy depend upon the tangible evidence item itself.

To complicate credibility matters even further, we have to give attention to the persons who actually assess both the credibility and meaning of all forms of tangible evidence, whether it is real or demonstrative. There are questions concerning who will perform these assessments: intelligence analysts or other persons who serve them in various ways such as the collectors, processors, and interpreters of evidence. An analyst considering some item of HUMINT from a foreign asset rarely, if ever, can confront this asset directly and will also be privy to all the sources of evidence we might have about this asset. These problems are especially difficult concerning tangible evidence, whether real or demonstrative. In many or most cases, analysts will never see tangible items themselves but will receive written accounts of these items that are provided by other intelligence professionals who process and evaluate this evidence in various ways.

4.3.1 Authenticity

The word authenticity stems from the Greek *authentikos* meaning "genuine". So, something authentic is genuine or something verified as being what it is claimed to be. The term authenticity is fundamental in intelligence analysis as it is in other contexts such as law. This term is widely used with reference to tangible things and also to people; is this person who he claims to be? In the affairs of individuals, as well as of nations or organizations, the variety of ways in which one person, nation, or organization has tried to mislead other persons, nations, or organizations is endless. The production of counterfeit currency goes back to the earliest ages in which tangible money was used as a medium of exchange. Who knows how many forged documents or other contrived exhibits have given rise to military disasters or erroneous verdicts in trials at law. One of the most recent examples is a document recording the request for certain nuclear materials from the country of Niger, allegedly made by Saddam Hussein. At least some persons used this letter as evidence of Saddam's development of nuclear weapons, when the letter was discovered to be fraudulent. In the Sacco and Vanzetti case, a 32-caliber bullet was exhibited at trail that was represented as being the bullet that caused the death of a payroll guard the two defendants were charged with killing. This 32-caliber bullet may have come from a 32-caliber automatic that Sacco was carrying when he was arrested. However, arguments exist to this day about whether the bullet shown at trial was the same one extracted from the guard's body. It is possible that the prosecution showed a different bullet that had been test-fired through Sacco's Colt automatic while the trial was in progress.

There are very interesting and difficult authenticity questions that are raised by information obtained from web sites. It is now well known that digital photos and other records on a web site can easily be altered in various ways. Photos can be easily changed and tabled records can be altered. The problem is that detection of these alterations is extremely difficult. So, when we observe a photo of known terrorist at a certain location on an Jihad web site, it may be very difficult to determine the authenticity of this photo. If we have any record that has been digitized, there is no way of telling, without having other evidence, whether this record has been subtly or radically altered. One writer, citing conclusions of a seventy-member international consortium on digital imaging, says that one conclusion was that: "digital photographs, as evidence of reality, are dead" (Paul, 2007, p.4). There are ways of preventing digital records from being altered, but they require special software systems. These systems employ what is termed Publlc Key Infrastructure [PKI] for encrypting digital records so that their authenticity can be preserved. Unfortunately, these PKI systems do not allow us to detect changes in records that have been made by others such as the Jihad organization that are displayed on their own web sites. So, the authenticity of any digitized records from adversarial groups will be difficult to establish.

The bottom line on authenticity is that it concerns all imaginable forms of tangible evidence, real or demonstrative. We are as much concerned about the authenticity of the objects or things themselves as we are about various representations or illustrations of things. In the examples given above, we are as much concerned about the authenticity of the IED detonator we received from "Rambo" as we are about the hand-drawn map we received from the captured Shiite militia member.

4.3.2 Reliability

Here is a credibility-related term that we believe has been often misused in intelligence analysis and elsewhere when it is applied to human sources of testimonial evidence. It is true that common synonyms for "reliability", applied to people, often include: "dependable", "devoted", "loyal", "staunch", and "dedicated". But what counts as far as the credibility of human sources of evidence is concerned is the credibility of what they tell us on particular occasions. The credibility of a human source depends on her competence, veracity, objectivity, and observational sensitivity as far as this particular report is concerned. In short, the credibility of a source of HUMINT depends on more things than a source's reliability.

In many contexts, science and engineering for example, a reliable process of any kind is one that produces consistent or repeatable results. A test of any kind is reliable to the extent that it provides the same or nearly the same results on repeated occasions under similar conditions. The *Oxford English Dictionary* says that: In statistics, reliability means the extent to which a measurement made repeatedly in identical circumstances will yield concordant results. You say your car is reliable to the extent that it will take you where you wish to go for some time in the future. This assumes that you keep your car well maintained and that you do not operate it in foolish or unusual ways. Applied to sensing devices we ask: would this device have provided us with the same information again and again on other similar occasions on which this device could have been employed? How well this device has been designed and maintained surely bears directly on its reliability.

Unless we are wrong, the attribute reliability applies mainly to any source that provides us with demonstrative tangible evidence in the form of images or various other representations of things. Here are a few examples. An analyst does not have the soil samples themselves that were allegedly obtained near Leninsk in Russia; he only has an account of a computer based MASINT chemical analysis of this soil sample. The reliability issue here concerns the extent to which this computer based chemical analysis would yield the same results if it were done over on other occasions. Here is an ELINT image of the probable locations of radar installations within a hundred-mile radius of Estafan in Iran. Would additional images taken by the ELINT sensors give us the same locations?

As I briefly mentioned above, there are a few problems concerning what we take to be real or demonstrative evidence. For example, here is the COMINT recording of a phone conversation between the high-ranking government official in Iraq and the Iranian IRGC member. We listed this as being real evidence even though we have only a recording of this conversation and did not hear the actual conversation ourselves. Since an agency we trust made this recording, we should probably not be concerned about its authenticity, unless the two persons involved in this conversation were misidentified. The reliability issues here concern the recording process itself, the identification of the persons involved, and the process by which this conversation was revealed to analysts. Very likely analysts would only be provided with a written account of this recording. We also gave an example of real evidence being the IMINT photo of military aircraft at an Iranian air force base outside of Estafan in Iran. Analysts might receive this photo but would also receive a written account by an image analyst regarding what this photo reveals. The reliability of the camera as well as the reliability of the image analyst would be of concern here. Would the analyst make the same identifications on repeated examinations of this photo?

There are no rules prescribing whether a tangible item is real or demonstrative in nature. In some cases, such as those just mentioned, it may not be obvious whether a tangible item is real or demonstrative. It may call for careful judgments on the part of the analysts.

4.3.3 Accuracy

This attribute seems to have variations depending on the tangible item whose credibility is being assessed. In some cases we may be concerned about the accuracy, sensitivity, or resolving power of the sensor that recorded events of concern, whether we take this tangible evidence to be real or demonstrative. But in cases in which the item is patently demonstrative we will have other concerns. An example of patently demonstrative evidence is the hand-drawn map captured from the Shiite militia member. What we may term the representational accuracy of a tangible item is the extent to which it accurately depicts the matters it is allegedly representing. This will involve such matters as the scale, contours, dimensions, and details of the demonstrative item. In the case of the hand-drawn map, we would also be concerned about whether the location depicted on the map was accurately described. If not, the map would be quite useless. Presumably, we could have access to assessments of the accuracy, as well as the reliability, of our own sensing devices. These records might well be considered by the experts who assess the credibility of images coming from these sensors and what events these sensor images reveal.

4.4 Credibility of Human Sources of Tangible Evidence

We believe that the greatest threats to our being intentionally misled come from foreign assets who supply both tangible and testimonial evidence. We can mislead ourselves in many ways by mishandling tangible evidence from our own sensors; but we can also take various steps to reduce such possibilities. Concerning real and demonstrative tangible evidence provided by foreign assets or our own military or other personnel in the field, the basic protection we have against being misled comes from very careful assessments of the credibility of these tangible items and the credibility of their human sources.

4.4.1 Competence Attributes

So far we have identified five attributes of asset competence we are asking the user to consider for testimonial evidence. We believe two of these attributes are most important as far as tangible evidence is concerned.

1) What evidence do we have that an asset was actually in a position to obtain the tangible item he has just provided for us? Stated another way, what evidence do we have that this asset could have had access to this tangible item? The asset will either volunteer, or be required to tell us where, when, and perhaps how, he obtained this item. We would have immediate grounds for suspecting the authenticity of this item if we had credible evidence that this asset was not at the location where and when he says he obtained this item. If we also had evidence that this item was in an inaccessible location, we would naturally wonder how he obtained it. Naturally, of course, we have the authenticity of the item itself to consider. The asset my be entirely truthful about where, when, and how he obtained this tangible item but has knowingly or unknowingly passed an inauthentic and misleading item on to us.

2) Does the asset have an understanding of the significance of the tangible item he has just provided for us? Answering this question may depend in part upon whether the asset was

instructed by us to obtain this tangible item. For example, if we asked him to try to obtain a certain document, the asset may wonder why we are interested in it. But if the asset obtains some tangible item on his own, he must have thought it significant enough to provide it for us. What's interesting is that this attribute may involve the question: Why are you giving us this tangible item? Earlier, we gave an example involving an asset who passes a grocery list on to us. This may appear to be insignificant to us until the asset shows us how it contains a coded message intended for some insurgent group. This would indicate considerable understanding on the part of the asset, particularly if he can decode this message correctly.

Regarding the example in item 2) above, we ran across an interesting real example of problems we face in dealing with our current adversaries. Someone recorded the following e-mail message. The message reads, in part:

I would like to clarify the following with relation to the birthday.

(a) Don't think of showering because it may harm your health.

(b) We can't make a reservation for you, but they usually don't mind making reservations for guests. Those who wish to make a reservation should go to Quwedar.

(d) I don't have any gravel.

This message was actually sent from bin Laden's chief deputy Ayman al Zawahiri to al Qaeda members in Yemen on February 1, 1999. The word *birthday* means *attack*, *Quwedar* is a pastry shop in Cairo, and *gravel* means *ammunition or bomb-making material* (Emerson, 2006, pp. 468 – 469). Someone with inside understanding must have examined this message to determine its real meaning. Lack of authenticity has many faces. It often takes persons having an unusual degree of understanding to detect it.

4.4.2 Veracity attributs

We will argue that the veracity of the asset providing us with real or demonstrative tangible evidence is the only other credibility attribute we should consider. There is an exception that we will note in a minute. What we are mainly concerned about is the truthfulness of the asset's account of where, when, how, and [perhaps] why he acquired this tangible item. Evidential answers to these questions are of vital importance.

Here comes the exception we mentioned. In at least some cases an asset will accompany a tangible item with testimony about events he observed while acquiring this item. In this case what he provides is a mixture of tangible and testimonial evidence. In such cases we have the asset's veracity to consider as far as the tangible item is concerned, as well as his veracity, objectivity, and observational sensitivity as far as his testimonial assertion is concerned. In section 4.4.1 we have already considered the asset's competence as far as the tangible item is concerned. But we also have his competence as far as concerns any additional testimony he provides. There are some very interesting issues here. Suppose we obtain answers to veracity and competence questions regarding the tangible item that differ from anwers to these question concerning the testimony he provides. This invites some very difficult and interesting questions. Should we accept the authenticity of the tangible item or the credibility of the testimonial assertion? Or should we discard both as not being worthy of belief? The tangible and testimonial items appear linked together in such cases and make the credibility assessment process very difficult.

4.5 Credibility of Chains of Custody

4.5.1 What Is a Chain of Custody?

In the previous sections we have discussed the different types of evidence (such as testimonial or tangible), and the ingredients of their credibility assessment. However, very rarely if ever does the analyst have access to the original evidence. Most often, what is being analyzed is an item of evidence that has undergone a series of transformations through a chain of custody. Here we have borrowed an important concept from the field of law, where a chain of custody refers to the persons or devices having access to the original source evidence, the time at which they had such access, and what they did to the original evidence when they had access to it. The original evidence may be altered in various ways at various links in chains of custody. The important point here is to consider the extent to which what the analyst finally receives is an authentic and complete account of what an original source provided. Uncertainties arising in chains of custody of intelligence evidence are not always taken into account. One result is that analysts are often misled about what the evidence is telling them.

Basically, establishing a chain of custody involves identifying the persons and devices involved in the acquisition, processing, examination, interpretation, and transfer of evidence between the time the evidence is acquired and the time it is provided to intelligence analysts. Lots of things may have been done to evidence in a chain of custody that may have altered the original item of evidence, or have provided an inaccurate or incomplete account of it. In some cases, original evidence may have been tampered with in various ways. Unless these difficulties are recognized and possibly overcome, intelligence analysts are at risk of drawing quite erroneous conclusions from the evidence they receive. They are being misled, not by our original sources of evidence, but by the activities of our own persons or devices who do various things to incoming intelligence evidence.

In civilian and military courts, proponents of evidence, for either side of the matter in dispute, are required to verify the chain of custody of tangible evidence before it is admitted to trial. In many cases, evidence gathered is passed from one person to another, each of whom may examine and process the evidence in various ways. In many situations, proponents are required to select experienced and credible persons who serve as evidence custodians. The major task for these persons is to establish the chain of custody of evidence keeping records of who gathered the evidence, the persons who had access to the evidence, the times at which they had access, and what these persons did with the evidence while they had access to it. A very good account of questions regarding the chains of custody of evidence in our courts is to be found in (Lempert et al., 2000, pp. 1167 – 1172).

Now, we are of course not privy to the actual chains of custody of various forms of intelligence evidence in any of our intelligence organizations. And we do not know whether there are any appointed evidence custodians in these organizations as there are in our legal system. However, we can offer accounts of chains of custody of evidence that seem reasonable and necessary for different forms of evidence. In our examples of how Cogent can assist analysts to establish the credibility of evidence, we have established conjectural accounts of chains of custody for an item of testimonial and an item of tangible evidence in order to illustrate the virtues of Cogent and how it can assist the intelligence analyst to make these very difficult credibility assessments.

First, suppose we have an analyst who is provided with an item of testimonial evidence by an informant who speaks only in a foreign language. We assume that this informant's original testimony is first recorded by one of our intelligence professionals then translated into English by a paid translator. This translation is then edited by another intelligence professional and then the edited version of this translation is transmitted to an intelligence analyst. So, there are four links in this conjectural chain of custody of this original testimonial item: recording, translation, editing, and transmission. Various things can happen at each one of these links that can prevent the analyst from having an authentic account of what our source originally provided.

Then, suppose that an analyst is provided with an account of a tangible item in the form of a digital photo. This photo has been taken by one of our foreign assets. We note that the analyst may see a copy of the photo itself, or just a written account of the events recorded in this photo. Suppose in this case the analyst only receives a written account of what this original photo revealed. We have supposed that this digital photo is first transferred to the computer of one of our intelligence professionals; it is then transmitted to a photo interpreter; the image is interpreted by this person; and then the written interpretation of this photo is transmitted to the analyst. So, in this conjectural chain of custody there is a transference link, an interpretation link, and two transmission links. At any of these links there are possible reasons why what the analyst receives is not an authentic account of what the asset's original photo depicted.

There are many possible chains of custody, for different types of evidence, as illustrated in 0. However, they can all be characterized by a chain of basic evidence transformation processes (such as translation, editing, or transmission). Moreover, for each such process, one can identify the ingredients and the arguments of its credibility assessment, just as for the different types of the evidence. An earlier system, Disciple-LTA (Tecuci et al., 2005a; 2007b; 2008a), employs a systematic approach to the assessment of the credibility of items of evidence obtained through a chain of custody (Schum et al., 2009). The same approach may be used with Cogent.

With these evidential and chain of custody ideas in mind, we can now show how Cogent can assist intelligence analysts to assess the many sources of uncertainty associated with the authenticity of evidence they receive.

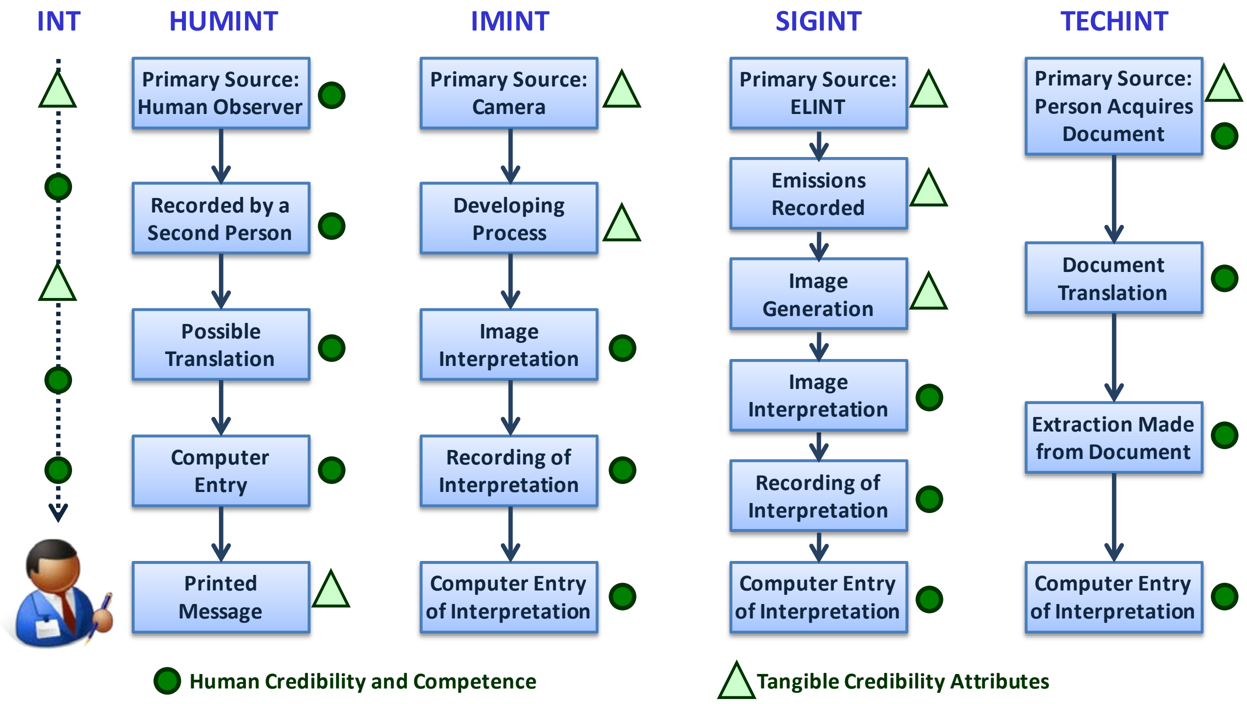
INT
HUMINT
IMINT
SIGINT
TECHINT
Primary Source: Human Observer
Recorded by a Second Person
Possible Translation
Computer Entry
Printed Message
Primary Source: Camera
Developing Process
Image Interpretation
Recording of Interpretation
Computer Entry of Interpretation
Primary Source: ELINT
Emissions Recorded
Image Generation
Image Interpretation
Recording of Interpretation
Computer Entry of Interpretation
Primary Source: Person Acquires Document
Document Translation
Extraction Made from Document
Computer Entry of Interpretation
Human Credibility and Competence
Tangible Credibility Attributes

Figure 40. Typical chains of custody for different INTs.

4.5.2 A Case Involving Chains of Custody

The cover story for this hypothetical case involves an experienced analyst named Clyde who is involved with intelligence analyses concerning matters in Iraq. Clyde's present inferential problem involves an Iraqi named Emir Z., a respected official of the government in Iraq. Emir Z. has publicly argued on many occasions about the necessity of stopping the sectarian violence that has plagued Iraq and coalition efforts to achieve stability in this country since the U.S. and coalition intervention in 2003. Clyde's present problem is that he wonders how respectable Emir Z. really is. Clyde has evidence that Emir Z. has been in contact in Iran with representatives of the Iranian Islamic Revolutionary Guards Corp (IRGC). We already have a variety of strong evidence that the IRGC has been involved in supplying weapons, training, and intelligence to various Shiite militia groups in Iraq. This has certainly not contributed to stability in Iraq. So, Clyde entertains the hypothesis H_1: "Emir Z. is collaborating with the IRGC" (i.e., he is not the respected official we have believed him to be).

So far, Clyde has two items of evidence bearing on H_1. He first has an item of testimonial evidence from a source code-named "Wallflower" who reports that five days ago he saw Emir Z. leaving a building in Ahwaz, Iran in which the IRGC has offices. Wallflower, an Iranian national, issued his report in the Farsi language that was recorded and then translated into English by a paid interpreter. Then an edited version of this translation is recorded and transmitted to Clyde. Then Clyde receives an item of tangible evidence in the form of a photograph taken of Emir Z. eight days ago at an IRGC Qods Force base outside Dezful in Iran. This photo was taken by another source, code-named Stovepipe. The identification of Emir Z. in this photo was verified by one of the U.S. intelligence professionals who has had contact with Emir Z. We will assume that Clyde receives this photo together with a written account of what this photo revealed. But we also allow for the possibility that Clyde receives only a written account of the contents of this photo.

The following sections present the inferential problems Clyde faces as he attempts to assess how believable he imagines these two items of evidence to be.

4.5.3 A Chain of Custody for Testimonial Evidence

We begin by examining the credibility of the information Clyde has received from Wallflower. This information has passed through a chain of custody that is illustrated in 0. The top part of 0 shows the successive transformations suffered by Wallflower's original testimony (E001-Wallflower-testimony) until it reaches Clyde, who actually receives the item of evidence E005-Emir-Iran. Wallflower provides testimonial evidence in the form of an assertion he made concerning Emir leaving the IRGC building in Ahwaz, Iran. Wallflower says he based this assertion on his direct observation of these events five days ago. First, Clyde has not heard Wallflower's original testimony and could not possibly have understood it unless Clyde speaks Farsi. We have identified a case officer named Bob who may only speak limited Farsi. Bob records what Wallflower has testified on a Sony STEMQ recording device. This recording is transmitted to a paid foreign national named Husam who speaks fluent Farsi. Husam's written translation of Wallflower's testimony about Emir is then transmitted to a reports officer named Marsha. Marsha edits Husam's translated version of Wallflower's testimony and prepares her written edited version for transmission over a (fictitious) system we will call SN 247. What Clyde receives is Marsha's edited version of Husam's translation of Wallflower's original testimony. This is a tangible item that is only an account of what Wallflower originally said.

In this example we have three identified persons involved in the chain of custody of Wallflower's report: case officer Bob, translator Husam, and reports officer Marsha. Now one thing about this fictitious example is that the analyst, Clyde, may or may not know the identities of these three persons. Obviously, in an actual situation, there may be more or different persons involved in a chain of custody.

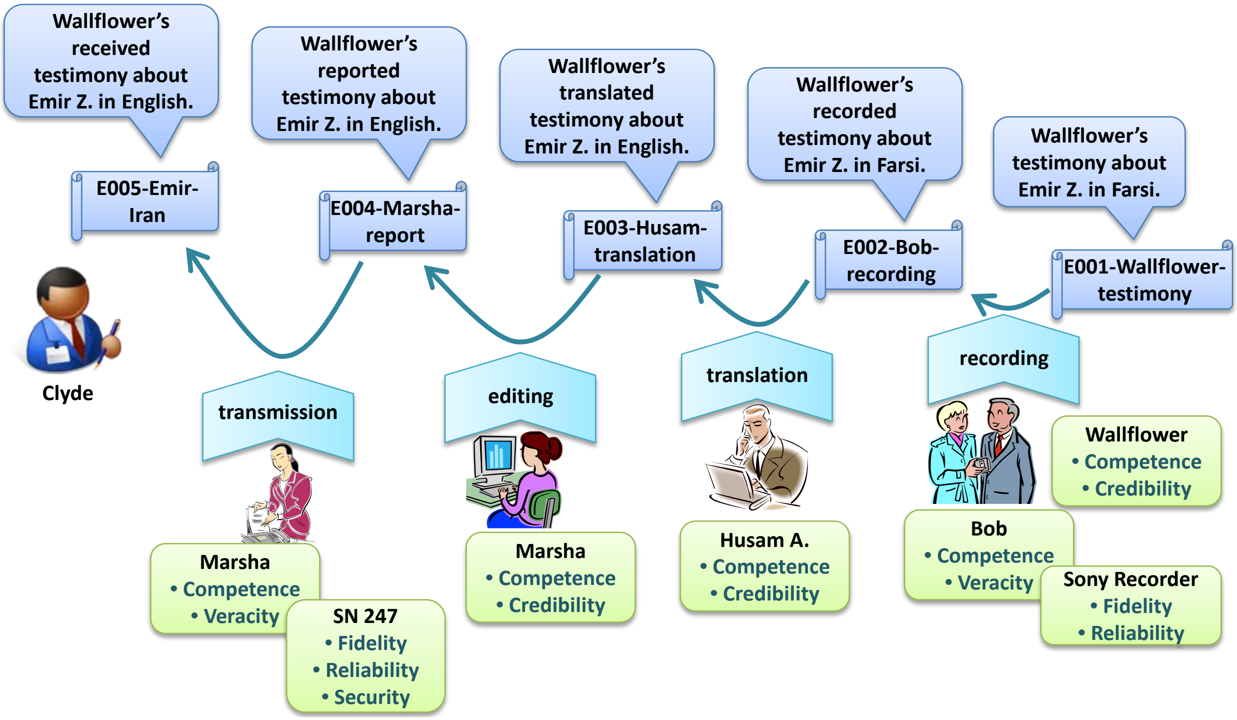

Figure 41. The chain of custody of the Wallflower's testimony, and the processes involved in this chain.

The reader may quickly note other links in this chain of custody that we have overlooked, or have different labels for the ones we have included. Not being privy to any actual chains of custody of intelligence evidence, our conjectural chains of custody may lack realism. But our claim is that they are plausible enough to illustrate the kinds of uncertainties encountered in chains of custody of intelligence evidence and also to illustrate how Cogent can assist analysts in making these kinds of assessments.

The major issue here is the extent to which Clyde can believe what he is told Wallflower said. This belief rests not only on evidence regarding Wallflower's credibility, but also on the credibility issues raised by what was done to Wallflower's original report about Emir before Clyde received a report of what Wallflower testified. The bottom part of 0 shows the credibility issues that can arise regarding the persons and the devices involved in this example.

Let's start with the original source, Wallflower, who provides testimonial evidence in the form of an assertion he made concerning Emir leaving the IRGC building in Ahwaz, Iran. Wallflower says he based this assertion on his direct observation of these events five days ago. In deciding whether to believe Wallflower, Clyde must first consider Wallflower's credibility, as discussed in Section 4. Wallflower's competence involves his access and his understanding. The basic access question involves asking whether Wallflower was actually in a position to observe what he tells Clyde. The understanding question asks whether Wallflower knew enough about what he was observing to give Clyde an intelligible account of what he observed. His credibility also involves the veracity, objectivity, and observational sensitivity attributes.

Then Wallflower's testimony (E001-Wallflower-testimony) was tape-recorded by the case-officer Bob, who interacts with Wallflower. So, we have natural concerns about the fidelity and reliability of this tape recording, as well as about Bob's competence and veracity. Among other things, is the recording understandable and complete? Was all of Wallflower's testimony on the recording and did no gaps appear in it? Clyde would also be interested to know whether Wallflower provided his report voluntarily or whether he was asked to report on Emir by Bob. This is quite important since if Wallflower gave this report voluntarily and we believe he is being untruthful, Clyde has to ask why Wallflower told Bob this particular lie in preference to any of the others he might have told. The same question will arise for Stovepipe's evidence.

As we have specified, Wallflower speaks only Farsi and not English, and Husam has translated his recorded testimony (i.e., E002-Bob-recording) into English. We have concerns here about the credibility attributes of Husam. Competence involves not only knowledge of Farsi and English, but also knowledge of the subject matter being translated. The translated account of Wallflower's original testimony (i.e., E003-Husam-translation) is then edited by Marsha. We may also have concerns about Marsha's credibility.

Finally, this recorded, translated, and edited account of Wallflower's testimony (i.e., E004-Marsha-report) is transmitted through a computer network to Clyde and possibly many other interested persons. What we have here are concerns about the credibility of the person who performed the transmission, and also about the fidelity, reliability, and security of the transmission.

4.5.4 A Chain of Custody for Demonstrative Tangible Evidence

Let us now consider the tangible evidence supplied by Stovepipe: a photo he says he took eight days ago of Emir at an IRGC base outside Dezful, Iran. Clyde must first consider Stovepipe's competence. What evidence does Clyde have that Stovepipe was actually in Dezful at the time he says he took the photo? In addition, what evidence does Clyde have that Stovepipe knew the person he was photographing? Clyde also has one credibility attribute for Stovepipe to consider, namely his veracity. Was Stovepipe truthful in telling us when, where, how, and why he took this photo? We will assume that we are treating the photo provided by Stovepipe as being demonstrative tangible evidence. There are several reasons why this makes sense: First, we have the credibility of the photo itself to consider. Is this photo authentic (Is it what it is represented as being, namely Emir at the IRGC base outside Dezful, Iran)? Second, has this evidence come from a reliable sensing device that would supply us with repeatable information? Third, is the evidence accurate in allowing Clyde or anyone else to tell whether it was really Emir in the photo?

These three maters all concern the credibility of the photo itself.

Unfortunately, regarding the photo of Emir Z., allegedly taken outside of Dezful, Iran, eight days ago, we have two possibilities to consider. We have to consider whether Clyde was given a copy of this photo to examine himself, or whether he was just given a written account of what this photo depicted. The major trouble, of course, is that the chains of custody will be different in these two cases. So, we will make some conjectures about what the chains of custody might look like in both of these cases. The first is where Clyde is given the photo to examine; the second is where he is just given a written account of what is in Stovepipe's photo.

4.5.5 Chain of Custody for a Photo Given Directly to the Analyst

We will begin with the case in which Clyde sees Stovepipe's photo itself or, more than likely, a copy of Stovepipe's original photo. The corresponding chain of custody involves the persons and processes shown in 0. The top part of 0 shows the successive transformations suffered by Stovepipe's original photo of Emir (called E006-Emir-photo-in-camera) until it reaches Clyde (as E010-Emir-Iran). First, suppose Stovepipe used an Olympus #AB1 digital camera to take the photo of Emir Z. We suppose that Stovepipe had some means for transferring this digital photo to case officer Bob's laptop computer. Bob then transmits this stored digital photo to a photo interpreter we shall name Mike. Mike examines the photo to assess its authenticity, and he also verifies that the person depicted in the photo is Emir Z. Mike prepares a written account of his analysis of this photo. Mike then transmits his written account of this photo, and possibly a copy of the photo itself over our (fictitious) SN 247 system. This written account of the photo, and possibly a copy of the photo, is eventually routed to Clyde for his analysis. The bottom part of 0 shows the credibility issues that can arise regarding the persons and the devices involved in this example.

In this example, we have chosen to have Stovepipe use a digital camera to eliminate the necessity of considering by whom and where the photo was developed. We will also assume that Stovepipe, who knows Emir, was instructed to follow Emir eight days ago in Dezful, to see where he went that day. Additionally, we also will assume that Stovepipe was able to visit Bob, bringing his digital camera with him. There are, of course, other means by which Stovepipe might have communicated with Bob. If Stovepipe had a laptop computer, he could have e-mailed the message to Bob, who could have been anywhere. But this would add additional risks involving the e-mail being intercepted by the Iranian IRGC. Surely, there are risks associated with Stovepipe's meeting personally with Bob. It took Stovepipe eight days between the time he took the photo and the time he delivered it to Bob. How Stovepipe actually got the photo to Bob is interesting and will eventually bear on Stovepipe's competence and the authenticity of the photo. Maybe the procedures necessary for Stovepipe to communicate directly with Bob are very complex and have been designed to reduce the risks associated with this direct communication. For example, suppose Bob is in Baghdad, Iraq, but meets with Stovepipe in Al Amara, Iraq. Both cities are near the Iraq-Iran border, about 80 miles from each other. Perhaps it takes several days for Stovepipe to communicate with Bob and then arrange the cover necessary to go into Iraq. These are all matters that bear upon Stovepipe's competence and may also concern the authenticity of the photo.

But the major assumption underlying this scenario concerns whether or not we can say that Clyde is provided with the original photo that Stovepipe allegedly took. This involves the assumption that the image Bob uploaded on his computer and then transmitted to Mike was not altered in any way, nor was the image that Mike transmitted to Clyde altered. If the image that Clyde received had exactly the same pixels as the one on Bob's computer, we could probably say that Clyde received the same original photo that Bob received from Stovepipe.

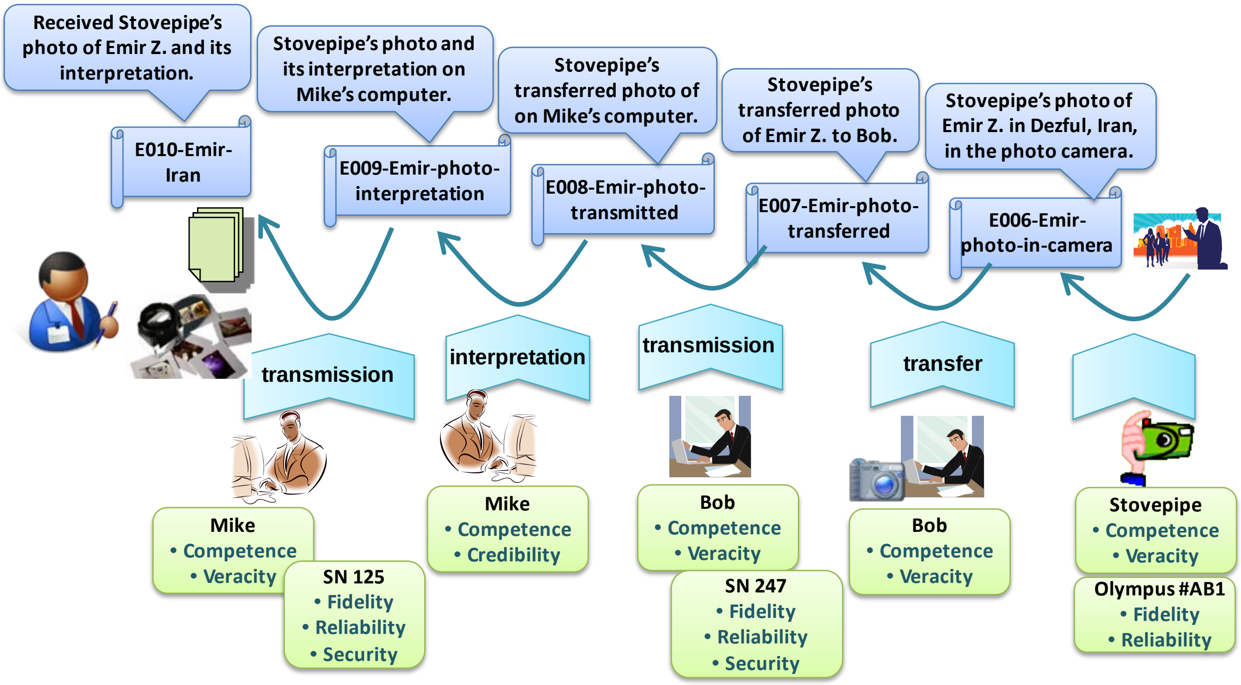
Received Stovepipe's photo of Emir Z. and its interpretation.
Stovepipe's photo and its interpretation on Mike's computer.
Stovepipe's transferred photo of on Mike's computer.
Stovepipe's transferred photo of Emir Z. to Bob.
Stovepipe's photo of Emir Z. in Dezful, Iran, in the photo camera.
E010-Emir-Iran
E009-Emir-photo-interpretation
E008-Emir-photo-transmitted
E007-Emir-photo-transferred
E006-Emir-photo-in-camera
transmission
interpretation
transmission
transfer
Mike
• Competence
• Veracity
SN 125
• Fidelity
• Reliability
• Security
Mike
• Competence
• Credibility
Bob
• Competence
• Veracity
SN 247
• Fidelity
• Reliability
• Security
Bob
• Competence
• Veracity
Stovepipe
• Competence
• Veracity
Olympus #AB1
• Fidelity
• Reliability

Figure 42. The chain of custody of the demonstrative tangible evidence sent to Clyde.

4.5.6 Chain of Custody for a Written Description of a Photo Given to the Analyst

This second case is interesting for the following reason: We are treating the photo as being demonstrative tangible evidence since it is just a possible representation of Emir Z. being outside of Dezful eight days ago. Then, if Clyde is only given a written account of this photo, this is only demonstrative evidence of Stovepipe's original demonstrative evidence. So, in cases such as this, we have a chain of demonstrative evidence involving two or more sources.

The above discussion shows how complex the processes of analyzing the probability of the hypothesis "H_1: Emir Z. is collaborating with the IRGC" are, even when we have just a small amount of evidence. However, for a good analysis, one would have to consider many more items of evidence, and to consider both favoring and disfavoring evidence. Unfortunately, the complexity of the analytic process grows so much with the addition of new items of evidence that all the involved probabilities cannot be assessed. There is simply not enough time for the analyst to assess them, or evidence necessary to support these assessments is not available.

4.5.7 Drill-Down Analysis of Chains of Custody

As indicated above, the analyst may not have the time or the evidence to assess all the factors involved. In such a case, Cogent allows him/her to make assumptions with respect to the solutions of the unsolved problems. For example, the analyst may assume that the credibility of most of the processes involved in the chain of custody is almost certain, and concentrate his analysis on assessing the credibility of Wallflower. He may further make some assumptions about Wallflower’s competence, veracity, objectivity, and observational sensitivity, and then Cogent will automatically estimate the overall credibility of E005-Emir-Iran.

But the analyst can also drill-down to analyze each of the processes from the chain of custody. For example, the based on the general fidelity and reliability of a Sony STEMQ recorder, the recording performed by Bob may be assessed as “certain.” Cogent also allows the analyst to investigate all sorts of “what-if” scenarios. For example, Clyde may consider alternative values for the veracity of Husam A., the translator of Wallflower’s testimony.

4.6 Review Questions

92. Given a hypothesis that a new missile called X was flight tested on 1 July, which of the items of evidence below would have the highest relevance to this hypothesis:
    E1: The missile that was flight tested demonstrated the same range as missile X.
    E2: A source reported in August that missile X was flight tested on 1 July.
    E3: A source reported in April that a flight test of missile X was scheduled for 1 July.

93. Given a hypothesis that Mark robbed the bank, which of the items of evidence below would have the lowest relevance to this hypothesis:
    E1: The first five digits of the license plate number of Mark’s car matches the first five digits of a six-digit license plate number of the car that the bank robber used in his escape
    E2: Mark was not at home at the time of the bank robbery
    E3: Mark told a friend a week before the bank robbery that he was planning to rob a bank

94. Consider the following arguments:

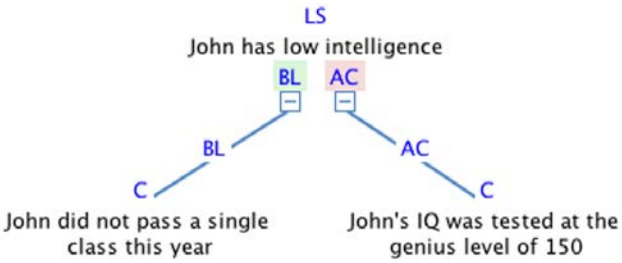
LS
John has low intelligence
BL
AC
BL
AC
C
C
John did not pass a single class this year
John's IQ was tested at the genius level of 150

Figure 43. Cogent argumentation.

The relevance of the sub-hypothesis *"John did not pass a single class"* was assessed as BL because:

a) If John did not pass a single class, it shows conclusively that John has low intelligence.
b) There may be a host of other reasons why John did not pass a single class.

95. However credible you believe a person might be, you also must give consideration to this person's competence. Provide an example of a credible source which is not believable.

96. How does intelligence analysis differ from evidential analyses in law trials as far as the completeness of evidence is concerned?

97. Can you think of instances in which you might say you have conclusive evidence when this is actually not correct?

98. The leadership in Country T has embarked upon an aggressive track regarding its relationship with neighboring countries. We are presently assessing the capability of Country T to wage war on a country with whom we have very friendly relations. We suspect that policy makers in T are considering the development of a certain tactical weapon system we will call W. If they are successful in developing system W, this would give T a decided advantage in any armed conflict they might have with this friendly country. We presently have a source S, a national of Country T, who is an engineer and an expert on the design of weapon systems such as W. Further, she meets regularly with policy makers in Country T regarding the development of tactical weapon systems. Source S has agreed to inform us about deliberations made by policy makers in T regarding the development of system W.

Does what we know so far about S bear on her competence or credibility?

Source S now reports to us the following information. She says she was just told by a ranking policy maker in Country T that all plans to develop system W have been suspended because it was thought that such development would be far too expensive.

What S has told us seems to be good news, but can we believe what she says? What general kinds of evidence should we consider about the credibility of what she has just told us?

99. Why is the competence of sources of testimonial evidence so important and why is

competence not the same as credibility?

100. The credibility attribute, veracity or truthfulness, is widely discussed and often widely misunderstood or mistakenly attributed. It seems obvious that the veracity attribute is a property of human sources of evidence. It is very hard to imagine a mechanical or electronic sensor attempting to mislead us, willfully or not, in providing a report. Such reports can of course be incorrect but for reasons not involving truthfulness. Provide some examples of the uses and misuses of the attribute veracity.

101. The objectivity attribute of the credibility of a human source is widely overlooked in spite of its importance. What is odd is that lack of objectivity is much discussed in common discourse. We so often hear that we all from time to time believe things because we want to believe them in spite of having little or no evidence for them. One matter of interest concerns the possibility that objectivity is not only a property exclusive to human sensors but also relevant to sensing devices. Provide some examples of the objectivity attribute in various situations.

102. Show how evidential dissonance and selectivity are related in ways that can be inferentially hazardous.

103. It can be argued that of all the inferential issues involved in intelligence analysis, the most important and interesting ones involve the believability of evidence and its sources. Give some reasons why this is so.

104. The most important attribute of the credibility of tangible evidence is its authenticity: Is this evidence what it is claimed to be? Provide some examples of real and demonstrative tangible evidence items that are not authentic.

105. Source T now reports to us the following information. She says she was just told by a ranking policy maker in Country T that all plans to develop system W have been suspended because it was thought that such development would be far too expensive. What S has told us seems to be good news, but can we believe what she says? What general kinds of evidence should we consider about the credibility of what she has just told us?

106. Give some examples from your own experience when you have heard people providing information about which they hedge or equivocate.

107. What inferences might we draw from Omar al-Massari’s refusal to provide us with his laptop computer?

108. Can you provide some examples of mixtures of evidence from your own experience?

109. Human source Y reports to us that the morale among combat troops in Country B is at an

all-time low. We ask Y to give us some specifics. He then reports seeing a classified document at a military installation in B that describes the increasing rate of defections and AWOL (Absent Without Official Leave) over the past year. What kind of evidence is this and how should it be analyzed?

110. It has been noted for years, and by many persons, that intelligence analysts are hampered by not being the persons who assess the credibility or credibility of much of their evidence; this is particularly true of HUMINT evidence. In many cases, sources of HUMINT are under deep cover and their identities are not revealed to analysts. In addition, evidence bearing on the credibility of these HUMINT sources is not always made available to intelligence analysts who will use this HUMINT evidence. Show how this credibility burden on intelligence analysts is made so much heavier when we consider the chains of custody discussed in this section.

111. Intelligence analysts may choose to ignore the heavy burden mentioned in Question **Error! Reference source not found.**. Analysts might prefer to accept versions of HUMINT reports they receive without questioning anything about chains of custody of these reports. Show some of the consequences of failure to assess possible sources of doubt that may lurk in a chain of custody.

112. One thing analysts are trained to do is to assess the consistency of one item of evidence with other items of evidence they may also have. Show how even this consistency assessment is affected by ignoring chains of custody.

113. The veracity, objectivity, and observational sensitivity need to be considered for all of the human sources in our dirty bomb example. Pick a human source in this example and state what kinds of questions you would ask about the veracity, objectivity, and observational sensitivity of the person you have chosen.

5 METHODS OF ASSESING UNCERTAINTY

5.1 General Classes of Probability and Uncertainty

Pierre-Simon Laplace
Probability theory is nothing but common sense reduced to calculation.

We have now considered the major sources of uncertainty related to evidence and have shown how uncertainty arises in chains of reasoning linking evidence to hypotheses we entertain. A major credential of evidence we have only mentioned briefly is its inferential force or weight. As we noted, this credential is always expressed in probabilistic terms but is a source of great controversy. In Figure 37 we showed how the force or weight of evidence concerns the strength of all the credibility and relevance links in our chains of reasoning. But this illustration just concerns the chains of reasoning from two items of evidence. In any intelligence analysis there will be many items of evidence to consider and very many sources of uncertainty or doubt that will be associated with complex arguments linking this mass of evidence to hypotheses at issue in the analysis. It would not be uncommon to be able to identify hundreds or even thousands of sources of doubt arising from masses of evidence being considered.

There are other matters apart from assessing the force or weight of evidence in which uncertainty arises. One way of describing evidence-based reasoning is to say that it involves the revision of probabilistic beliefs about hypotheses based on the evidence we have obtained. To say that we are revising these beliefs suggests that they must have had some initial state in order for them to be revised. The term prior probability is used to indicate the initial conditions of our uncertainty before we consider evidence that begins to emerge. In truth, there has been considerable controversy about prior probabilities and how they can be assessed. When we consider our evidence and its force or weight, we can begin the process of revising these prior beliefs based upon evidence. In the process we revise our prior beliefs to form what are usually termed posterior beliefs, those assessed after we receive and incorporate the evidence we have. But we must take a bit of care concerning the process of belief revision just described.

In the process of discovery or investigation in intelligence analysis and elsewhere, we may have evidence in search of hypotheses and hypotheses in search of evidence all going on at the same time. In other words, it would be quite wrong to suggest that an intelligence analysis always begins with a complete set of hypotheses having been identified. The generation of hypotheses rests on potential evidence we begin to accumulate. Further evidence may suggest new hypotheses or revisions in ones we have generated. In short, evidence not only causes revisions in our probabilistic beliefs about a collection of hypotheses, but it also causes mutations or changes in this collection itself. How this probabilistic belief-revision process proceeds and what its major ingredients are depend vitally upon our view of probability and uncertainty. In the following sections we consider different probability systems.

There is little consensus about how uncertainty should be expressed, combined, and reported. Most analysts have learned in school about the conventional system of probability in which uncertainty is expressed by percentages, or odds and odds ratios, as the only way in which uncertainty can be expressed about evidence-based conclusions. Many analysts will have taken courses in statistics in which the probability of events is estimated based on the relative frequency of their observed occurrence in the past. This approach involves events that are the result of replicable or repeatable processes. In such instances, probability estimation involves events that can be counted. For example, counts can be made of the frequency of occurrence of certain types of vehicles entering or leaving a certain military installation that is under observation. Voting patterns are examined in a certain country to see how strongly in the past voters in a certain country have supported candidates whose interests seem favorable to our own interests. Counts are made of the number of instances of car bomb attacks in Iraq or in other places. There are thus many instances in intelligence analysis in which statistical estimates of probabilities can be made.

However, the problem is that there are many more instances in intelligence analysis in which we have uncertainty about certain past or future events when we will have nothing to count because these events are unique, singular or one-of-a-kind. If they happened in the past, they did so on just one occasion. If they will occur in future, they will do so on just one occasion. In such instances we will have no statistics to back our uncertainty assessments. Examples of uncertainty about these unique or singular events abound in intelligence analysis. Did the Iranians supply the shaped explosive devices that destroyed the two Humvees in Iraq? Was this HUMINT informant or asset truthful or accurate when he identified the person who drove the truck carrying the bomb that destroyed the hotel in Baghdad? Did the Saudis supply the weapons found in the possession of the Sunni militia group? Who will be the successor to Putin in assuming the leadership in Russia? Because there is nothing to count and therefore no statistics to support the probabilities associated with answering such questions, the assessments must be judgmental or subjective in nature. This also means that different analysts may assess these probabilities differently and arrive at different probabilities regarding major answers.

The point here is very simple: all statistical reasoning is probabilistic in nature, but not all probabilistic reasoning is statistical in nature. There are some very interesting but difficult issues concerning the extent to which the concepts and methods so useful in statistical analyses continue to apply in situations in which we have uncertainty but no statistics. There are many issues regarding the assessment, combination, and reporting of uncertainty in these non-statistical situations that are very important but frequently go unrecognized.

Accurately representing uncertainty is not a challenge unique to intelligence analysis. Conclusions reached in law, medicine, science, and history typically are qualified or hedged to convey the extent to which a conclusion falls short of being certain. It is common in any of these professions to qualify one's conclusion.

Some views of probabilistic reasoning are quite different from the conventional view, in which probabilities are thought of only in terms of replicable or countable events. We regard the task of making sense out of masses of evidence to be too rich an intellectual activity for us to expect that any single view of probability will capture all of this richness. Absent one unifying theory of probability, the intelligence analyst must grade and express his or her uncertainties differently based on the context in which that uncertainty arises. In the following sections we review different probability systems. We start with Table 9 which categorizes the alternative views of probability we will discuss as we proceed.

Table 9. Some alternative views of probability.

Enumerative	**Non-Enumerative**
Aleatory (Chances) Relative Frequency and Statistics Bayesian Statistics	Subjective Bayesian Belief Functions Baconian Probability Fuzzy Probability

Intelligence analysts will encounter works by some individuals who will argue that the term probability is only applicable to the enumerative situations just described above. In short, we are always out of luck applying Bayes' rule because of its requirement for subjective prior probabilities; and we are especially out of luck applying probabilities in situations in which the events of concern are singular, unique, or one-of-a-kind, and so we have nothing to count. However, some probabilists argue that subjective judgments are always necessary and that we can assess numerical probabilities provided that they obey the Kolmogorov axioms just described. We will also mention the views of probabilists who argue that not all of the Kolmogorov axioms make sense for subjective judgments of singular or unique events. They can point to the basic sources of uncertainty we discussd in Section 3.5 and argue that probabilities enumerated or judged in accordance with the Kolmogorov axioms cannot capture all of these sources of uncertainty. There are alternative methods for expressing uncertainty that do capture some of these sources but do not rest on the Kolmogorov axioms or his definition of a conditional probability. What follows is a brief account of the essentials four quite different views of uncertainty assessments: Subjective Bayesian view, Belief Functions, Baconian probabilities, and Fuzzy probabilities. We can only provide a look at the essentials of these four views. More extensive comparisons of these four views appear in (Schum, 1994/2001, pp. 200-269). In doing so, we have chosen to focus on what each one has to tell us about what the force or weight of evidence means. Remember that it is in the process of assessing the force or weight of evidence that uncertainties concerning evidence are first expressed.

We begin by discussing two views of probability that involve processes in which we can obtain probabilities or estimates of them by enumerative or counting processes: Aleatory probabilities and Relative Frequency and Statistics.

5.2 Aleatory Probability

The first conception of probability that involve counting operations the aleatory probability. This term has its roots in the Latin term *alea,* meaning chance, game of chance, or devices such as dice involved in games of chance. Games of chance have two important ground rules:

- There is a finite number n(S) of possible outcomes.
- All outcomes in S are assumed to have equal probability.

For example, in a game involving a pair of fair six-sided dice, where we roll and add the two numbers showing up, there are 36 ways in which the numbers showing up will have sums between 2 and 12, inclusive. So, in this case, n(S) = 36. Suppose you wish to determine the probability that you will roll a 7 on a single throw of these dice. There are exactly 6 ways in which this can happen. If E = “the sum of the numbers is 7”, then n(E) = 6. The probability of E, P(E), is simply determined by dividing n(E) by n(S), which in this example is P(E) = 6/36 = 1/6. So, aleatory probabilities are always determined by dividing n(E) by n(S), whatever E and S are, as long as E is a subset of S.

We can dismiss aleatory probability as not being interesting in intelligence analysis since there seem to be no instances in which the two aleatory ground rules will apply. However, we should note that there are frequent instances in which analysts may use the term “chance“, when it may not be appropriate. For example, here is an analyst who says, “The chances are 9 in 10 (probability equals 0.9) that country Green is supplying arms to insurgents in country Orange.” This judgment cannot have arisen by any counting operation in which the two ground rules for aleatory probabilities apply. One rather unfortunate occurrence is that most people are initially introduced to probability theory by use of games of chance to illustrate various concepts. People then often believe that these concepts occur and retain the same meaning when they are used in entirely different contexts in which uncertainty assessments are required.

5.3 Relative Frequency and Statistics

Another way of assessing probabilities involves the many situations in which aleatory ground rules will not apply but when we do have empirical methods at hand to estimate probabilities. These situations arise when we have replicable or repeatable processes in which we can count the number of times events have occurred in the past. Suppose that, employing a defensible method for gathering information about the number of times event E has occurred, we determine the relative frequency of an occurrence of E by counting the number of times E has occurred, n(E), and then dividing this number by N, where N is the number of observations we have made, or the sample size we have taken. In this case, the relative frequency of E, f(E), equals n(E)/N. You recognize that this is a statistical process that can be performed in many situations, provided that we assume processes that are replicable or repeatable. It is true, of course, that a relative frequency f(E) is just an estimate of the true probability of E, P(E). The reason, of course, is that the number N of observations we have made is always less than the total number of observations that could be made. In some cases there may be an infinite number of possible observations. If you have had a course in probability theory you will remember that there are several formal statements, called the laws of large numbers, for showing how f(E) approaches P(E) when N is made larger and larger.

Probability theory presents an interesting paradox. It has a very long history, but a very short past. There is abundant evidence that people as far back as Paleolithic times used objects resembling dice either for gambling or, more likely, to foretell the future (David, 1962). But attempts to calculate probabilities only date back to the 1600s, and the first attempt to develop a theory of mathematical probability only dates back to 1933 in the work of Russian A. N. Kolmogorov (1933). Kolmogorov was the first to put probability on an axiomatic basis. The three basic axioms he proposed are the following ones:

Axiom 1: For any event E, P(E) ≥ 0.
Axiom 2: If an event is sure or certain to occur, which we label S, P(S) = 1.0.
Axiom 3: If two events, E and F, cannot occur together, or are mutually exclusive, the probability that one or the other of these events occurring is the sum of their separate probabilities. In symbols, P(E or F) = P(E) + P(F).

All Axiom 1 says is that probabilities are never negative. Axioms 1 and 2, taken together, mean that probabilities are numbers between 0 and 1. An event having 0 probability is commonly called an "impossible event." Axiom 3 is called the additivity axiom and it holds for any number of mutually exclusive events.

Certain transformations of Kolmogorov's probabilities are entirely permissible and are often used. One common form involves odds. The odds of event E occurring to its not occurring (written as E^C), which we label Odds(E, E^C), is determined by Odds(E, E^C) = P(E)/P(E^C) = P(E)/(1 - P(E)). For any two mutually exclusive events E and F, the odds of E to F, Odds(E, F), are given by Odds(E, F) = P(E)/P(F). Numerical odds scales range from zero to an unlimited upper value. The person who said the chances that country Green is supplying weapons to insurgents in country Orange is 9 in 10, might better have said that the odds favoring Green supplying the weapons, to their not supplying weapons, are 9 to 1.

What is very interesting, but not always recognized, is that Kolmogorov had only enumerative probability in mind when he settled on the above three axioms. He makes this clear in his 1933 book and in his later writings (Kolmogorov, 1969). It is easily shown that both aleatory probabilities and relative frequencies obey these three axioms. But Kolmogorov went an important step further in defining conditional probabilities that are necessary to show how the probability of an event may change as we learn new information. He defined the probability of event E, given or conditional upon some other event F, as: P(E given F) = P(E and F)/P(F), assuming that P(F) is not zero. P(E given F) is also written as P(E|F). He chose this particular definition since conditional probabilities, so defined, will also obey the three axioms just mentioned. In other words, we do not need any new axioms for conditional probabilities.

Now comes a very important concept you may have heard about. It is called Bayes’ rule and results directly from applying the definition of the conditional probability. From P(E* and H) = P(H and E*) you obtain P(E*|F) P(F) = P(H|E*)P(E*). This can then be written as shown in Figure 44. This rule is named after the English clergyman, the Reverend Thomas Bayes (1702 – 1761), who first saw the essentials of a rule for revising probabilities of hypotheses, based on new evidence (Dale, 2003). He had written a paper describing his derivation and use of this rule but he never published it; this paper was found in his desk after he died in 1761 by Richard Price, the executor of Bayes’ will. Price realized the importance of Bayes’ paper and recommended it for publication in the *Transactions of the Royal Society*, in which it appeared in 1763. He rightly viewed Bayes’ rule as the first canon or rule for inductive or probabilistic reasoning. Bayes’ rule follows directly from Kolmogorov’s three axioms and his definition of a conditional probability, and is entirely uncontroversial as far as its derivation is concerned. But this rule has always been a source of controversy on other grounds. The reason is that it requires us to say how probable a hypothesis is before we have gathered evidence that will possibly allow us to revise this probability. In short, we need prior probabilities on hypotheses in order to revise them, when they become posterior probabilities. Persons wedded to enumerative conceptions of probability say we can never have prior probabilities of hypotheses since, in advance of data collection we have nothing to count. Statisticians are still divided today about whether it makes sense to use Bayes’ rule in statistical inferences. Some statisticians argue that initial prior probabilities could only be assessed subjectively and that any subjective assessments have no place in any area that calls itself scientific. Bayes’ rule says that if we are to talk about probability revisions in our beliefs, based on evidence, we have to say where these beliefs were before we obtained the evidence.

It is time for us to consider views of probability in situations where we will have nothing to count, either a priori or anywhere else.

5.4 Subjective Bayesian View

We refer to our first non-enumerative view as an epistemic view, since it assumes that probabilities in any case are based on some kind of knowledge, whatever form it may take. In short, probabilities are the result of informed judgments.

5.4.1 Likelihood Ratios

Many statisticians now favor the use of Bayes' rule in enumerative or frequentistic situations and have no objection to subjective assessments of prior probabilities. Bayes' rule requires the assessment of two basic probabilistic ingredients: prior probabilities on hypotheses, and likelihoods or their ratios. As we will illustrate, it is these likelihoods and their ratios that are the ingredients of Bayes' rule that concern the inferential force of evidence. Furthermore, many persons favor the use of Bayes' rule for combining subjective assessments of all the prior and likelihood ingredients of Bayes' rule. *But what these persons require is that these assessments be entirely consistent with Kolmogorov's three axioms and his definition of conditional probabilities we noted above*. Since Bayes' rule rests on these axioms and definition, we must adhere to them in order to say that our assessment process is coherent or consistent.

Here is a simple explanation of how ratios of likelihoods express the force of evidence in Bayes' rule. Suppose we have two hypotheses H and $\mathrm{H^c}$ (the complement of H, i.e., $\neg H$), and a single item of evidence $\mathrm{E^*}$ saying that event E occurred. What we are interested in determining are the posterior probabilities: $\mathrm{P(H|E^*)}$ and $\mathrm{P}(\neg H \,\big|\, \mathrm{E^*})$. Using the Bayes' rule from Figure 44, we can express these posterior probabilities as:

$$\mathrm{P(H|E^*)} = \frac{\mathrm{P(E^*|H)P(H)}}{\mathrm{P(E^*)}} \qquad \mathrm{P}(\neg H\,|\mathrm{E^*}) = \frac{\mathrm{P}(\mathrm{E^*}|\neg H\,)\mathrm{P}(\neg H\,)}{\mathrm{P(E^*)}}$$

The next step is to divide $\mathrm{P(H|E^*)}$ by $\mathrm{P}(\neg H\,|\mathrm{E^*})$ that will produce three ratios; in the process the term $\mathrm{P(E^*)}$ will drop out. Here are the three ratios that result:

$$\frac{\mathrm{P(H|E^*)}}{\mathrm{P}(\neg H\,|\mathrm{E^*})} = \frac{\mathrm{P(H)}}{\mathrm{P}(\neg H\,)} \frac{\mathrm{P(E^*|H)}}{\mathrm{P(E^*|H^c)}}$$

The left-hand ratio $\frac{\mathrm{P(H|E^*)}}{\mathrm{P}(\neg H\,|\mathrm{E^*})}$ is called the posterior odds of H to $\neg H$, given evidence $\mathrm{E^*}$. In symbols we can express this ratio as: $\mathrm{Odds}(\mathrm{H}{:}\neg H\,|\mathrm{E^*})$. This first ratio on the right, $\frac{\mathrm{P(H)}}{\mathrm{P}(\neg H\,)}$, is called the prior odds of H to $\neg H$. In symbols we can express this ratio as: $\mathrm{Odds}(\mathrm{H}{:}\neg H)$. The

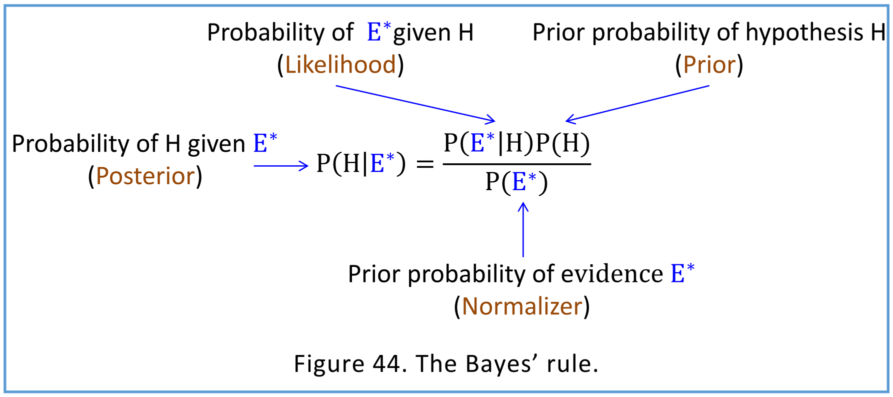
Probability of E* given H
(Likelihood)
Prior probability of hypothesis H
(Prior)
Probability of H given E*
(Posterior)
P(H|E*) = P(E*|H)P(H) / P(E*)
Prior probability of evidence E*
(Normalizer)

Figure 44. The Bayes' rule.

remaining ratio on the right, $\frac{\mathrm{P(E^*|H)}}{\mathrm{P(E^*|\neg \mathit{H}\ )}}$ is called the likelihood ratio for evidence E^*; we give this ratio the symbol $\mathrm{L}_{\mathrm{E}^*}$. In terms of these three ratios, Bayes' rule applied to this situation can be written simply as:

$$\mathrm{Odds(H{:}\neg \mathit{H}\ |E^*) = Odds(H{:}\neg \mathit{H})L_{E^*}}$$

This simple version of Bayes' rule is called the odds-likelihood ratio form. It is also called, somewhat unkindly, "idiots" Bayes. If we divide both sides of this equation by the prior odds, $\mathrm{Odds(H{:}\neg \mathit{H})}$, we observe that the likelihood ratio $\mathrm{L}_{\mathrm{E}^*}$ is simply the ratio of posterior odds to prior odds of H to $\neg H$. This likelihood ratio shows us how much, and in what direction (toward H or $\neg H$), our evidence E^* has caused us to change our beliefs toward H or toward $\neg H$ from what they were before we obtained evidence E^*. In short, likelihood ratios grade the force of evidence in Bayesian analyses. But this is the simplest case possible. Likelihood ratios become much more complex when we attempt to show how relevant E^* is to H and $\mathrm{H^c}$, and to capture the credibility of the source of evidence E^*.

Example 35. Here is an example of how likelihoods and their ratios provide a method for grading the Bayesian force of an item of evidence. This is an example of a situation involving a singular evidence item where we have nothing to count. Suppose you are an analyst whose interest concerns whether or not the Greens are supplying parts necessary for the construction of shaped explosive devices to a certain insurgent militia group in the neighboring country Orange. Thus you are entertaining the following binary hypotheses:

H: The Greens are supplying parts necessary for the construction of shaped explosive devices.

H^c:The Greens are not supplying parts necessary for the construction of shaped explosive devices.

Suppose you believe, before you have any evidence, that the prior probability of H is $P(H) = 0.20$. Because you must obey the rules for enumerative probabilities, you must also say that $P(H^c) = 0.80$. This follows from the third axiom we discussed in Section 5.2. So, your prior odds on H relative to H^c, have a value $Odds_0 = \frac{P(H)}{P(H^c)} = \frac{0.20}{0.80} = \frac{1}{4}$.

Your first item of evidence E_1^* is a report that a member of the Green’s military was captured less than one kilometer away from a location in Orange at which parts necessary for the construction of these shaped explosives were found. You ask yourself how likely is this evidence E_1^* if H were true, and how likely is this evidence E_1^*, if H were not true. Suppose you say that $P(E_1^*|H) = 0.80$ and $P(E_1^*|H^c) = 0.10$. You are saying that this evidence is eight times more likely if H were true than if H were not true. So, your likelihood ratio $L_{E_1^*}\ = \frac{P(E_1^*|H)}{P(E_1^*|H^c)} = \frac{0.8}{0.1} = 8$. You now have all the ingredients necessary in Bayes' rule to determine the posterior odds

$Odds_1 = \frac{P(E_1^*|H)}{P(E_1^*|H^c)}$ and posterior probability of hypothesis H, $P(H|E_1^*)$. In this case:

$$Odds_1 = Odds_0 \times \ L_{E_1^*} = \frac{1}{4} \times 8 = 2.$$

This means that you now believe the posterior odds favoring H over H^c are two to one. But you started by believing that the prior odds of H to H^c were one in four, so the evidence E_1^* changed your belief by a factor of 8, which is just what $L_{E_1^*}$ says. As we have noted, the ingredient in Bayes' rule that indicates the force or weight of evidence are likelihood ratios.

When we have just one hypothesis together with its complement, we can easily convert odds to probabilities using the familiar rule: $Probability = \frac{Odds}{1+Odds}$. This rule follows from the fact that we have defined $Odds$ to be $\frac{P(H)}{P(H^c)} = \frac{P(H)}{1-P(H)}$. Suppose we wish to determine the posterior probability $P(H|E_1^*)$ from $Odds_1$. In this case we have $(H|E_1^*) = \frac{Odds_1}{1+Odds_1} = \frac{2}{1+2} = \frac{2}{3} = 0.67$. So, in terms of probabilities, evidence E_1^* caused you to increase the probability of H by 0.47.

There are various difficulties associated with grading the Bayesian force of evidence in terms of difference between prior and posterior probabilities. The next example shows what the difficulties are.

Example 36. Suppose a critic argues that your assessment of the prior probability of H was foolishly low; she argues that the prior probability of H is more like $P(H) = 0.75$. So her prior odds are $Odds_0 = \frac{0.75}{0.25} = 3$. But she agrees entirely with your assessment of the likelihood ratio of E_1^* being $L_{E_1^*} = 8$. From Bayes rule we now have:

$$Odds_1 = Odds_0 \times \ L_{E_1^*} = 3 \times 8 = 24.$$

Notice first that the ratio of posterior to prior odds is still 8:1. But now consider the posterior probability of H, it is $P(H|E_1^*) = \frac{24}{1+24} = 0.96$. So the difference between her posterior and prior probabilities is now just 0.96 - 0.75 = 0.21, a much smaller difference than it was using your prior probabilities. But the ratio of posterior odds to prior odds remains the same in her case and in yours; i.e., her posterior to prior odds ratio is $\frac{24}{3} = 8$. The likelihood ratio indicates the same force in both situations when the problem is construed in odds. This is the major reason why the odds-likelihood ratio form of Bayes' rule is so helpful and informative about what the force or weight of evidence means in Bayesian terms.

But you have only one item of evidence so far; what happens when you have additional items of evidence? Suppose you now have a new item of evidence that we label E_2^*. How do you combine this new evidence to revise or update your posterior Odds$_1$, based on this new item of evidence? Suppose this new item of evidence E_2^* is a report that a fragment of a shaped explosive device was found on a road leading to the Sand city, Orange. This fragment carries a serial number that we learn identifies this device as having been made in a munitions factory outside the capital of Green. To illustrate how Bayes' rule allows us to combine our uncertainties based on new evidence we must take a short detour to illustrate a most important concept in the Bayesian view of evidence-based reasoning; this concept is termed conditional dependence or, sometimes, conditional nonindependence. This concept allows us to capture an amazing array of complexities or subtleties in evidence. Here is a brief account of conditional dependence.

The probabilistic independence of two events A and B means that P(A|B) = P(A), or equivalently that P(B|A) = P(B). That is, the occurrence of B has no influence on the probability of A, and vice versa. Events A and B are then nonindependent if these equations do not apply. But there are many situations in which the independence of two events depends on some other event, say event C. It might be the case that events A and B are independent only if C is true, in which case we can say that P(A|BC) = P(A|C). This equality implies a product rule for conditional probabilities. If P(A|BC) = P(A|C), this also means that P(AB|C) = P(A|C)P(B|C). Either expression says that A and B are independent, given event C. Now, what is of interest are situations in which A and B are not independent, given event C. In this case we have P(AB|C) ≠ P(A|C)P(B|C). This is a probabilistic expression of the fact that considering two events jointly in light of event C, as we do in considering P(AB|C), means something different than they would mean if we considered them independently or separately as we do in considering P(A|C) and P(B|C).

We now give an example of the importance of conditional dependence. Suppose we have evidence for events A and B. In symbols, we have evidence items A* and B*. These two items are synergistic in nature, given event C, if P(A*B*|C) > P(A*|C)P(B*|C). This means that these two items have greater probability, given C and when taken together, than they would have if we considered them separately or independently, given event C. What is interesting is that, when P(A*B*|C) > P(A*|C)P(B*|C), this also means that P(B*|A*C) > P(B*|C). What this means is that B* has greater probability, given C, when we also consider A*, than it would have if we did not consider A*.

We now turn again to your inference concerning whether the Greens are supplying shaped explosive devices to insurgent groups in Orange. So far, you have determined the posterior odds

$$Odds_1 = Odds_0 \times \ L_{E_1^*} = \frac{1}{4} \times 8 = 2.$$

But now you have the new evidence E_2^* saying that a fragment of a shaped explosive device found on a road near the Sand city in Orange carried a serial number that identifies this device having been made in a munitions factory outside the capital of Green. The question is how do you combine evidence items E_1^* and E_2^* together in Bayes' rule?

What we must now do is to describe the recursive nature of Bayes' rule. What this says is that our new posterior odds, $Odds_2$, depends on our old posterior odds, $Odds_1$. One way of saying this is to say that yesterday's posterior odds become today's prior odds in light of today's new evidence. In short, we never begin from scratch with each new item of evidence but incorporate old evidence with new evidence. The first thing we must do is to carefully define our new posterior odds, $Odds_2$. Since we now have two items of evidence and so $Odds_2 = \frac{P(H|E_1^* E_2^*)}{P(H^c|E_1^* E_2^*)}$. When we decompose these two conditional probabilities using the rules provided above, we obtain:

$$Odds_2 = \frac{P(H|E_1^* E_2^*)}{P(H^c|E_1^* E_2^*)} = \frac{P(H)}{P(H^c)} \frac{P(E_1^*|H)}{P(E_1^*|H^c)} \frac{P(E_2^*|E_1^* H)}{P(E_2^*|E_1^* H^c)}$$

The first thing to notice is that the first two ratios on the right-hand side of this equation form the old posterior odds, Odds_1 that just concerns E_1^*, our first item of evidence. This illustrates the recursive nature of Bayes' rule. The posterior odds for the first item of evidence become the prior odds in determining the posterior odds for the combined evidence E_1^* and E_2^*.

But now we come to the most interesting ingredient of Bayes' rule, it concerns the likelihood ratio $\frac{P(E_2^*|E_1^* H)}{P(E_2^*|E_1^* H^c)}$. Here is where our discussion of conditional independence and dependence becomes so important. What Bayes' rule requires us to answer is whether E_1^* and E_2^* are independent, given H, and given H^c. Another way of posing this question is to ask: Does the force of evidence E_2^* on H and H^c, depend on our first item of evidence E_1^*? Suppose that E_1^* acts to increase the probability of E_2^*, given H, and also acts to decrease the probability of E_2^*, given H^c. In this case our two evidence items would be synergistic in their effects; taken together they point more strongly toward H than they would do if they were considered separately or independently. If they were independent under both H and H^c, then we would have: $\frac{P(E_2^*|E_1^* H)}{P(E_2^*|E_1^* H^c)} = \frac{P(E_2^*|H)}{P(E_2^*|H^c)}$.

Here are two examples illustrating matters we just discussed concerning what Bayes' rule requires concerning the combination of evidence.

Example 37. Suppose in the example involving whether the Greens are supplying shaped explosive devices to the insurgents in Orange, our analyst first decides that there's no connection between E_1^* and E_2^*, as defined above. He says to himself, "I don't see the connection between a Green being found a kilometer away from a shaped explosive device and a fragment of another such device bearing a number showing that the device was made in some factory no one ever heard of in Green. So, I am going to judge $P(E_2^*|H) = 0.6$, and $P(E_2^*|H^c) = 0.2$. My reasoning here is that the Oranges could have purchased the parts for this device from some other country that had purchased these parts from Green. So, the Oranges made this device themselves." So now we can use these likelihood assessments for evidence E_2^* and determine:

$$Odds_2 = Odds_0 \times L_{E_1^*} \times L_{E_2^*} = \frac{1}{4} \times 8 \times 3 = 6.$$

In this case the analyst is saying that the posterior probability of H, given E_1^* and E_2^* is $\frac{6}{1+6} =$

0.86.

Example 38. Our critic appears again and this time she argues strongly that E_1^* and E_2^* are not at all independent given either H or H^c. She says, "I can't believe you don't see any connection between E_2^* and E_1^*. If there was a member of the Green military just a short distance from where a shaped explosive device was found, does this not suggest to you that we will eventually find some of these devices that can be identified as having been made in Green, which we have just discovered in E_2^*? I'm going to judge $P(E_2^*|E_1^*H) = 0.90$ and $P(E_2^*|E_1^*H^c) = 0.05$. Now, let's see what Bayes' rule says your posterior odds and probability should be using my assessments." So we calculate:

$$Odds_2 = Odds_0 \times \ L_{E_1^*} \times \ L_{E_2^*} = \frac{1}{4} \times 8 \times 18 = 36.$$

In this case the posterior probability of H, given E_1^* and E_2^* is $\frac{36}{37} = 0.97$. What the critic has done is to say that because of the conditional dependence of E_1^* and E_2^*, given H and given H^c, E_2^* has six times as much force or weight when we consider E_1^* than it would have if we did not consider E_1^*.

Whose view of the weight of evidence in the above two examples makes the most sense to you? Only you can answer this question.

These examples using "idiots" Bayes are simple because we have not constructed any arguments showing the believability of E_1^* and E_2^* and showing their relevance on our hypotheses H and H^c. This is where matters become very complex indeed. The next section discusses the use of Bayesian Networks to perform such analyses.

5.4.2 Bayesian Networks

There is a variety of software systems, referred to as Bayesian Networks Systems, that are designed to perform probabilistic analyses in complex inference networks, such as those discussed in the previous sections of this book. An illustration of the use of Bayesian networks for evidence-based analysis of the hypothesis “The cesium-137 canister is missing from the STEMQ warehouse”, analyzed in Section 0. We are going to analyze this hypothesis using a Bayesian network.

Constructing the Argument Structure

Bayesian Network analysis goes from the top-down, from hypotheses to evidence. Following are the steps necessary in constructing a Bayesian Network analysis.

Step 1: Construct the chain of questions or the argument structure to be analyzed. In this case

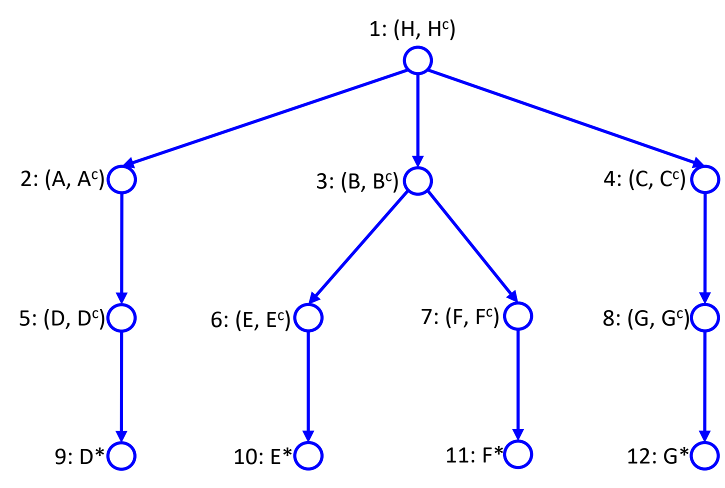

Figure 45. Bayesian network for the hypothesis H: The cesium-137 canister is missing from the STEMQ warehouse.

we are using the argumentation from Figure 12, resulting in the Bayesian network from Figure 45.

Forming the Key List

Step 2: Form the key list for the chart in 0 to identify its ingredients, as shown in Table 10. Notice that we must describe an event and its complement at all stages above the evidence. This is necessary for probabilistic analyses. Notice also that we designate the items of evidence by D*, E*, F*, G*, and the events described by these items of evidence as D, E, F, G, respectively. But just because evidence D*, for example, says that the event D has happened, does not mean that D has actually happened. At issues here is the believability of the source of that item of evidence. Therefore we have to consider both the event D and its negation or complement D^c. Notice that we must describe an event (e.g., A) and its complement (i.e., A^c) at all stages above the evidence. This is necessary for probabilistic analyses.

Table 10. Key list for Figure 45.

1: H = The cesium-137 canister is missing from the STEMQ warehouse.
H^C = The cesium-137 canister is not missing from the STEMQ warehouse.
2: A = The cesium-137 canister was in the STEMQ warehouse before being reported as missing.
A^C = The cesium-137 canister was not in the STEMQ warehouse before being reported as missing.
3: B = The cesium-137 canister is no longer in the STEMQ warehouse.
B^C = The cesium-137 canister is still in the STEMQ warehouse.
4: C = No one has checked the cesium-137 canister out of the STEMQ warehouse.
C^C = Someone has checked the cesium-137 canister out of the STEMQ warehouse.
5: D = The cesium-137 canister is registered as being in the STEMQ warehouse.
D^C = The cesium-137 canister is not registered as being in the STEMQ warehouse.
6: E = A canister containing cesium-137 was missing from the STEMQ warehouse in Baltimore, MD.
E^C = A canister containing cesium-137 was not missing from the STEMQ warehouse in Baltimore, MD.
7: F = The cesium-137 canister is not located anywhere in the hazardous materials locker.
F^C = The cesium-137 canister is located somewhere in the hazardous materials locker.
8: G = No one at the STEMQ Company had checked the cesium-137 canister out.
G^C = Someone at the STEMQ Company had checked the cesium-137 canister out.
9: D* = E002-Ralph testimony to D in 5.
10: E* = E001-Willard testimony to E in 6.
11: F* = E004-Ralph testimony to F in 7.
12: G* = E003-Ralph testimony to G in 8.

Finally, notice that the construction of our network chart and its accompanying key list proceeds concurrently. In other words, we must say what all the ingredients of our chart mean when we are constructing it. Our inference network contains various lines of argument linking our evidence to hypotheses we are considering. This task involves imaginative reasoning as we have discussed in this book. Our Bayesian Network system cannot perform this imaginative reasoning as we will note in a minute.

Identifying the Likelihoods and Prior Probabilities

Step 3. The next step is to identify the probabilities Bayes' Rule says are required in the network as constructed. These are listed in Table11and explained in the following.

Table11.Required probabilities for the Bayesian network in Figure 45
(assuming no conditional dependencies).

1: $P(H)$, $P(H^C)$. Priors

2: $P(A|H)$, $P(A|H^C)$. Likelihoods.

3: $P(B|H)$, $P(B|H^C)$. Likelihoods.

4: $P(C|H)$, $P(C|H^C)$. Likelihoods.

5: $P(D|A)$, $P(D|A^C)$. Likelihoods.

6: $P(E|B)$, $P(E|B^C)$. Likelihoods.

7: $P(F|B)$, $P(F|B^C)$. Likelihoods.

8: $P(G|C)$, $P(G|C^C)$. Likelihoods.

9: $P(D^*|D)$ = Ralph's Hit probability in reporting D (also likelihood, but related to Ralph's believability).
$P(D^*|D^C)$ = Ralph's False Positive probability in reporting D
(also likelihood, but related to Ralph's believability).

10: $P(E^*|E)$ = Willard's Hit probability in reporting E (also likelihood, but related to Willard's believability).
$P(E^*|E^C)$ = Willard's False Positive probability in reporting E
(also likelihood, but related to Willard's believability).

11: $P(F^*|F)$ = Ralph's Hit probability in reporting F (also likelihood, but related to Ralph's believability).
$P(F^*|F^C)$ = Ralph's False Positive probability in reporting F
(also likelihood, but related to Ralph's believability).

12: $P(G^*|G)$ = Ralph's Hit probability in reporting G (also likelihood, but related to Ralph's believability).
$P(G^*|G^C)$ = Ralph's False Positive probability in reporting G
(also likelihood, but related to Ralph's believability).

The first thing to note here is that Bayes' rule requires two basic forms of probabilistic ingredients: priors and likelihoods. The ones listed above are particular to the inference network shown in Step 1. Look again at this network and first observe that the top level node #1 contains our major question in the form of hypotheses: H and H^C. The question is: Is the cesium-137 canister missing from the STEMQ warehouse? The two possible answers are yes (H) or no (H^C). Then notice that this node contains no parents above it. What this means is that, in our network as shown, there are no questions or sources of influence on our basic hypotheses. Actually, there are but we have not included them to make this example simple. There being no sources of influence on H and H^C means that we must assign prior probabilities to the hypotheses at this node indicating the strength of our beliefs in them before we consider any evidence relevant to them.

All the other ingredients, called likelihoods, allow us to assess the strength of our arguments on these questions or hypotheses based on the evidence we have. What we are trying to determine are the posterior probabilities of H and H^C, based on the four items of evidence we have. The bottom four pairs of likelihoods shown above concern the believability of the sources (Ralph and Willard) of the four items of evidence we have. In any inference network, believability considerations always form the basic foundations of our arguments. What Bayes' rule shows us is how to combine all these prior and likelihood ingredients in order to determine the posterior probabilities of H an H^C, based on the evidence we have.

Here's a problem associated with computer-assisted Bayesian Network analyses regardless of the software system you are using. We cannot show you the exact equations the computer is using to calculate posterior probabilities for H and H^C, under various conditions. The reason is that the equations are always buried below the surface and the computer does not reveal them to us. So we have to be confident that the computer knows which equations to use under various conditions that we specify. Being unable to see the exact equations being used often causes difficulties for us as we try to explain the results of our Bayesian Networks calculations.

Given an inference network structure such as the one we are considering, our Bayesian Network system will tell us what probabilities we need in order for us to determine posterior probabilities for major questions or hypotheses on this network, based on the evidence we have. No Bayesian Networks system will tell us how to ask these questions or construct this network. Remember that the construction of an inference network is an imaginative reasoning task followed by critical reasoning in which we evaluate the logical consistency of the arguments we have constructed on this network. Now, the network we have constructed above has a hierarchical structure. The top-level question at node 1 suggests three questions at the next lower level. Then these three second-tier questions suggest one or more questions at the third tier that can be answered by the four items of evidence at the bottom.

Now, notice that we could have made the network in 0 more complex by adding additional links having the following very general meaning. Basically, we have to ask whether the answers we could get to one question influence the probability of getting answers to a different question. Here's an example. Suppose the answer to the question at # 2 is A: "The cesium-137 canister was in the STEMQ warehouse before being reported as missing." If so, then we also ought to ask whether the occurrence of this event would influence the probability of answers, B and B^C, at node #3. Specifically, if A: "The cesium-137 canister was in the STEMQ warehouse before being reported as missing", does this make B: "The cesium-137 canister is no longer in the STEMQ warehouse" more or less likely? In other words, are the questions we are asking dependent in various ways?

What we have not done in constructing this network is to include what are termed conditional dependencies (also called conditional nonindependencies). The existence of conditional dependencies allows us to capture a wide assortment of evidential and inferential subtleties or complexities in a probabilistic analysis. Here is a very brief description of what conditional dependence involves. First, at a very basic level, two or more events taken together may mean something quite different than they would do if considered separately or independently. Second, if the first is true, there is an equivalent statement we could make. The second of two events means something different if we took the first event into account than it would if we failed to take the first event into account. Following is an example illustrating conditional dependence of the events we have just considered.

Here is our basic question or hypothesis: H = "The cesium-137 canister is missing from the STEMQ warehouse." Two events relevant to inferences about H are: A ="The cesium-137 canister was in the STEMQ warehouse before being reported as missing;" and B = "The cesium-137 canister is no longer in the STEMQ warehouse." Here's the issue we could address: Are events A and B more/less forceful in inferences about H if we took them together or jointly than they would be if we considered them separately or independently. If A and B mean more in inferring H if we took A and B together, we would say that A and B are dependent, conditional on H. To capture this dependence involving A and B, we would draw an arc or arrow linking nodes 2 and 3 in the network from Figure 45.

If we did so, our Bayesian Networks system would recognize this and inform us about the new and additional probabilistic assessments we must make. As you see there are other such dependence linkages we might consider such as ones between nodes 3 and 4 and between 6 and 7. The topic of conditional dependence is complex but very important in Bayesian analyses of any sort and requires careful study by anyone contemplating such analyses. This is one of the high points about Bayesian analyses; they can capture a wide assortment of evidential complexities or subtleties, more than any other probabilistic system. For more on conditional dependencies, see Chapter 7 in (Schum, 1994/2001).

Using the Bayesian Network

Step 4. This fourth step involves putting a Bayesian Networks analysis to work in the probabilistic hedging of conclusions regarding inferences about our major hypotheses. As we have discussed, given an inference network structure such as the one for Figure 45, a Bayesian Networks system will show us what probabilities we must assess in order to calculate posterior probabilities for hypotheses of interest. We might say first that the network structure defines the major plot of stories we can tell about the inference of concern to us. When we assess the probabilities Bayes' rule says are required, we breathe life into our story plot and consider the "actors" in this story and their roles in it. By "actors" we mean the events of concern in our inference task. Of course we can give the actors different identities depending on the probabilities we assign to them. This can be done in an unlimited number of ways and so we can tell an unlimited number of different stories about the same inference network. But, once we have decided upon a specific set of probabilities for the actors in a story, a calculation involving Bayes' rule tells us how this story ends in terms of the posterior probabilities for hypotheses of interest. In the following we will tell five different stories about the Figure 45 network and offer an explanation of them. In doing so we are essentially doing what is called a sensitivity analysis.

Given the Bayesian Networks analyses here, there is an infinite number of stories that might be told, one for each possible combination of probabilistic ingredients. We started in Story 1 by supposing that H and H^C are equally likely a priori. Then, we assessed the first eight likelihood pairs with values that seemed to make sense. We next assessed hits and false-positive likelihoods for Ralph in his testimony of D*, F*, and G*. We picture Ralph as being a very credible source of evidence; his h/f ratio is 95: 1. But we have pictured Willard as being less credible than Ralph; Willard's h/f ratio is only 7:1. The row labeled H Posterior (all evidence) is how Bayes' rule says Story 1 should end if we considered all four items of evidence we have. The ERGO Bayesian Networks system (developed by Noetic Systems Inc.) calculated posterior probabilities: P(H|all evid) = 0.92, and P(H^C| all evid) = 0.08. But then we asked: How would this story end if we only had Willard's testimony? The ERGO system lets us determine posterior probabilities of major hypotheses for various combinations of evidence. The next row shows P(H|Only Willard) = 0.73; and P(H^C|Only Willard) = 0.27. The last row shows what happens when we leave out Willard's evidence; the posterior probability of H would be 0.90.

As you see, Story 2 has the same ingredients as Story 1 with the exception of Willard's hits and false positives for his testimony E*. We thought what would happen if we made Willard at least as credible as Ralph. What happens here is that only very small increases occur in the posterior of H when we consider all the evidence or when we ignore Willard's evidence. These increases don't show up in our table since we are only carrying these numbers to two places. The only noticeable change occurs when we only consider Willard's evidence. The posterior of H increases over what it was in Story 1.

Now, in Story 3, the only ingredient change involves Willard again. But this time we have supposed that Willard may not be truthful. Notice that his false-positive probability is 8 times larger than his hit probability. This results in decreases in the posterior probability of H when we consider all the evidence. When we consider only Willard's evidence, the posterior on H^C is 0.73 and on H it is only 0.27. What this says is that, if Willard is lying, we can come to believe the opposite of what he tells us. Finally, if we ignore Willard our beliefs about the posterior probability of H don't change over what they were in Story 2.

Table 12. Network stories.

Probabilities	**Story 1**	**Story 2**	**Story 3**	**Story 4**	**Story 5**
$P(H); P(H^C)$	0.5; 0.5	0.5; 0.5	0.5; 0.5	**0.25; 0.75**	**0.25; 0.75**
$P(A\|H); P(A\|H^C)$	0.6; 0.15	0.6; 0.15	0.6; 0.15	0.6; 0.15	0.6; 0.15
$P(B\|H); P(B\|H^C)$	0.8; 0.1	0.8; 0.1	0.8; 0.1	0.8; 0.1	0.8; 0.1
$P(C\|H); P(C\|H^C)$	0.05; 0.8	0.05; 0.8	0.05; 0.8	0.05; 0.8	0.05; 0.8
$P(D\|A); P(D\|A^C)$	0.95; 0.1	0.95; 0.1	0.95; 0.1	0.95; 0.1	0.95; 0.1
$P(E\|B); P(E\|B^C)$	0.95; 0.1	0.95; 0.1	0.95; 0.1	0.95; 0.1	0.95; 0.1
$P(F\|B); P(F\|B^C)$	0.98; 0.03	0.98; 0.03	0.98; 0.03	0.98; 0.03	0.98; 0.03
$P(G\|C); P(G\|C^C)$	0.8; 0.4	0.8; 0.4	0.8; 0.4	0.8; 0.4	0.8; 0.4
$P(D^*\|D); P(D^*\|D^C)$	0.95; 0.01	0.95; 0.01	0.95; 0.01	0.95; 0.01	0.95; 0.01
$P(E^*\|E); P(E^*\|E^C)$	0.7; 0.1	**0.98; 0.01**	**0.10; 0.80**	**0.10; 0.80**	**0.01; 0.99**
$P(F^*\|F); P(F^*\|F^{C)}$	0.95; 0.01	0.95; 0.01	0.95; 0.01	0.95; 0.01	0.95; 0.01
$P(G^*\|G); P(G^*\|G^C)$	0.95; 0.01	0.95; 0.01	0.95; 0.01	0.95; 0.01	0.95; 0.01
H Posterior(all evd)	0.92; 0.08	0.92; 0.08	0.82; 0.18	0.60; 0.40	0.42; 0.58
Only Willard	0.73; 0.27	0.80; 0.20	0.27; 0.73	0.11; 0.89	0.08; 0.91
No Willard	0.90; 0.10	0.90; 0.10	0.90; 0.10	0.75; 0.25	0.42; 0.58

Now, the only change we have made in Story 4 over Story 3 is to decrease the prior probability of hypothesis H from 0.5 to 0.25. As you see, the posterior on H is noticeably reduced when consider all the evidence, only Willard's evidence, or without Willard's evidence. Finally in Story 5 we have just made Willard much more probably untruthful. This causes the posterior on H to be noticeably smaller in all three of the evidence cases: all evidence, only Willard, or no Willard.

A good question is: Why did we choose to tell these five particular stories in preference to others we might have told? The answer is that changes in only credibility-related values or prior probabilities allow us quite easily to explain what happens to the endings of stories. We could have told stories involving changes in the other likelihood ingredients but we would have had a much more difficult time accounting for the endings of stories. As we noted, the mathematics in Bayesian Networks systems are always buried below the surface and never revealed to the user. We could have told stories that were much more interesting, and in many cases counterintuitive, by incorporating the conditional dependencies we mentioned above. The five stories we have told simply illustrate one of the virtues of Bayesian analyses. We can show how different stories might end if we use different conventional probability values in telling these stories.

Utility and Feasibility of Bayesian Network Analyses

As we noted, Bayesian Networks analysis are quite well known among many analysts and there are persons who stoutly advocate this form of analysis for inferences in intelligence analysis and in other contexts. We have also mentioned that Bayesian Networks analyses are unexcelled in their ability to capture and exploit an array of evidential and inferential complexities or subtleties, provided that we recognize them and adjust our networks to allow us to capture them.

The feasibility of Bayesian Networks analyses raises many questions we must consider. The most obvious matter concerns the time and effort it can take to perform such an analysis when we drill down to levels at which we try to capture as many sources of doubt or uncertainty we believe to be important. The example we have used in this account has been part of a very complex inference involving the major question or hypothesis concerning a dirty bomb being set off in the Washington DC area, discussed in section 1.5, to be finalized in the case study from Section 6.5. As you know by now, such a network is very large and complex. Our Bayesian Networks example in 0 just described involves only one very small sector or fragment of this very large and complex network. As you see in Table11, even this small sector requires 24 probability assessments. If we linked together a large number of such sectors or fragments, some of which would be much more complex than the Figure 45 network, the number of assessment would almost certainly be many times more than any analyst, or group of analysts, might have the time or inclination to make.

But Bayesian Networks advocates might argue that there are many ways in which we can drill down to shallower levels of detail in order to simplify our structural analysis and reduce the number of required probability assessments. But doing so would require us to ignore or suppress valuable sources of uncertainty of great importance that may influence our final inferences about the major hypotheses in the complex inference just described.

5.5 Belief Functions

Beliefs about Uncertainty

Both the enumerative and the subjective Bayesian interpretations of probability conform to Kolmogorov’s three axioms. We asserted that these axioms rest on the investigation of replicable or repeatable processes such as statistical analysis of the results obtained in a sample of observations. But there are many reasons for wondering whether these three axioms remain self-evident concerning subjective probability judgments we all make from time to time involving unique events for which no enumerative process can be involved. In a very influential work, the probabilist Professor Glenn Shafer pointed to an array of difficulties associated with Axiom 3 concerning the additivity of enumerative probabilities for mutually exclusive events (Shafer, 1976). In particular, Shafer asserts that this axiom places various constraints on our judgments or beliefs about uncertainty that we may not be willing to accept. Here it is only necessary to mention two of the difficulties Shafer mentions:

- Indecisions we routinely face concerning ambiguities in our evidence.
- Instances in which we encounter what historically has been called “pure” evidence.

In so many instances we may not be sure what evidence is telling us, and so we wish to be able to *withhold* a portion of our beliefs and not commit it to any particular hypothesis or possible conclusion. A very important element in what Shafer terms *belief functions* is that the *weight of evidence means the degree of support* evidence provides to hypotheses we are considering. Shafer allows that we can grade degree of support **s** on a 0-1 scale similar to the scale for Kolmogorov probabilities; but we can do things with support assignments **s** that the Kolmogorov additivity Axiom 3 does not allow.

Table 13. Sample support assignment.

	{H}	{¬H}	{H or ¬H}
s	0.5	0.3	0.2

To illustrate, consider a situation involving a singular evidence item where we have nothing to count. Suppose you are an analyst whose interest concerns whether or not the Greens are supplying parts necessary for the construction of shaped explosive devices to a certain insurgent militia group in the neighboring country Orange. Thus you are entertaining the following binary hypotheses:

H: The Greens are supplying parts necessary for the construction of shaped explosive devices.

¬H: The Greens are not supplying parts necessary for the construction of shaped explosive devices.

At some stage we are required to state our beliefs about the extent to which the evidence supports H or ¬H. Our assessment: is shown in Table 13.

What does this support assignment mean? We are saying that we believe the evidence supports H exactly to degree **s** = 0.5, and that this evidence also supports ¬H exactly to degree **s** = 0.3. But there is something about this evidence that makes us unsure about whether it supports H or ¬H. So, we have left the balance of our **s** assignment, **s** = 0.2, *uncommitted among* H or ¬H. In other words, we have withheld a portion of our beliefs because we are not sure what some element of our evidence is telling us.

Table 14. Support assignment if required to obey Kolmogorov Axiom 3.

	{H}	{¬H}
s	a	1-a

If we were required to obey Kolmogorov Axiom 3, we would not be allowed to be indecisive in any way in stating our beliefs. Table 14 is what our support assignment would have to look like. In this case, we would be required to say that the evidence supports H to degree **s** = a, and supports ¬H to degree **s** = (1 – a) in agreement with Axiom 3 since H and ¬H are mutually exclusive and exhaustive. In short, Kolmogorov Axiom 3 does not permit us any indecision in stating our beliefs; we must commit all of it to either H and to ¬H. This, we believe, would not be a faithful or accurate account of our beliefs.

Mixed and Pure Evidence

But Shafer's belief function approach allows us to cope with another difficulty associated with Kolmogorov's axioms. For centuries it has been recognized that a distinction is necessary between what has been termed *mixed evidence* and *pure evidence*. *Mixed evidence has some degree of probability under every hypothesis we are considering.* But *pure evidence may support one hypothesis but say nothing at all about other hypotheses*. In other words, we may encounter evidence that we believe offers zero support for some hypothesis. Here is another example involving our Green-Orange situation. Suppose we encounters an item of evidence we believes supports H to a degree, but we believes offers no support at all for $\neg H$. Our support assignment **s** for this evidence is shown in Table 15.

Table 15. Sample support assignment for pure evidence.

	{H}	{¬H}	{H or ¬H}
s	0.5	0	0.5

In this situation we are saying that the evidence supports H to degree **s** = 0.5, but offers no support at all to ¬H. The rest of our support we leave uncommitted between H and ¬H. But now we have to examine what **s** = 0 for ¬H means; does it mean that ¬H could not be supported by further evidence? The answer is *no*, and the reason why it is no allows us to compare what ordinary probabilities mean in comparison with what support **s** means. This comparison is shown in Figure 46.

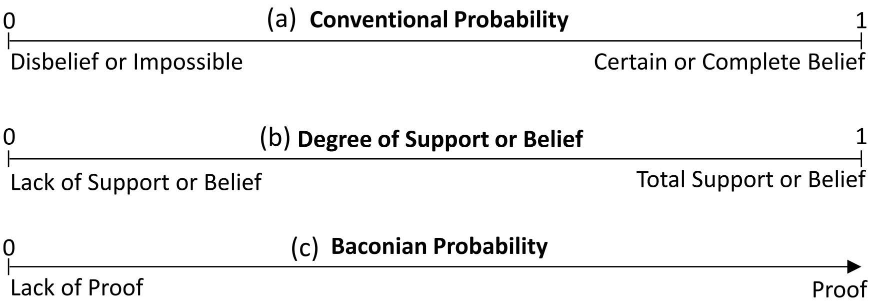
0
(a) Conventional Probability
1
Disbelief or Impossible
Certain or Complete Belief
0
(b) Degree of Support or Belief
1
Lack of Support or Belief
Total Support or Belief
0
(c) Baconian Probability
Lack of Proof
Proof

Figure 46. Different probability scales.

The (a) scale in Figure 46, for conventional or Kolmogorov probabilities, has a lower boundary with a meaning quite different from the meaning of this lower boundary on Shafer's support scale shown in (b). The value 0 in conventional probability refers to an event judged to be impossible and one you completely disbelieve. We will refer to this scale again in Section 5.6 when discussing Baconian probability. But all 0 means on Shafer's **s** scale is *lack of belief*, not disbelief. This is very important, since we can go from lack of belief to some belief as we gather more evidence. But we cannot go from disbelief to some belief. On a conventional probability scale, a hypothesis once assigned the probability value 0 can never be resuscitated by further evidence, regardless of how strong it may be. But some hypothesis, assigned the value **s** = 0, can be revised upward since we can go from lack of belief to some belief in this hypothesis when and if we have some further evidence to support it. Thus, **s** allows us to account for pure evidence in ways that ordinary probabilities cannot do.

Exampls of using the Belief Functions approach are provided in (Tecuci et al., 2016, pp.191-196). More details about this approach are provided in (Schum, 1994/2001a, pp. 222-243).

5.6 Baconian Probability

Variative and Eliminative Inferences

Here is a view of probabilistic reasoning that puts particular emphasis on a very important matter not specifically addressed in any other view of probability. In this view the weight of evidence depends on how much relevant and believable evidence we have and upon how complete is our coverage of existing evidence on matters we ourselves have recognized as being relevant in the analysis at hand. This Baconian view of probability and the weight of evidence is perhaps the least well known of any current views of probability, but it deserves much wider recognition. This Baconian view rests on the work of Professor L. Jonathan Cohen from Queens College, Oxford (Cohen, 1977, 1989).

The label "Baconian" on this system of probability acknowledges the work of Sir Francis Bacon (1561 – 1626) who revolutionized the process of inference in science. Bacon argued that attempting to prove some hypothesis by gathering instances favorable to it is mistaken, since all it would take to refute the generality of this hypothesis was one unfavorable instance of it. What Bacon argued was that we ought to design research with the objective of eliminating hypotheses. The hypothesis that best resists our eliminative efforts is the one in which we should have the greatest confidence. As this eliminative process proceeds, it is obvious that we should not keep performing the same test over and over again. What we need is an array of different tests of our hypotheses. The hypothesis that holds up under the most varied set of tests is the one having the greatest probability of being correct. So, Baconian inferences are eliminative and variative in nature.

Several attempts were made in the past, without success, to relate conventional probability to eliminative and variative inferences. Jonathan Cohen was the first person to generate a system of probabilities expressly congenial to this. His Baconian system of probabilities has properties unlike the two we have examined so far. Baconian probabilities have only ordinal properties and cannot be combined algebraically in any way. The Baconian probability scale is shown as (c) inFigure 46, to be compared with the conventional probability scale shown as (a) in Figure 46. On the conventional probability scale, 0 means disproof; but on the Baconian scale, 0 simply means lack of proof. A hypothesis now having zero Baconian probability can be revised upward in probability as soon as we have some evidence for it. As noted, we cannot revise upward in probability any hypothesis disproved, or having zero conventional probability.

Importance of Considering Evidential Completeness

Suppose that at, in a counterterrorism analysis, four hypotheses are being considered involving various actions a particular terrorist organization might take in the near future. These hypotheses describe the possible locations and times at which these actions might take place. We suppose that this terrorist group only has the means to take one of these actions. These competing hypotheses are: H_1, H_2, H_3, and H_4. Other hypotheses might be generated as time passes, but our "customer" in Homeland Security wishes to know which of these four possibilities now appears most likely since the decision maker in this agency wishes to be prepared to counter or prevent this terrorist action. Also imagine that a Bayes' Nets method is now being used by us for aggregating our probabilistic judgments in determining the relative probabiliies of hypotheses H_1, H_2, H_3, and H_4. On the available collection of evidence gathered so far, a posterior probability on H_3 of 0.998 is calculated using a Bayes' Nets system. Given the widespread confidence in this system and the very high probability it has assigned to H_3, the Homeland Security customer is given assurance by our intelligence group that H_3 is where and when the terrorist action will occur. Suppose H_3 says that the terrorist action will occur at a large shopping mall in Minneapolis, MN, on 12 November.

Time passes and it is now 11 November. Our customer has acted on our situation assessment and chooses a course of action that assumes the truth of H_3. As this course of action begins to unfold it becomes painfully apparent to our customer that H_3 is not true after all. In fact, from the most recent reports they are receiving, it appears that the situation should best have been described by our hypothesis $H_{1,}$ saying that the terrorist attack will occur at a large shopping mall in Duluth, MN. on 12 November. But our customer has already committed security forces to Minneapolis and must now try to make the best of what seems to a bad situation. If our customer cannot assemble security forces in Duluth in time and the terrorist attack occurs there, this tragic terrorist action will surely be tied to an "intelligence failure", to which operational miscarriages have been so readily tied in the past. A post-mortem analysis is now in progress regarding our apparent miscarriage in favoring H_3 so strongly.

At this post-mortem analysis a variety of evidential and inferential issues are examined. During this analysis attention turns to the Bayesian probabilistic scheme, which favored H_3 that now appears not to have been true. Someone asks: "How could we have gone wrong, given the strength with which our Bayes' Nets favored H_3? Everyone knows that the Bayesian approach is the optimal one we could employ" Pondering this question, another person finally asks: "Regardless of the strength and agreement of our current probabilistic assessments, was our conclusion based on enough evidence? Perhaps there were questions about the terrorist group, and other relevant matters bearing on the credibility and authenticity, of our evidence that we did not answer". This question generates considerable efforts to discover what matters were not taken into account in the analysis that led to H_3 being so strongly favored. The list grows quite long

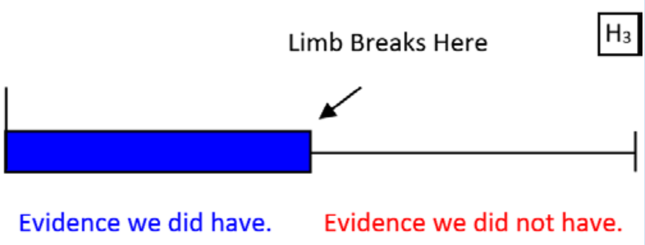

Figure 47. Our inferential limb.

and includes several questions concerning the terrorist group. The list also includes questions regarding the credibility and authenticity of our evidence that were not answered by any ancillary evidence.

As this list of unanswered questions grows, discussion turns to an examination of how the Bayesian probabilistic method employed to combine probabilistic assessments takes account of matters concerning the completeness or sufficiency of evidence. No one can think of a way in which a Bayesian posterior probability by itself reveals information about evidential completeness. All agree, however, that the Bayesian system offers a defensible way of combining our probabilistic beliefs regarding evidence we do have. What it lacks is a means for showing how final probabilistic conclusions should be additionally hedged by considering questions we have not answered and evidence we do not have. Someone in this post-mortem group says: “Maybe the Bayesian approach is not so optimal after all”.

One of the post-mortem participants then steps to the white board, draw Figure 48 and says: “In favoring H_3 as we did, based on the Bayesian probabilistic scheme we employed, we were actually out on a much longer and slender inferential limb than we realized. Part of this limb was strong, based as it was on the evidence we did take into account. But the rest of the limb was very weak since it includes matters we now recognize, but about which we had no evidence at the time we concluded that H_3 was so probable. As far as concerns our conclusion that H_3 was true, this is where our limb broke. So, it looks like high Bayesian posterior probabilities do not say all there is to say about the weight of evidence on conclusions we reach. Perhaps we must also consider how much evidence we had and how completely it covered matters we believe relevant in drawing a conclusion. I know of a probability system that does taken evidential completeness very seriously, it is called the Baconian probability system”.

The Bayesian system answers the question: How strong is the evidence we do have about this hypothesis? But the Baconian system answers the question: How much evidence do we have about this hypothesis, and how many questions about it remain unanswered?

Apart from the Baconian system, no other view of evidential reasoning focuses on evidential completeness and the importance of taking into account questions recognized as being relevant that remain unanswered by the evidence we do have. This is why Jonathan Cohen’s Baconian system is so important in intelligence analysis. What we do not take account of in intelligence analyses can hurt us very badly.

All inferences, made by intelligence analysts or anyone else, require generalizations that license inferential steps, and also require ancillary evidence to support the applicability of the generalization in the particular case in which it is being invoked. This is why we have said that generalizations and ancillary evidence form the “glue” that holds our arguments together. On some occasions, of course, this glue will fail to hold an argument together. As we all know, intelligence analysts frequently have to make inferences about matters for which they often have scant evidence, or no evidence at all. In other instances in which there may be available evidence, analysts may have no time to search for it or consider it carefully. The Baconian view being discussed offers much guidance in such matters.

But there is another important consequence of lack of time and evidence. This involves failure to decompose complex arguments into logically consistent or defensible stages. All arguments from evidence to hypotheses, or matters to be proved or disproved, involve chains of reasoning. Each link in such chains involves a source of doubt or uncertainty. Under the pressures of time or lack of evidence, an analyst may not even articulate in any way what these interposed sources of doubt may be. This results in a suppression of uncertainties and can lead to inferential miscarriages of many sorts. In our Cogent system we have used the term “drilling down” to indicate how much detail will be captured by an analyst in the construction of arguments from evidence. Cogent allows an analyst to drill down to various levels depending upon the time and evidential resources available to an intelligence analyst. In many cases this drilling down will be only very slight or nonexistent, in which cases the analysts must make generalizations that are not supported in any way. But this amounts to giving a generalization the benefit of the doubt, or to believing as-if some conclusion were true (absent any evidential support for it), or to taking something for granted without testing it in any way. All of these situations involve the suppression of uncertainties.

It happens that only the Baconian probability system provides any guidance about how to proceed when we must give benefit of doubt, believe as-if, or take things for granted. The major reason is that it acknowledges what almost every logician says about the necessity for asserting generalizations and supplying tests of them in evidential reasoning. Search the Bayesian or Belief Function literature and you will find almost no discussion of generalizations and ancillary tests of them. Suppose we are interested in inferring F from E. Bayes’ rule grinds to a halt when we have no basis for assessing the likelihoods $P(E|F)$ and $P(E|F^C)$. Bayesians counter by saying that we will always have some evidence on which to base these judgments. But they never say what this evidence is in particular cases and how credible it might or might not be. The Belief Function approach comes closer by saying that we can assess the evidential support for a body of evidence that may include both directly relevant and at least some ancillary evidence. Following is an account of the Baconian license for giving a generalization benefit of doubt, believing as-if it were true, or taking it for granted, provided that we are willing to mention all of the uncertainties we are suppressing when we do so. Stated another way, we must try to account for all of the questions we can think of that remain unanswered by the absence, or very scant amount, of evidence. This will be crucial in assisting an analyst to defend the generalization this analyst says is being made.

Here are the essentials of Cohen's Baconian approach to reasoning based on little or no ancillary evidence to either support or undermine a generalization. The first step, of course, is to make sure the generalization is not a non sequitur, i.e., it makes logical sense. In the simplest possible case, suppose we are interested in inferring proposition or event F from proposition or event E. The generalization G in doing so might read, "If E has occurred, then probably F has occurred." We recognize this if-then statement as an inductive generalization since it is hedged. Generalization G might also be stated in the future tense: "If E has occurred, then F will probably occur." Second, we consider various tests of this generalization using relevant ancillary evidence. Third, we consider how many evidential tests of this generalization there might be; suppose we identify N such tests. The best case would be when we perform all N of these tests and they all produce results favorable to generalization G. But we must not overlook generalization G itself; we do so by assigning it the value 1; so we have N+1 things to consider. Now we are in a position to show what happens in any possible case.

First, suppose we perform none of these N evidential tests. We could still proceed by giving generalization G the benefit of the doubt and detach a belief that F occurred (or will occur) just by invoking this generalization G regarding the linkage between events E and F. To do this we assign G the value 1 so that we are considering N + 1 things: the N evidential tests and our generalization G. So, when no evidential tests are performed, we are saying: "Let's believe as-if F occurred based on E and generalization G." This would amount to saying that the Baconian probability of event F is B(F) = 1/(N+1). This expression is never a ratio, all it says is that we considered just one thing in our inference about F from E, namely just the generalization G. We could also say, "Let's take event F for granted and believe that F occurred (or will occur) because E occurred as our generalization G asserts." However, note that in doing so, we have left all N ancillary evidential questions unanswered. This we represent by saying that our inference of F from E has involved only one of the N+1 considerations and so we have (N + 1 – 1) = N, the number of questions we have left unanswered. As far as evidential completeness is concerned, this is when the evidence we have is totally incomplete. But the Baconian system allows us to proceed anyway based on giving a generalization benefit of doubt. But our confidence in this result should be very low.

Now suppose we have performed some number k of the N possible ancillary evidential tests of generalization G, as asserted above, and they were all passed. The Baconian probability of F in this situation is given by B(F) = (k + 1)/(N +1). The difference between the denominator and numerator in such an expression will always equal the number of unanswered questions as far as the testing of G is concerned. In this case we have (N + 1) – (k + 1) = N – k questions that were unanswered in a test of generalization G. How high is our confidence that F is true depends on how high k+1 is as compared to N+1.

But now suppose that not all answers to these k questions are favorable to generalization G. Under what conditions are we entitled to detach a belief that event F occurred, based on evidence E, generalization G, and the k tests of G? The answer requires a subjective judgment by the analyst about whether the tests, on balance, favor or disfavor G. When the number of the k tests disfavoring G exceeds the number of tests favoring G, we might suppose that we would always detach a belief that event F did not occur, since G has failed more tests than it survived. But this will not always be such an easy judgment if the number of tests G passed were judged to be more important than the tests it failed to pass. In any case, there are N – k tests that remain unanswered. Suppose that k is quite large but the difference between the number of tests favorable to G is only slightly larger than the number of tests unfavorable to G. In such cases the analyst might still give event F the benefit of the doubt, or believe, at least tentatively, as-if F occurred pending the possible acquisition of further favorable tests of G. And in this case, the confidence of the analyst in this conclusion should also be very low.

Whatever the basis for an assumption or a benefit of doubt judgment there is, one of the most important things about the Baconian approach is that the analyst must be prepared to give an account of the questions that remain unanswered in evidential tests of possible conclusions. This will be especially important when analysts give generalizations benefit of doubt, draw as-if conclusions, or take certain events for granted. These are situations in which analysts are most vulnerable and in which Baconian ideas are most helpful.

Baconian Probability of Boolean Expressions

Some of the most important properties of Baconian probabilities concern their application to Boolean combinations of propositions such as hypotheses. Because the probabilities in the Baconian system have only ordinal properties, we can say only that hypothesis H_1 is more likely than H_2, but we cannot say how much more likely H_1 is than H_2. Also, in the Baconian system, it is never necessary to assess subjective probabilities. In our saying that H_1 is more probable than H_2, all we are saying is that there is more favorably relevant evidence on H_1 than there is on H_2. What counts most in the Baconian system is the completeness of our evidence and the extent to which we have questions that remain unanswered by the evidence we have. Here are the three most important Baconian properties of interest to us concerning intersections, unions, and negation.

Baconian Intersections: Suppose we have some events of interest like events F, G, and H. Suppose we have some favorably relevant evidence about each one of these events and have also considered how complete the evidence is for these events. So we determine that the Baconian probabilities (B) for these three events are $B(F) \geq B(G) \geq B(H)$. Here's what these probabilities say: we have more favorably relevant and complete evidence for event F than we do for event G, and more favorably relevant and complete evidence for event G than we have for event H. So, asked what the Baconian probability is for their intersection (F and G and H), we must say that B(F and G and H) = B(H). What this says is that the Baconian probability of the intersection of these three events is equal to the Baconian probability of the event with the least favorably relevant and complete evidence. This is an example of the MIN rule for Baconian intersections. We might compare this with the conventional probability of the intersection of these three events. Suppose that events F, G, and H are independent events where P(F) = 0.8, P(G) = 0.6, and P(H) = 0.4. In this case P(F and G and H) = 0.8×0.6×0.4 = 0.192 < P(H) = 0.4. In the Baconian system the probability of a conjunction of events or propositions can never be less than that of the event having the smallest Baconian probability.

Baconian Unions: Now consider the same case involving events F, G, and H. Again, suppose that B(F) ≥ B(G) ≥ B(H). Now what we wish to determine is the Baconian probability B(F or G or H). In this case, B(F or G or H) ≥ B(F), where B(F) is the largest of the Baconian probability for the events we are considering. This is the MAX rule for Baconian probability, and what it says is that the probability of a disjunction of events is at least as large as the largest Baconian probability of any of the individual events.

Baconian Negation. Baconian negation is not complementary. The Baconian rule is quite complex; here's what it says: If we have A and A^C, if B(A) > 0, then B(A^C) = 0. What this means essentially is that we cannot commit beliefs simultaneously to two events that cannot both occur.

What is quite interesting is that the Baconian treatment of conjunctions and disjunctions is the same as in Zadeh's Fuzzy probability system (Zadeh, 1963), namely they both make use of min-max rules for these connectives.

An example of the use of Baconian Probabilities is provided in (Tecuci et al., 2016, p.202). More details about this approach are provided in (Schum, 2001a, pp. 222-243). Again, the Baconian system is the only system we know about that is concerned about the completeness of evidence and the importance of considering how many questions remain unanswered by the evidence we have.

5.7 Fuzzy Probability

Fuzzy Force of Evidence

One can also express the uncertainty about a conclusion reached by using words, such as "likely", "almost certain", or "much less certain", rather than numbers, as illustrated by the following fragment from the letter sent by Albert Einstein to the United States President Franklin D. Roosevelt, on the possibility of constructing nuclear bombs (Einstein, 1939):

> *… In the course of the last four months it has been made probable — through the work of*

> *Joliot in France as well as Fermi and Szilárd in America — that it may become possible to set up a nuclear chain reaction in a large mass of uranium, by which vast amounts of power and large quantities of new radium-like elements would be generated. Now it appears almost certain that this could be achieved in the immediate future.*
>
> *This new phenomenon would also lead to the construction of bombs, and it is conceivable — though much less certain — that extremely powerful bombs of a new type may thus be constructed. …*

Years ago, Sherman Kent tried his best to relate intelligence analysts' verbal expressions of uncertainty such as these to specific ranges of numbers on a conventional probability scale (Kent, 1994). One instance we have all heard about involving words was use of the term "slam dunk" to indicate virtual certainty. This illustrates how no less care must be taken in verbal assessments of uncertainty than in the numerical methods just discussed.

Verbal expressions of uncertainty are common in many areas. In the field of law, for example, forensic standards of proof are always employed using words instead of numbers. We all know about standards such as: "beyond reasonable doubt" (in criminal cases); "preponderance of evidence" (in civil cases); "clear and convincing evidence" (in many senate and congressional hearings); and "probable cause" (employed by magistrates to determine whether a person should be held in custody pending further hearings).

All the verbal examples just cited have a current name; they can be called fuzzy probabilities. Words are less precise than numbers. There is now extensive study of fuzzy inference involving what has been termed approximate reasoning, which involves verbal statements about things that are imprecisely stated. Here is an example of approximate reasoning: "Since John believes he is overworked and underpaid, then he is probably not very satisfied with his job." All the underlined ingredients of this statement are imprecise or fuzzy. We are indebted to Professor Lofti Zadeh (University of California, Berkeley), and his many colleagues, for developing logics for dealing with fuzzy statements, including fuzzy probabilities (Zadeh, 1983; Negoita and Ralescu, 1975). It is, of course, entirely reasonable to grade the force of evidence in fuzzy terms, such as: "strong force", "weak force", "very strong force", and so on.

We have mentioned that Sherman Kent was concerned with fuzzy probabilities long before Lofti Zadeh began to study them carefully (Kent, 1994). But an American jurist named John H. Wigmore was well ahead of Sherman Kent as far as using words to grade uncertainty in evidential reasoning (Wigmore, 1913; 1937). Wigmore understood that the linkages between propositions in chains of reasoning reveal sources of doubt or uncertainty. What Wigmore was concerned about was the inferential force of one proposition on another in these chains of reasoning. But Wigmore was no probabilist and did not consider any numerical methods for grading evidential force. Instead, he used words such as "strong force", "weak force", and "provisional force" to indicate the strength of linkages in chains of reasoning. The point here is that verbal assessments of uncertainty in intelligence analysis have a long and very respectable lineage.

Zadeh went a bit farther than Kent did in his methods for relating verbal assessments of uncertainty with numerical equivalents. Zadeh employed what he termed a possibility function, μ, to indicate ranges of numerical probabilities a person might associate with a verbal expression of uncertainty. Zadeh reasoned that a person might not be able to identify a single precise number he/she would always associate with a verbal statement or fuzzy probability. Here is an example of a possibility function for the fuzzy probability "very probable."

Asked to grade what numerical probabilities might be associated with an analyst's fuzzy probability "very probable", the analyst might respond as follows:

For me, "very probable" means a numerical probability of at least 0.75 and at most 0.95. If it were any value above 0.95, I might use a stronger term such as "very, very probable." I would further say that I would not use the term "very probable" if I thought the probability was less than 0.75. In such cases, I would weaken my verbal assessment. Finally, I think it is most possible (μ = 1.0) that my use of the verbal assessment "very probable" means something that has about 0.85 of occurring. If the analyst decides that "very probable" declines linearly on either side of μ = 1.0, we would have the possibility function shown inFigure 48.

As an example of using fuzzy probabilities, suppose we have three events or propositions A, B, and C. We consider the following Fuzzy probabilities (F) for these events, and we say the following:

- Event A is very likely.
- Event B is likely.
- Event C is very unlikely.

We express this by saying that F(A) > F(B) > F(C).

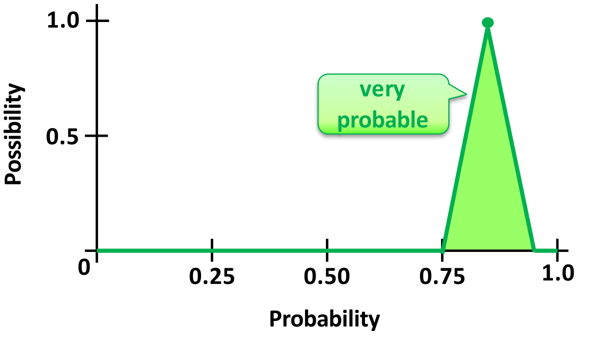

Figure 48. Possibilities and fuzzy probabilities.

Fuzzy Probability of Boolean Expressions

Fuzzy Intersections. In the situation just described, the Fuzzy conjunction of these three events F(A and B and C) = F(C), which is the minimum Fuzzy probability of the three events.

Fuzzy Unions. In the situation just described, the Fuzzy disjunction of these three events F(A or B or C) = F(A), which is the maximum Fuzzy probability of these three events.

So, both in the Baconian system and in the Fuzzy system we have MIN/MAX rules for combining probabilities for complex events. These systems give us formal license to use the MIN and MAX rules in Cogent. But there is one other Boolean connector we have not yet discussed: it is negation, where we have some event and its complement, like A and A^C. It is here that the Baconian and Fuzzy views depart.

Fuzzy Negation. Fuzzy negation is complementary: $F(A) = 1 – F(A^C)$.

On Verbal Assessments of Probabilities

Let us also consider the critics who sneer at verbal assessments of probabilities, saying that only numerical assessments, conforming to the Kolmogorov axioms, are the only ones acceptable. As a top-ranking analyst, you are asked by an equally high-ranking customer for the probability of a crucial hypothesis H_K. All the evidence in this case is for one-of-a-kind event and so your assessment is necessarily subjective. You tell the customer, “Sir, the probability of H_K, on our analysis is 78%”. The customer asks, “This is a very precise number, how did you arrive at it, given the subjective nature of your assessment”. You reply, “Yes sir, what I really should have said was that my probability is between 73% and 83%, and 78% seemed like a good figure to quote.” The customer then says, “But the limits to the probability interval you quoted are also precise, how did you arrive at them?” You might say, “Well my lower limit is really between 70% and 76% and my upper limit is between 80% and 86%.” Your customer says, “But these are also precise numbers”. There is, as you see, an infinite regress of similar questions regarding the basis for subjective numerical assessments.

There are many places to begin a defense of verbal or fuzzy probability statements. The most obvious one is law. All of the forensic standards of proof are given verbally: “beyond reasonable doubt”, “clear and convincing evidence”, “balance of probabilities”, “sufficient evidence”, and “probable cause.” Over the centuries attempts have been made to supply numerical probability values and ranges for each of these standards, but none of them have been successful. The reason, of course, is that every case is unique and rests upon many subjective and imprecise judgments. Wigmore (1913) understood completely that the catenated inferences in his Wigmorean networks were probabilistic in nature. Each of the arrows in the chain of reasoning describe the force of one hypothesis on the next one, e.g., $E \rightarrow F$. Wigmore graded the force of such linkages verbally using such terms as “strong force”, “weak force”, “provisional force”, etc. Toulmin (1963) also used fuzzy qualifiers in the probability statements of his system which grounds Rationale (van Gelder, 2007). There are many other examples of situations in which it is difficult or impossible for people to find numerical equivalents for verbal probabilities they assess. Intelligence analysis so often supplies very good examples in spite of what Sherman Kent said some years ago. Indeed, using words is quite often necessary in analyses based on masses of evidence that are so complex that they resist even the most devoted attention to the construction of inference networks. Couple this with the fact that different analysts might disagree substantially about what specific probability should be assigned to a conclusion. In addition an analyst might assign a different probability to the same conclusion, based on the same evidence, on different occasions. What this says is that there will be inter-analyst and intra-analyst variation in the assessment of probabilities. Words are less precise than numbers and so there will often be less disagreement about a verbal or a fuzzy probability.

We conclude this discussion by recalling what the well-known probabilist Professor Glenn Shafer said years ago (Shafer, 1988): “Probability is more about structuring arguments than it is about numbers. All probabilities rest upon arguments. If the arguments are faulty, the probabilities however determined, will make no sense.”

5.8 Complementarity of the Prbability View

As was discussed in more details in Section 3.5, there are five sources of uncertainty and corresponding challenges to drawing accurate conclusions from evidence: our evidence is always incomplete, usually inconclusive, frequently ambiguous, commonly dissonant, and with various degrees of credibility. As presented in Table 16 and discussed below, each of the discussed probability view can cope best with some of these challenges, but no single view copes best with all of them. Therefore the intelligence analysts should know about all of them.

The *Incompletness* entry in Table 16 lists a major strength that is exclusive to the Baconian system, its concern about how much favorable evidence was taken into account in an analysis, and how completely this evidence covered matters judged relevant to conclusions that could be reached. A major question this form of analysis allows us to address is the extent to which questions that have not been answered by existing evidence could have altered the conclusion being reached. It would be quite inappropriate to assume that answers to the remaining unanswered questions would, if they were obtained, all favor the conclusion that was being considered. This, of course, requires analysts to consider carefully matters relevant to any conclusion that are not addressed by available evidence. We acknowledge that completeness matters are difficult to manage in current intelligence in which analysts are asked to provide conclusions on very short order. The shorter the time available for the assessment of evidence the more unanswered questions there will be. We hope our customers requiring quick analyses appreciate this fact.

The *Inconclusiveness* entry in Table 16 notes that all four of the uncertainty methods have very good ways for dealing with the inconclusive nature of most evidence; but they do so in different ways. The Subjective Bayesian does so by assessing non-zero likelihoods for the evidence under every hypothesis being considered. Their relative sizes indicate the

Table 16. A summary of non-enumerative uncertainty methods and what they best capture.

Evidence Characteristic	Subjective Bayes	Belief Functions	Baconian	Fuzzy
Incompleteness			☑	
Inconclusiveness	☑	☑	☑	☑
Ambiguity		☑		☑
Dissonance	☑	☑	☑	☑
Credibility	☑		☑	

force the evidence is judged to have on each hypothesis. But the Belief Functions advocate assigns numbers indicating the support evidence provides for hypotheses or subsets of them. We should be quick to notice that Bayesian likelihoods do not grade evidential support, since in belief functions an analyst can say that an item of evidence provides no support at all to some hypothesis. But a Bayesian likelihood of zero under a particular hypothesis would mean that this hypothesis is impossible and should be eliminated. Offering no support in belief functions does not entail that this hypothesis is impossible, since some support for this hypothesis may be provided by further evidence. The Baconian acknowledges the inconclusive nature of evidence by assessing how completely, as well as how strongly, the evidence favors one hypothesis over others. In Fuzzy probabilities it would be quite appropriate to use words in judging how an item or body of evidence bears on several hypotheses. For example, an analyst might say, “This evidence is indeed consistent with H_1 and H_2, but I believe it *strongly favors* H_1 over H_2.”

The *Anbiguity* entry in Table 16 first acknowledges the Belief Functions and Fuzzy concerns about ambiguities and imprecision in evidence. In the Belief Functions approach, an analyst is entitled to withhold belief for some hypotheses in the face of ambiguous evidence. In such cases the analyst may not be able to decide upon the extent to which the evidence may support any hypothesis being considered, or even if the evidence supports any of them. Judgmental indecision is not allowed in the Bayesian system since it assumes the analyst can say precisely how strongly evidence judged relevant favors every hypothesis being considered. Judgmental indecision, as allowed in the Belief Functions system, seems a natural attribute of many evidential matters encountered by intelligence analysts. Ambiguities in evidence may be commonly encountered. The Fuzzy advocate will argue that ambiguities or imprecision in evidence hardly justifies precise numerical judgments. In the face of fuzzy evidence we can only make fuzzy judgments of uncertainty.

The *Disonance* entry in Table 16 shows that all four probability systems have very good mechanisms for coping with dissonant evidence in which there are patterns of contradictory and divergent evidence. Recall that dissonant evidence is directionally inconsistent; some of it will favor certain hypotheses and some of it will favor others. In resolving such inconsistencies, both the Bayesian and Belief Functions approaches will side with the evidence having the strongest credibility, though the mechanisms for doing so differ in the Bayes’ rule used for Bayesian probabilities and the Dempster’s rule used for Belief Functions. The Bayesian approach to resolving contradictions is especially interesting since it shows how “counting heads” is not the appropriate method for resolving contradictions. In times past, “majority rule” was the governing principle. Bayes’ rule shows that what matters is the aggregate credibility on either side of a contradiction. The Baconian approach also rests on the strength and aggregate credibility in matters of dissonance, but it also rests on how much evidence is available on either side and upon the questions that remain unanswered. In Fuzzy terms, evidential dissonance, and how it might be resolved, can be indicated in verbal assessments of uncertainty. In such instances an analyst might say, “We have dissonant evidence favoring both H_1 and H_2, but I believe the evidence favoring H_1 predominates because of its very strong credibility.”

Row five in Table 16 concerns the vital matter of assessing the credibility of intelligence evidence. From considerable experience, we find that the Bayesian and Baconian systems are especially important when they are combined. In many cases these two radically different schemes for assessing uncertainty are not at all antagonistic but are entirely complementary. Let us consider a body of evidence about a HUMINT asset or informant. Ideas from the Baconian system allow us to ask, *“How much evidence do we have about this asset, and how many questions about his asset remain unanswered?”* Ideas from the Bayesian system allow us to ask: *How strong is the evidence we do have about this asset* (Schum, 1991)?

Finally, notice from Table 16 that these probability views are complementary in their ability to drawing accurate conclusions from evidence. In Section 6.2 we present a mixed probability view hat can cope wih all the five characteristics of evidence.

The following sections describes a variety of ways intelligence analysts can build strong arguments and responsibly qualify them with accurate expressions of uncertainty. This discussion also deconstructs the close relationship between the strength of evidence and argument on the one hand, and uncertainty on the other.

5.9 Review Questions

114. As we noted, the subjective Bayesian view of probability lets us assess probabilities for singular, unique, or one-of-a-kind events, provided that our assessed probabilities obey the three Kolmogorov axioms we discussed above regarding enumerative probabilities. First, is there any way of showing that these axioms for enumerative probabilities also form the basis for ideal or optimal probability assessments in the non-enumerative case? Second, can this really be the rational basis for all probability assessments based on evidence?

115. Show how Bayes’ rule supplies no method for incorporating “pure evidence” as does the Belief Function system.

116. Provide an example showing how an analyst’s numerical assessment of a probability applied to a conclusion can invite criticism.

117. Think back to the very first time you were ever tutored about probability, what it means, and how it is determined. What were you told about these matters? Then, describe your present views about these probability matters.

118. Think back to the very first time you were ever tutored about probability, what it means,

and how it is determined. What were you told about these matters? Then, describe your present views about these probability matters.

6 HYPOTHESIS ANALYSIS

6.1 Uncertainty and Arguments

Questions about uncertainty are naturally linked to views about probability. Some analysts may have studied probability quite extensively; many others will have little or no formal tutoring on the subject. Still others may dislike the sometimes complicated formulas that the study of probability often requires. However, words and pictures - rather than mathematics - can express very useful ideas about probability and uncertainty. One does not need a background in mathematics or statistics to reason with uncertainty and draw conclusions from evidence.

The following sections describes a variety of ways intelligence analysts can build strong arguments and responsibly qualify them with accurate expressions of uncertainty. This discussion also deconstructs the close relationship between the strength of evidence and argument on the one hand, and uncertainty on the other.

6.1.1 Wigmore’s Legacy on Evidentiary Reasoning

We owe a great debt to the field of law for the very rich legacy of experience and scholarship it has provided concerning the credentials of evidence and our drawing conclusions from it. As far as our present purposes are concerned, no one has contributed more to this rich legacy than an American evidence scholar named John H. Wigmore [1863 - 1943]. Wigmore's views on evidence are still widely cited in current discussions of evidence in the field of law. Our major interest in Wigmore's work concerns a book he wrote on what he termed the science of judicial proof (Wigmore, 1937). Wigmore was the very first person to study the construction of chains of reasoning in arguments such as those we have been considering. He called such reasoning catenated inferences, from which we have drawn the metaphor of "links" in chains of reasoning. Today we call such inferences cascaded, hierarchical, or multi-stage. More importantly perhaps Wigmore was the first person known to us to study systematically the task of drawing conclusions from masses of different kinds of evidence. Today we refer to such a process as involving inference networks. In 1913 Wigmore wrote the very first work on this subject in which he described an analytic and synthetic method for constructing complex arguments from evidence in what we term today as inference networks (Wigmore, 1913).

We have drawn heavily on Wigmore's ideas concerning the process of constructing defensible arguments from masses of evidence. We know of no work in another field that has greater relevance to intelligence analysis than Wigmore's work on what is involved in constructing defensible arguments from masses of different kinds of evidence that come to us from a variety of different sources. In fact, Wigmore's work contains many insights not evident in more recent work on inference networks.

The intelligence analyst is quite accustomed to using expert judgment as a basis for analyses delivered to senior U.S. policymakers. The analyst might cull the judgments of several other analysts to integrate several lines of investigation and analysis into a complex assessment. In any case, suppose that the analyst has generated and considered several possible conclusions and has decided that the relevant and credible evidence most strongly favors a particular conclusion. Presumably, the analyst reached this conclusion by considering many sources of doubt or uncertainty about how strongly the evidence favored a particular conclusion. But now the question arises: How certain can the analyst be that the chosen conclusion reached is the correct one, and how to communicate uncertainty about this conclusion to supervisors and policymakers? Failing to accurately report uncertainty about a conclusion risks misleading customers which could have grave consequences for U.S. foreign policy.

6.1.2 Decomposed Versus Holistic Analyses

There are alternative ways of describing the complexity of the tasks we face in the analyses based on masses of different forms of evidence coming from different sources. One way, presently very common, is to say that we have very difficult tasks of “connecting the dots”, where we have masses of dots to be connected, and these dots can be connected in various complex ways. Another way is to describe our inferences about major hypotheses as being exercises in drawing final conclusions about matters of interest involving many sources of doubt that arise from the evidence we are considering and the arguments we have constructed based on this evidence. Here come two questions of vital importance in such complex evidential reasoning tasks:

- How many sources of doubt should we try to identify and combine?
- How many dots should we try to capture and connect?

These questions have a virtually unlimited number of possible answers depending on our time for analysis, objectives, and resources. These same important questions arise in many contexts besides intelligence analysis, such as in law, medicine, science, and many other areas of interest.

We start by discussing two extremes in answers to the above questions. The first involves the Wigmorean analysis we have performed in Section 2.5 for the Case of General Alpha. The Wigmorean approach encourages us to try to identify all the sources of doubt, or dots, we can imagine that could affect the conclusions we could reach, based on evidence we can defend as being relevant to these conclusions. This approach encourages us to decompose our evidential reasoning task to as detailed a level as we can. Stated in other words, the Wigmorean approach encourages us to drill down to as deep a level as we can.

At the other extreme, we could entertain an extreme level of what is called holistic analysis. Holistic analyses involve various levels of suppression of doubts we may either ignore, or don't even attempt to identify, in forming our final conclusions. We do everything in our heads, according to methods we can never articulate, in expressions of our uncertainty about a stated conclusion. Here is the most extreme level of holistic analysis possible in our Case of General Alpha. In this extreme form of holistic analysis, we don't drill down at all. Here is the ultimate hypothesis we are considering:

H: The Blues will succeed in their insurgency against General Alpha's government in country Orange.

An analyst, or group of analysts, considers this hypothesis and evidence bearing on it, and says: "I [or we] have considered the evidence for this hypothesis H and can conclude that it is very likely to be true. If you require numbers, I can say that it has a probability of 0.7 of being true". That's all; no further words are provided that indicate how the strength of this conclusion was determined. Asked by a "customer" to state the basis for this fuzzy or exact probability, the analyst might experience varying levels of difficulty in coming up with answers to this question. What the analyst is doing here is saying: "Take my word for it, H is very likely, or having 0.7 probability, of being true". This would hardly be acceptable to most decision-making consumers of intelligence analyses.

Here is a slightly less extreme form of holistic analysis, and one that drills down to only a very shallow level of detail. Suppose the analyst, or analysts, says: "I have considered various arguments resulting from a parsing of our major hypothesis H that are as follows:

H1: The Blues enjoy popular support among the citizens of country Orange.
H2: The Blues will have the military capability necessary for the insurgency to succeed.
H3: The Orange military is vulnerable to an insurgency.
H4: The Blue group leadership is adequate to make the insurgency successful.

The analyst continues: "Based on the evidence we have so far, I think P_1 is extremely probable (0.95); P_2 and P_3 are very probable (0.75); and P_4 is only fairly probable (0.60). On this basis, I conclude that hypothesis H is quite probable, say 0.70". In this case, the analyst has only given the barest description of the probabilistic ingredients of theirhis/her final hedge on hypothesis H. But the analyst has given no further information about the basis for these four probabilities, and no information about how they were combined. This level of holistic analysis would very likely be no more appealing to a decision-making consumer than the most extreme holistic analysis mentioned above.

Our next task is to illustrate how drilling down to various levels of detail influences the very structure of any intelligence analysis based on patterns of evidence. This process also allows us to account for the number of sources of doubt or uncertainty we recognize at each level of drilling down. Starting with our most extreme level of holistic analysis, we have no inference network at all, and only one recognized source of doubt, namely our ultimate hypothesis H. At the shallowest level of drilling down, or task decomposition we do have the simple inference network from Figure 32. At this level we have five sources of doubt consisting of H and each one of the four main lines of argument we have considered.

Any number of further levels of task decomposition or drilling-down are possible. For example, we could marshal our available evidence to see what lines of argument we could generate on each one of the four major lines of argument we have constructed so far on our hypothesis H**,** as discussed in Section 2.6. This would result in the identification of more sources of doubt. But our arguments begin to be the most complex inference networks, and the largest number of sources of doubt will be identified, as we attempt to link specific patterns or combinations of our evidence to each one of these major arguments. The deeper we drill down or decompose our analysis of complex evidential reasoning tasks we face, the more sources of doubt or uncertainty we will identify and the more difficult are the probability assessment and combination tasks we will face, regardless of the probability system we employ. A basic fact about this complex activity is that no probabilistic or other formal theory associated with the analysis of inference networks will tell us what our constructed networks should be, and how many sources of doubt we must identify, for any particular complex evidential reasoning task we face.

But there is another vitally important imaginative reasoning task. The construction of an argument based on evidence is always grounded first on imaginative reasoning. In the Case of General Alpha, as in all other actual intelligence analyses, there will never be a reference of any sort that an analyst can consult to see what arguments should be constructed from any pattern of evidence. The analyst must imagine what chains of reasoning would plausibly link evidence to some matter to be proved, such as the four sub-hypotheses in our case.

But there is another vital form of reasoning in the construction of arguments that we now address; it involves critical reasoning.

The final very important item we must consider involves the defensibility and persuasiveness of our arguments from evidence to matters to be proved. We have always taken to heart something said years ago by the justifiably noted probabilist Professor Glenn Shafer. Shafer said that *probability is more about arguments than it is about numbers. All probabilities rest upon arguments. If the arguments are faulty, the probabilities however determined, will make no sense.* What this means is that we must subject our arguments to very critical analyses in which we attempt to find any disconnects or non sequiturs in them. So, careful argument construction involves both imaginative and critical reasoning on the part of analysts. Regarding the persuasiveness of our analyses, and the arguments upon which they rest, it seems easy to say that *if we have reached imaginative and productive conclusions, resting upon defensible arguments, and if our conclusions are hedged by some appropriate probabilistic method, then our analysis will be persuasive*. The trouble of course depends upon the persons we are trying to persuade. Known since the time of the ancient Greeks is that some persons fail to be persuaded by valid arguments and often are persuaded by invalid arguments. As we all learn, the "customers" of our intelligence services may form their own beliefs about matters of vital interest to our nation's security and may disregard even the most imaginative, defensible, and productive intelligence analyses. This is an argument that the "customers" of our intelligence services should also receive the same kind of training that you are now receiving in this book.

6.1.3 Divide and Conquer

Reading about the problems with holistic assessments in the example discussed in the previous section, you might believe that they have very little relevance to intelligence analysis. You might say that complex intelligence analysis requires teams of analysts working together to address some problem; it rarely involves just a single analyst thinking alone in some secluded office. You might easily give an example involving counterterrorism. In this situation we would have weapons experts, existing terrorist group experts, geopolitical experts, cultural affairs experts, and a whole host of other experts at work trying to predict the next occurrence of a terrorist incident. But this is actually the very first stage of a problem we now address concerning the structure of an intelligence problem and its specific ingredients. All we have done so far is to note the obvious multidimensional nature of so many intelligence problems. We have just begun the process called divide and conquer.

Within any of the areas of expertise just noted, there are questions raised that require answers, there are masses of evidence to be considered, and there are many dots to be connected. And there are further dots to be connected when conclusions from the experts in each area are brought together in order to draw any conclusion. Now consider any of the experts such as in the counterterrorism example just mentioned. Each one of these experts faces the problems discussed in Section 0 concerning the task of connecting the dots. We have no way of knowing about the extent to which intelligence experts employ holistic approaches to some degree such as that described in the example above. But we strongly suspect that most analysts take great care to be able to defend their conclusions. One well-established key to being more easily able to defend the conclusions in some inferential or decision problem you face is to decompose this problem into smaller elements; this is the basic strategy of divide and conquer.

It has been recognized for many years that decomposing a problem into smaller elements seems an obvious way of simplifying difficult problems. The study of task decomposition has been an object of study by psychologists, computer scientists and others for a long time. As obvious as the benefits of divide and conquer approaches are, there are some difficulties that are not always recognized. First, it is not always recognized how many pieces or smaller elements there are in some problem being decomposed. This is of course related to the issue: How deep do we wish to decompose a problem? There could be any number of levels or gradations of any problem decomposition. If we do not make our decomposition fine enough, or detailed enough, we may fail to make important distinctions that should be made and that will affect the accuracy of our conclusions. On the other hand, when we decompose a problem into too many smaller elements, we may easily be overcome by their number and never reach any conclusion at all. We have all heard the expression "paralysis by analysis." Too detailed an analysis is one reason for this paralysis. So, an abiding issue is: How detailed should task decomposition be and how many of the complexities of an evidential reasoning problem should we try to capture and analyze? Another way of stating this problem is to ask: How many sources of doubt should we try to expose in our intelligence analysis?

But there is a second problem often associated with task decomposition. It is often alleged that the smaller elements into which a complex problem is decomposed are easier to deal with. Unfortunately, this is not always the case. In making holistic, or even partially decomposed, judgments we may so often fail to recognize how many and how difficult these judgments actually are.

Analysts face different requirements in their efforts to serve their policy and decision-making "customers." In some cases, they are required to answer questions that are of immediate interest and that do not allow time for extensive research and deliberation on available evidence. In other cases, teams of analysts participate in more lengthy finished intelligence that combines evidence from every available source. Sometimes finished intelligence can refer to long-term assessments on matters of current and abiding interest. Each such requirement will have a direct influence on the depth of the problem decompositions performed by the analysts. That is, the analysts may decompose problems at various levels of detail (depending on the available time and evidence), evaluating the problems at those levels, as will be discussed in the following section.

6.2 Augmented Wigmorean Argumentations

6.2.1 Baconian Probabilities with Fuzzy Qualifiers

What is especially of interest to us is the fact that both the Baconian system and the Fuzzy system use the min/max rules for combining probabilistic assessments. This gives us the license to define a combined Baconian and Fuzzy probability system. We will use the simple problem from Table 17 to introduce the elements of this system.

Table 17. The chemical weapons problem.

Situation: Velovia, a technologically advanced democratic state, is home to Hakka, an apocalyptic sect suspected of plotting mass casualty terrorist attacks.

Question: Does Hakka have chemical weapons?

Available Information:

A source, who has reported accurately in the past, indicated that he knows a member of Hakka, Sulfurus, who has a bachelor's degree in chemistry.

According to official documents, Hakka has created two chemical companies. These companies have purchased technical equipment and substantial amounts of chemicals.

Regulatory and policing constraints are strong in Velovia. However, chemicals of the purity required for chemical weapons (e.g., sarin) can be procured at low visibility for plausibly legitimate business purposes.

A source from inside Hakka, who has reported accurately in the past, reported that Hakka has the ability to raise funds from its members but the source was not specific on how much. The source also reported that Hakka does not have chemical weapons.

Hakka also runs a hospital that conducts some biological research for which it has purchased biological materials.

Although synthetic biology and other developments enhance opportunities for biotoxin synthesis, creating or procuring seed stock for botulinum-based weapons is more difficult than procuring chemicals.

6.2.2 Probability Scale

This symbolic probability system may be used with different probability scales, but the Cogent system for this book uses the one presented in Table 18. Stating that hypothesis H is "likely" means that, based on the available evidence, H's probability of being true is between 55% and 70%. Stating that hypothesis H is "more than likely" means that, based on the available evidence, H's probability is between 70% and 80%. Stating that hypothesis H is "lacking support" means that the available evidence does not support its truthfulness with a probability greater than 50%. The probability of H could be anywhere between 0% and 50%.

Table 18. Probability scale.

100%	certain (C)
95-99%	almost certain (AC)
80-95%	very likely (VL)
70-80%	more than likely (ML)
55-70%	likely (L)
50-55%	barely likely (BL)
0-50%	lacking support (LS)

6.2.3 Evidence, Credibility, and Fact

It is important to distinguish between evidence about a fact and the fact itself (see Figure 49). Consider, for example, the following item of evidence: "A source, who has reported accurately in the past, indicated that Hakka has a member with a bachelor's degree in chemistry." Can we conclude from this that Hakka has a member with a bachelor's degree in chemistry? No. At issue here is the credibility of the source who may or may not be telling the truth.

The credibility of an item of evidence is the extent to which the evidence may be believed. We can assess the credibility of this item of evidence by answering the following question: What is the probability that the evidence is true? This assessment can be influenced by many factors, including doubts about the source's veracity or objectivity. The source of this evidence has reported accurately in the past. We can therefore assume that this current report is very likely to be true.

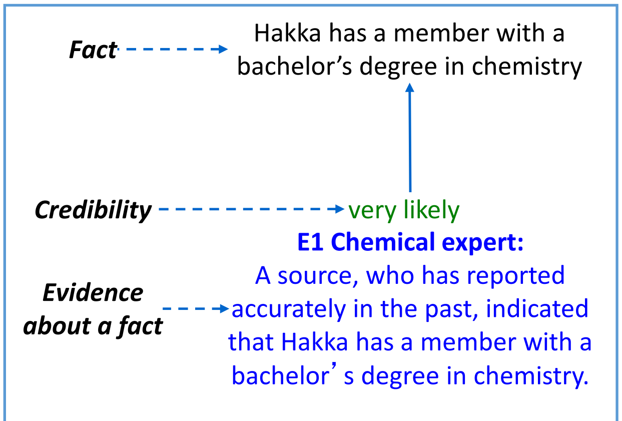
Fact
Hakka has a member with a bachelor's degree in chemistry
Credibility
very likely
E1 Chemical expert:
A source, who has reported accurately in the past, indicated that Hakka has a member with a bachelor' s degree in chemistry.
Evidence about a fact

Figure 49. Evidence about a fact and the fact itself.

6.2.4 Relevance of Evidence to a Hypothesis

Another credential or property of evidence is its relevance. The relevance of an item of evidence indicates how strongly does this item support a specific hypothesis in the argument. We can assess the relevance by answering the question:

> *Assuming that the evidence is true, what is the probability that the hypothesis is true?*

When we have evidence about a fact and the hypothesis is the fact itself, the relevance of evidence is certain, as shown in Figure 50. Indeed, if we assume that the evidence is true, then the hypothesis is true. But let us consider the hypothesis "Hakka has expertise to develop chemical weapons." If Hakka has a member with a bachelor's degree in chemistry, what is the probability that it has expertise to develop chemical weapons?

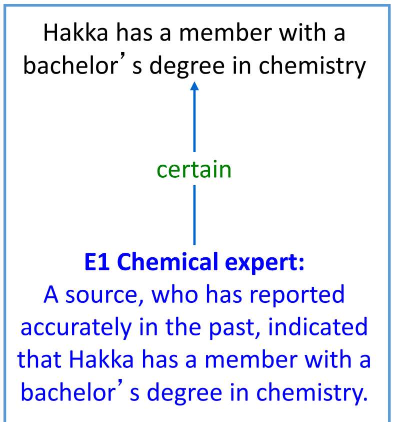
Hakka has a member with a bachelor' s degree in chemistry
certain
E1 Chemical expert:
A source, who has reported accurately in the past, indicated that Hakka has a member with a bachelor' s degree in chemistry.

Figure 50. Simple relevance.

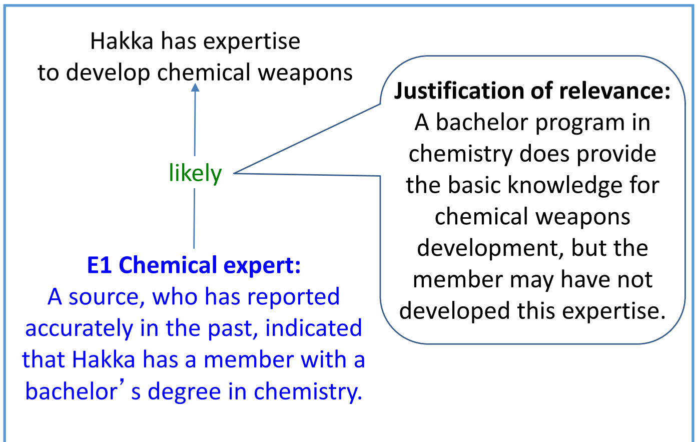
Hakka has expertise to develop chemical weapons
likely
E1 Chemical expert:
A source, who has reported accurately in the past, indicated that Hakka has a member with a bachelor' s degree in chemistry.
Justification of relevance:
A bachelor program in chemistry does provide the basic knowledge for chemical weapons development, but the member may have not developed this expertise.

Figure 51. Justification of a relevance assessment.

A bachelor program in chemistry does provide the basic knowledge for chemical weapons development, but this does not necessarily prove that Hakka has the expertise. Indeed, the Hakka member may have not developed this expertise. Thus, this item of evidence is not conclusive and we assess its relevance only as "likely." Explicit explanations (justifications) for assessments of relevance that are less than certain, as in this example, will clarify the reasoning and can be used in the completed argumentation to support the conclusion (see Figure 51).

When developing an argumentation, it is a good practice to consider, for each item of evidence, which is the corresponding fact, and then reason from that fact to the upper-level hypotheses, as illustrated in Figure 52.

The justification of the relevance assessment in *Figure 51* points to the missing information that, if present, would have allowed us to assess the relevance as certain. An implied assumption has to be made for the evidence E1 to support the hypothesis. The probability of this implied assumption being true, in effect, becomes the assessed relevance of the evidence E1, as shown in*Figure 53*.

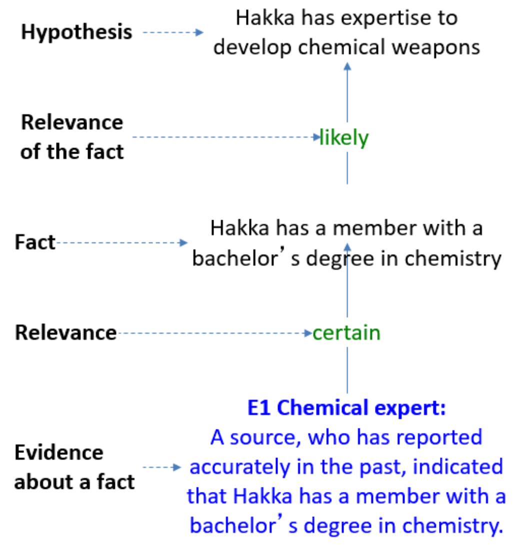
Hypothesis
Hakka has expertise to develop chemical weapons
Relevance of the fact
likely
Fact
Hakka has a member with a bachelor's degree in chemistry
Relevance
certain
Evidence about a fact
E1 Chemical expert:
A source, who has reported accurately in the past, indicated that Hakka has a member with a bachelor's degree in chemistry.

Figure 52. Recommended evidence-based reasoning chain.

6.2.5 Inferential Force of Evidence

The third credential of evidence is its inferential force, defined and illustrated in Figure 54. We have assessed the relevance of evidence item E1 as certain because it is an inference from evidence about a fact to the fact itself. We have also assessed the credibility of E1 as very likely because the source has reported accurately in the past. *Inferential force* answers the question:

> *What is the probability of the hypothesis above being true based only on this item of evidence below?*

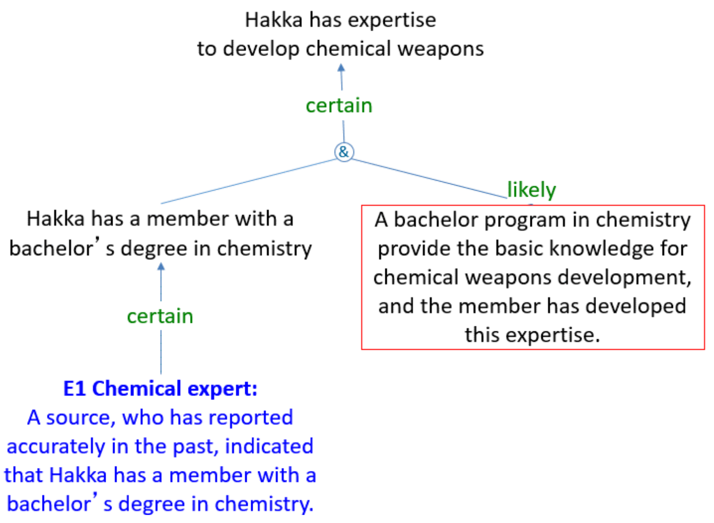
Hakka has expertise to develop chemical weapons
certain
&
likely
Hakka has a member with a bachelor's degree in chemistry
A bachelor program in chemistry provide the basic knowledge for chemical weapons development, and the member has developed this expertise.
certain
E1 Chemical expert:
A source, who has reported accurately in the past, indicated that Hakka has a member with a bachelor's degree in chemistry.

Figure 53. Relevance as pobability of an implied assumption.

In our example, the relevance of E1 is certain, but its credibility is only very likely. Therefore, the probability that Hakka does in fact have a member with a bachelor's degree in chemistry is only very likely. In general, *the inferential force of an evidence item is determined as the smaller between its credibility and its relevance.* Indeed, an evidence item that is not credible would not convince us that the hypothesis is true, no matter how relevant the provided information is.

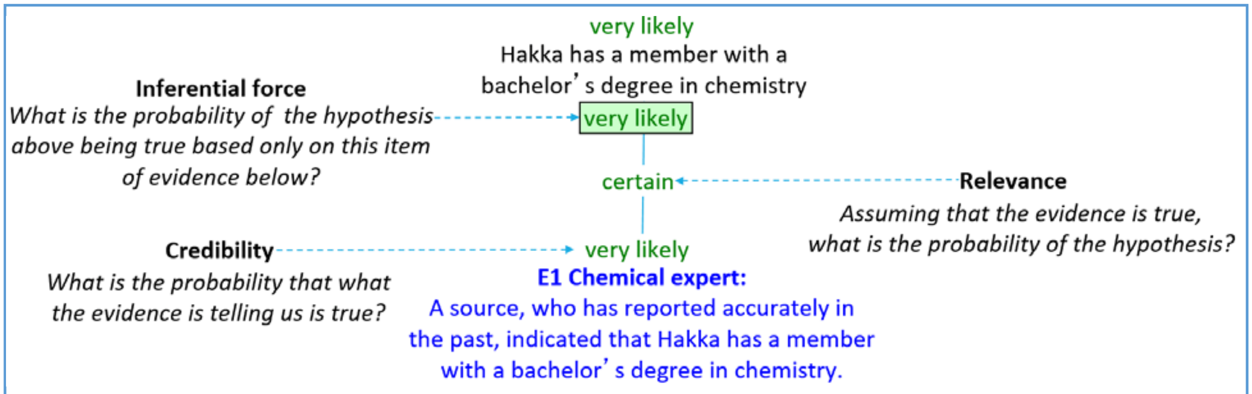
very likely
Hakka has a member with a bachelor's degree in chemistry
Inferential force
What is the probability of the hypothesis above being true based only on this item of evidence below?
very likely
certain
Relevance
Assuming that the evidence is true, what is the probability of the hypothesis?
Credibility
What is the probability that what the evidence is telling us is true?
very likely
E1 Chemical expert:
A source, who has reported accurately in the past, indicated that Hakka has a member with a bachelor's degree in chemistry.

Figure 54. Inferential force of evidence.

Therefore, the inferential force in this circumstance would be low. Similarly, it is not enough for the evidence item to be credible, if the information provided is not relevant to the hypothesis. The inferential force will be high only if the evidence item is both highly relevant and credible.

In this case, because we have only one item of evidence, the probability of the hypothesis "Hakka has a member with a bachelor's degree in chemistry" is given by the inferential force of this evidence item. However, if we have more items of evidence, some favoring the truthfulness of the hypothesis, and some disfavoring it, then the probability of the hypothesis will result from the combined inferential forces of all these items of evidence.

As another example, let's now consider the upper-level hypothesis "Hakka has expertise to develop chemical weapons" (see Figure 55). The relevance of "Hakka has a member with a bachelor's degree in chemistry" to this hypothesis was assessed as likely. Because the probability of the sub-hypothesis is very likely, its inferential force on the top hypothesis is likely, the minimum of its relevance and probability.

In this case we have just this one reason and one argument for the top hypothesis to be true. Therefore, the probability of the top hypothesis is the inferential force of this argument. In general, however, we may have multiple arguments, some favoring the truthfulness of the top hypothesis and some disfavoring it. In such a case the probability of the top hypothesis will be given by the combined inferential forces of all these arguments.

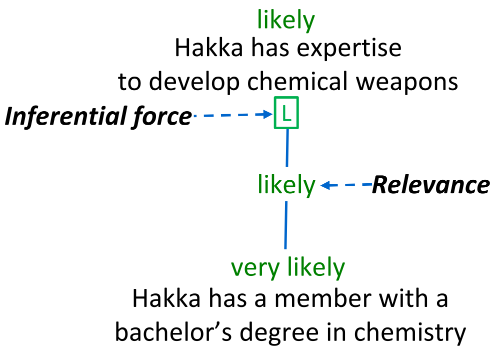
likely
Hakka has expertise to develop chemical weapons
Inferential force
L
likely
Relevance
very likely
Hakka has a member with a bachelor's degree in chemistry

Figure 55. Another relevance assessment.

6.2.6 Evidence-Based Hypothesis Assessment

In general, there may be several items of evidence that are relevant to a given hypothesis, some favoring it and some disfavoring it, each with a specific credibility, relevance, and inferential force. Figure 56 illustrates a situation where there are two items of evidence favoring hypothesis H (E1 and E2), and two items of evidence disfavoring it (E3 and E4). Each item of evidence has a specific inferential force on hypothesis H given by the smallest between its credibility and relevance. Thus, the inferential force of E1 is likely, that of E2 is almost certain, that of E3 is likely, and that of E4 is likely.

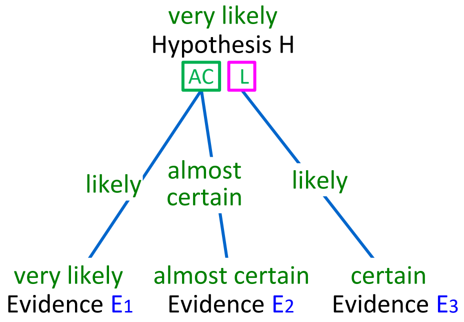

Figure 56. Hypothesis assessment based on the credentials of evidence.

The combined inferential force of the two favoring items of evidence (E1 and E2) is almost certain, the highest between the inferential force of E1 and that of E2. This combined inferential force is displayed in the left (green) box under the hypothesis H to indicate that it favors the truthfulness of H. The combined inferential force of the two disfavoring items of evidence (E3 and E4) is likely, the highest between the inferential force of E3 and that of E4. This combined inferential force is displayed in the right (pink) box under the hypothesis H to indicate that it disfavors the truthfulness of H.

The probability of the hypothesis H is determined by balancing the combined inferential force of the favoring evidence (almost certain), and the inferential force of the disfavoring evidence (likely).

The Baconian probability view (Cohen, 1977; 1989) requires considering either **H** or **not H** as probably true, but not both at the same time. To assess a hypothesis that has both favoring and disfavoring evidence, such as hypothesis H in Figure 56, we have introduced an <u>on-balance</u> function shown in Table 19 that balances the inferential force of the favoring evidence (almost certain) with that of the disfavoring evidence (likely), assessing a probability of more than likely for H.

Table 19. On-balance function.

		Inferential force of disfavoring evidence/arguments						
	H	lacking support	barely likely	likely	more than likely	very likely	almost certain	certain
Inferential force of favoring evidence/arguments	lacking support	lacking support	lacking support	lacking support	lacking support	lacking support	lacking support	lacking support
	barely likely	barely likely	lacking support	lacking support	lacking support	lacking support	lacking support	lacking support
	likely	likely	barely likely	lacking support	lacking support	lacking support	lacking support	lacking support
	more than likely	more than likely	likely	barely likely	lacking support	lacking support	lacking support	lacking support
	very likely	very likely	more than likely	likely	barely likely	lacking support	lacking support	lacking support
	almost certain	almost certain	very likely	more than likely	likely	barely likely	lacking support	lacking support
	certain	certain	almost certain	very likely	more than likely	likely	barely likely	lacking support

As indicated in the right and upper side of Table 19, if the inferential force of the disfavoring evidence is higher than or equal to that of the favoring evidence, then the hypothesis H is “lacking support.” If, however, the inferential force of the favoring evidence is strictly greater than that of the disfavoring evidence (and there is some force of the disfavoring evidence), then the probability of H is lowered, based on the inferential force of the disfavoring evidence (see the left and lower side of Table 19).

6.2.7 Assessing Complex Hypotheses

The previous section presented a simple way of directly assessing a hypothesis based on evidence. This works well when it is easy to assess the relevance of the evidence to the hypothesis. But it does not work well for assessing complex hypothesis. Such complex hypotheses are decomposed into simpler and simpler hypotheses, down to the level of very simple hypotheses for which the relevance of evidence can be more confidently assessed, as discussed in this section.

Evaluating the Basis for a Hypothesis

When building a favoring argument, we will be looking for conditions that would make a hypothesis true. Consider the argument:

- IF Hakka has a member with a bachelor’s degree in chemistry
- THEN Hakka has expertise to develop chemical weapons

However, if Hakka *does not have* a member with a bachelor’s degree in chemistry then it would be *incorrect* to conclude that Hakka *does not have* expertise to develop chemical weapons. Hakka may still have this expertise, for example, from someone who has learned by himself to build such weapons, without having a bachelor’s degree. From a strictly logical point of view, if the sub-hypothesis is not true we cannot conclude anything about the top hypothesis. Therefore, such an argument fragment is useful in our reasoning, only if the hypothesis “Hakka has a member with a bachelor’s degree in chemistry” is true (see Figure 57).

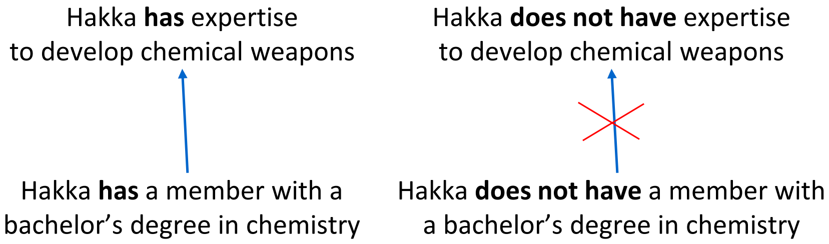
Hakka **has** expertise to develop chemical weapons
Hakka **does not have** expertise to develop chemical weapons
Hakka **has** a member with a bachelor’s degree in chemistry
Hakka **does not have** a member with a bachelor’s degree in chemistry

Figure 57. Examples of correct and incorrect arguments.

However, our arguments may contain negated hypotheses, as illustrated by the arguments in Figure 58

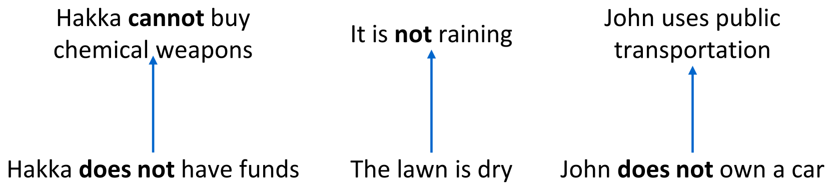
Hakka **cannot** buy chemical weapons
It is **not** raining
John uses public transportation
Hakka **does not** have funds
The lawn is dry
John **does not** own a car

Figure 58. Arguments with negated hypotheses.

The point we want to make here is that we will always be interested in showing that some hypothesis is true or some negated hypothesis is true. In the following we will investigate various types of arguments.

Simple Argument

Let us consider the hypothesis "Hakka has chemical weapons." What condition would make this hypothesis true? IF Hakka develops chemical weapons THEN Hakka has chemical weapons.

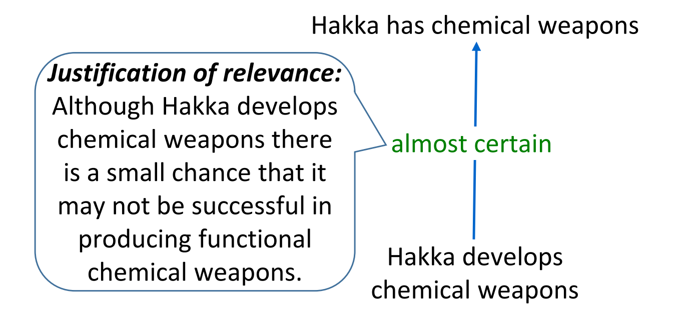
Hakka has chemical weapons
Justification of relevance: Although Hakka develops chemical weapons there is a small chance that it may not be successful in producing functional chemical weapons.
almost certain
Hakka develops chemical weapons

Figure 59. Simple argument with justification of relevance.

You also need to assess the relevance of the sub-hypothesis to the hypothesis, by answering the question: *Assuming that Hakka does indeed develop chemical weapons, what is the probability that Hakka has chemical weapons?* If Hakka develops chemical weapons then we would expect it to have chemical weapons. However, there is a small chance that it may not be successful in producing functional chemical weapons. Therefore, you may assess the relevance of the "develops" sub-hypothesis as "almost certain." When assessing the relevance, it is recommended that you record the justification of your assessment in the argument (see Figure 59).

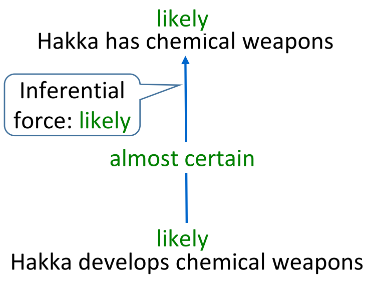
likely
Hakka has chemical weapons
Inferential force: likely
almost certain
likely
Hakka develops chemical weapons

Figure 60. The inferential force of a simple argument.

Once the probability of the sub-hypothesis is determined (for example likely), the system computes its inferential force as the smallest between its probability and its relevance. Thus, the inferential force of this simple argument is likely (see Figure 60). Because this is the only argument for the top hypothesis, the probability of the top hypothesis is given by the inferential force of this argument.

Alternative (OR) Arguments

Is there any other condition that would make the top hypothesis true? Yes, IF Hakka buys chemical weapons THEN it is certain that it has chemical weapons. This is an alternative argument for the top hypothesis "Hakka has chemical weapons" (see Figure 61).

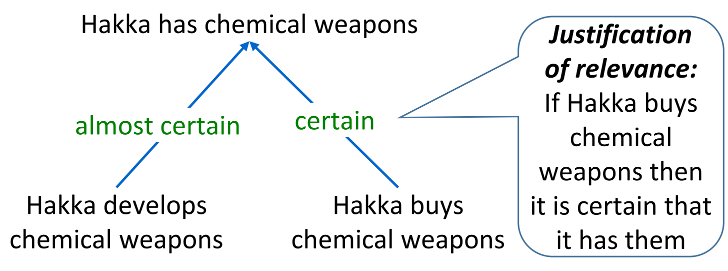
Hakka has chemical weapons
almost certain
certain
Hakka develops chemical weapons
Hakka buys chemical weapons
Justification of relevance: If Hakka buys chemical weapons then it is certain that it has them

Figure 61. An example of alternative (OR) arguments.

Let us assume that the probability of "Hakka buys chemical weapons" is barely likely. Then the inferential force of this alternative argument is also barely likely, the smallest between the probability of the sub-hypothesis and its relevance. Now we have two arguments for the truthfulness of "Hakka has chemical weapons", one with inferential force likely, and the other with inferential force barely likely. In such a case, the system assesses the probability of the top hypothesis as the highest of the two inferential forces. Thus, it is likely that "Hakka has chemical weapons" (see Figure 63). Here is a general rule: T*he probability is determined by the highest inferential force of the individual supporting arguments.*

Conjunctive (AND) Argument

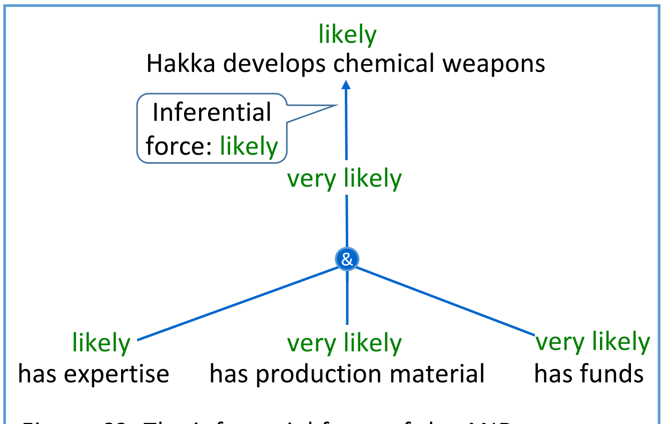

Figure 62. The inferential force of the AND argument.

It may be the case that multiple sub-hypotheses need to be true in order for the hypothesis to be true, as in the example from **Error! Reference source not found.**: IF Hakka has expertise AND production material AND funds THEN Hakka develops chemical weapons.

These three sub-hypotheses, taken together, represent an argument for the top hypothesis. You also need to assess the relevance of this AND argument, and it is recommended that you also justify your assessment of this relevance: There may be several additional conditions that are necessary to develop chemical weapons. Therefore, based on these indicators alone, we assess that it is only very likely that Hakka develops chemical weapons. The probability of the top hypothesis is determined by the inferential force of the AND argument which is the smallest among the probabilities of the sub-hypotheses and the relevance of the argument, as shown in Figure 62.

Indicator Argument

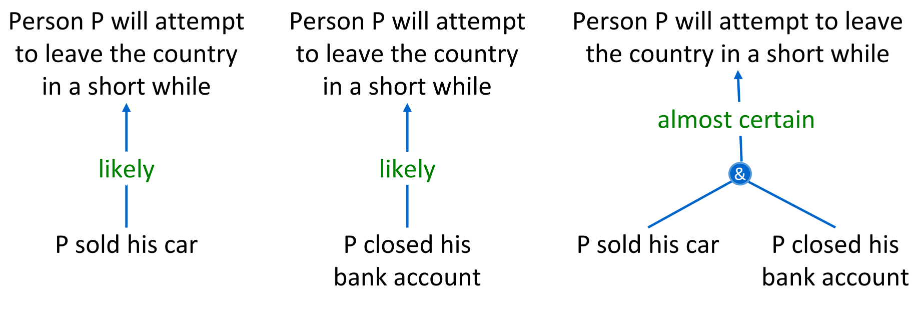
Person P will attempt to leave the country in a short while
likely
P sold his car
Person P will attempt to leave the country in a short while
likely
P closed his bank account
Person P will attempt to leave the country in a short while
almost certain
&
P sold his car
P closed his bank account

Figure 63. The relevances of individual indicators and of their combination.

Many times when we are assessing a hypothesis we only have a set of indicators. The more indicators are supported by evidence, the more likely the hypothesis is. As an example, consider Person P who has been under surveillance in connection with terrorist activities. We suspect that P will attempt to leave the country in a short while. Three days ago we received information that P sold his car. Today, we received information that he closed his account at his bank. Each of these is only a likely indicator of the hypothesis that P plans to leave the country. He could be planning to buy a new car, or he could be dissatisfied with his bank. But, taken together, these two indicators are almost certainly suggesting that P is planning to leave the country (see Figure 63).

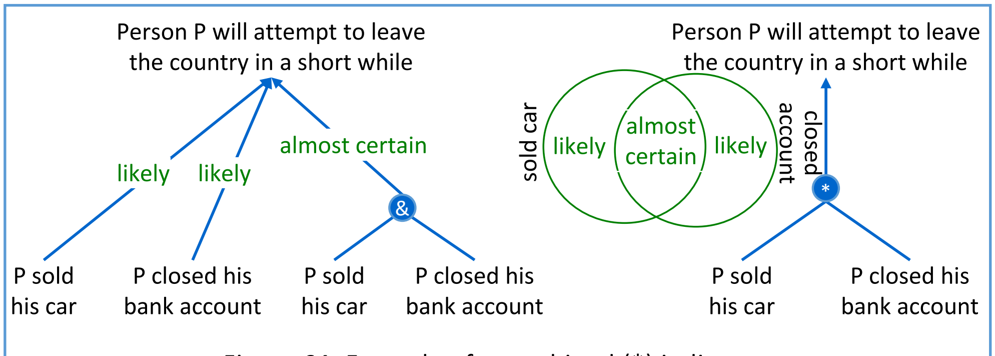
Person P will attempt to leave the country in a short while
likely
likely
almost certain
&
P sold his car
P closed his bank account
P sold his car
P closed his bank account
Person P will attempt to leave the country in a short while
sold car
likely
almost certain
likely
closed account
*
P sold his car
P closed his bank account

Figure 64. Example of a combined (*) indicator.

Notice that the indicator arguments from Figure 63 are, in fact, alternative favoring arguments of the hypothesis “Person P will attempt to leave the country in a short while”, as shown in the left-hand side of Figure 64. We can represent them more compactly by using the combined indicator operator, also called the “*” operator, as shown in the right-hand side of Figure 64.

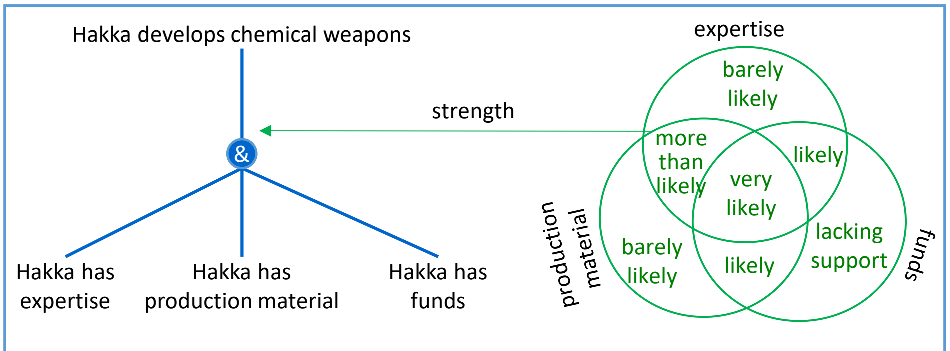

Figure 65. Another example of an indicator argument.

The combined indicator also alleviates a problem with the AND argument which requires all the sub-hypotheses to be true in order for the hypothesis to be true. The problem is that, sometimes, we do not have evidence for some of the subhypotheses, but might still have drawn some conclusions from individual conjuncts. In such cases we can replace the AND operator with a *operator, as illustrated in Figure 65. This indicator argument is a compact representation of seven alternative arguments with the following relevancies:

- The relevance of “expertise” alone is barely likely.
- The relevance of “production material” alone is barely likely.
- The relevance of “funds” alone is lacking support.
- The relevance of “expertise” AND “production material” is more than likely.
- The relevance of “production material” AND “funds” is likely.
- The relevance of “funds” AND “expertise” is likely.
- The relevance of “expertise” AND “production material” AND “funds” is very likely.

Disfavoring Evidence and Arguments

If you are rigorously developing and evaluating your argumentation you will be taking into account both favoring and disfavoring evidence and arguments for a given hypothesis. Let us consider the hypothesis: "Hakka buys chemical weapons."

A favoring argument or reason for this hypothesis (see Figure 68), whose relevance we judge as very likely is: "Hakka has funds." However there would be strong sanctions against the seller and we judge that this almost certainly will deter any potential seller.

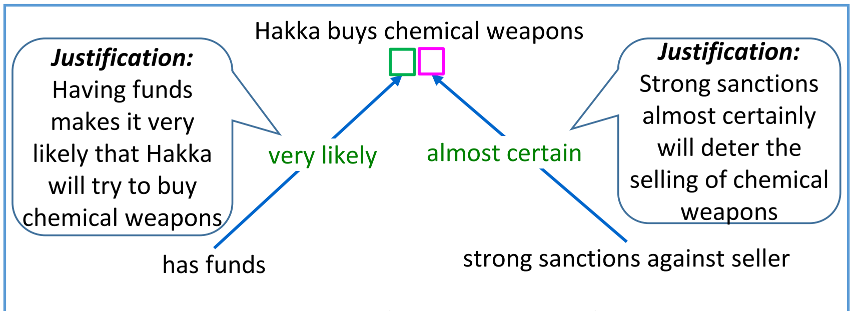
Hakka buys chemical weapons
Justification: Having funds makes it very likely that Hakka will try to buy chemical weapons
very likely
has funds
almost certain
Justification: Strong sanctions almost certainly will deter the selling of chemical weapons
strong sanctions against seller

Figure 66. A hypothsis with a favoring and a disfavoring argument.

Consider that very likely Hakka has funds. Then the inferential force of this favoring argument is computed as very likely. Now consider that, almost certain there will be strong sanctions against the seller. The inferential force of this disfavoring argument is computed as almost certain. As a result, based on the on-balance function from Table 19, the hypothesis "Hakka buys chemical weapons" is lacking support, as shown in Figure 67. Therefore, the evidence does not support the conclusion that Hakka buys chemical weapons, but it supports the conclusion that it does not buy them.

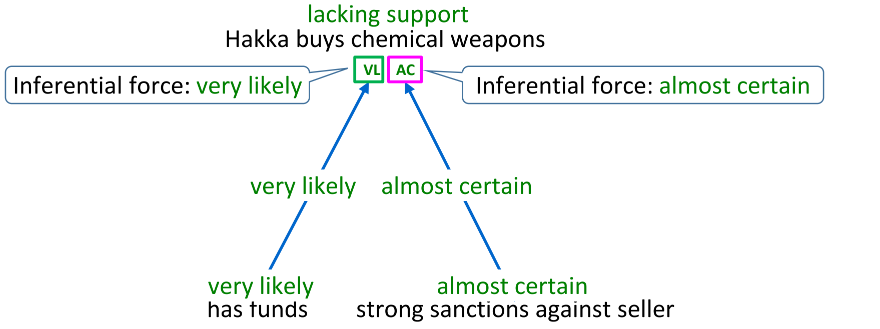
lacking support
Hakka buys chemical weapons
VL
AC
Inferential force: very likely
Inferential force: almost certain
very likely
almost certain
very likely
has funds
almost certain
strong sanctions against seller

Figure 67. Hypothesis assessment based on favoring and disfavoring arguments.

6.2.8 Summary of Augmented Wigmoran Argumentations

Figure 68 summarizes the hypothesis assessment process using an abstract example of an augmented Wigmorean argumentation which is probabilistic inferential network integrating logic with Baconian and Fuzzy probabilities. Hypothesis H is decomposed into simpler hypotheses by considering both favoring arguments (supporting the truthfulness of H), under the left (green) square, and disfavoring arguments (supporting the falsehood of H), under the right (pink) square. Each argument is an independent strategy of showing that H is true or false, and is characterized by a specific relevance or strength. The argument consists either of a single hypothesis (e.g., H_3) or a conjunction of hypotheses (e.g., H_1 & H_2). The sub-hypotheses from these arguments are further decomposed through other arguments, leading to simpler and simpler (sub-sub-)hypotheses that can be more accurately assessed based on evidence.

Consider, for example, sub-sub-hypothesis H_{2b}. There are two items of evidence relevant to this hypothesis, the favoring evidence item E^1_{2b}, and the disfavoring evidence item E^{21}_{2b}. Each item of evidence has three credentials that need to be assessed: credibility, relevance, and inferential force. The credibility of evidence answers the question: *What is the probability that the evidence is true?* The relevance of evidence to a hypothesis answers the question*: What would the probability of the hypothesis be if the evidence were true?* Based on these two credentials, the system computes the *inferential force or weight* of the evidence on the hypothesis that answers the question: *What is the probability of the hypothesis, based only on this evidence?* This is computed as the minimum between the credibility and relevance. For example, the inferential force of E^1_{2b} is almost certain (100%), that of E^{21}_{2b} is barely likely (50-55%).

The probability of sub-sub-hypothesis H_{2b} is determined as very likely (80-95%) by balancing the inferential force of the favoring evidence with that of the disfavoring evidence. Therefore, the probability of H is likely, based on the on-balance function from Table 19.

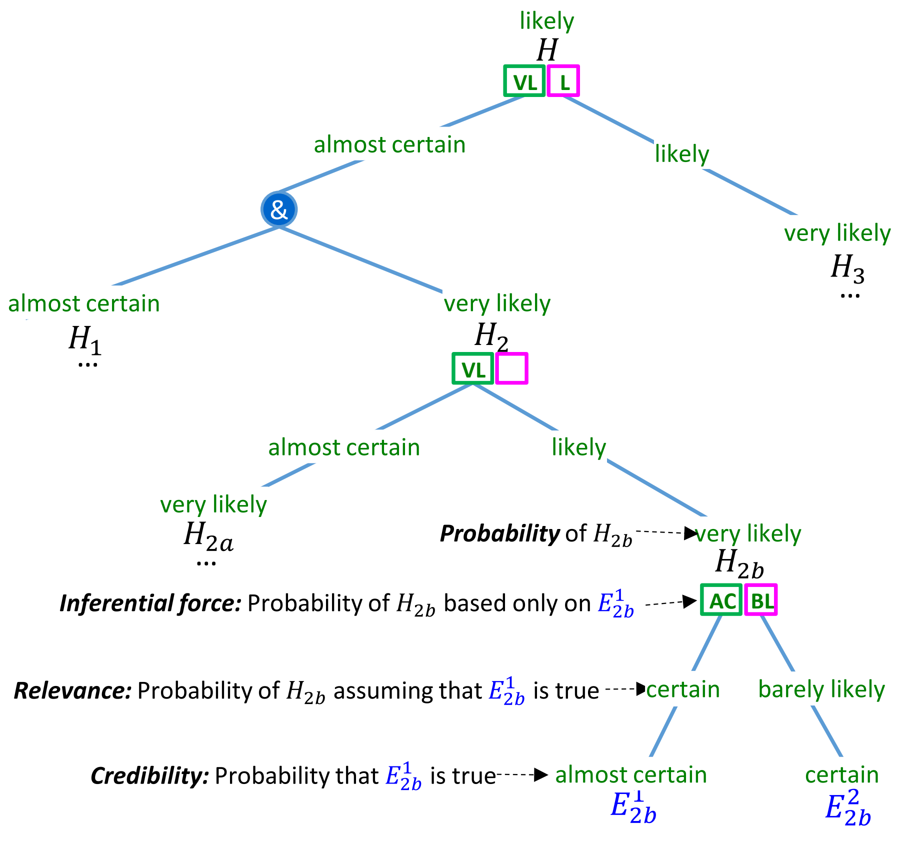
likely
H
VL
L
almost certain
likely
&
very likely
H_3
...
almost certain
H_1
...
very likely
H_2
VL
almost certain
likely
very likely
H_{2a}
...
Probability of H_{2b}
very likely
H_{2b}
Inferential force: Probability of H_{2b} based only on E^1_{2b}
AC
BL
Relevance: Probability of H_{2b} assuming that E^1_{2b} is true
certain
barely likely
Credibility: Probability that E^1_{2b} is true
almost certain
E^1_{2b}
certain
E^2_{2b}

Figure 68. Hypothesis assessment with Wigmorean argumentation.

Once the probabilities of the bottom-level hypotheses have been determined based on evidence, the probabilities of the upper level hypotheses are computed as explained below, based on the logical structure of the Wigmorean argumentation (conjunctions and disjunctions of hypotheses), using min-max probability combination rules common to the Fuzzy probability view

(Zadeh, 1983; Negoita and Ralescu, 1975; Schum 2001) and the Baconian probability view (Cohen, 1977; 1989; Schum, 2001). These rules are much simpler than the Bayes rule used in the Bayesian probability view (Schum, 2001), or the Dempster-Shafer rule in the Belief Functions probability view (Shafer, 1976). Hypothesis H_2 has two favoring arguments, H_{2a} with inferential force very likely (80-95%) and H_{2b} with inferential force very likely (80-95%). Its probability is very likely (80-95%), the maximum between the inferential force of H_{2a} (very likely, 80-95%) and that ofH_{2b} (likely, 55-70%).

Finally, hypothesis H has a favoring argument with inferential force very likely (the minimum of almost certain, almost certain and very likely), and a disfavoring argument with inferential force likely (the minimum of very likely and likely). Balancing these probabilities results in the probability of H as being assessed as likely. The augmented Wigmorean argumentation for the problem in Table 17 is shown in Figure 69.

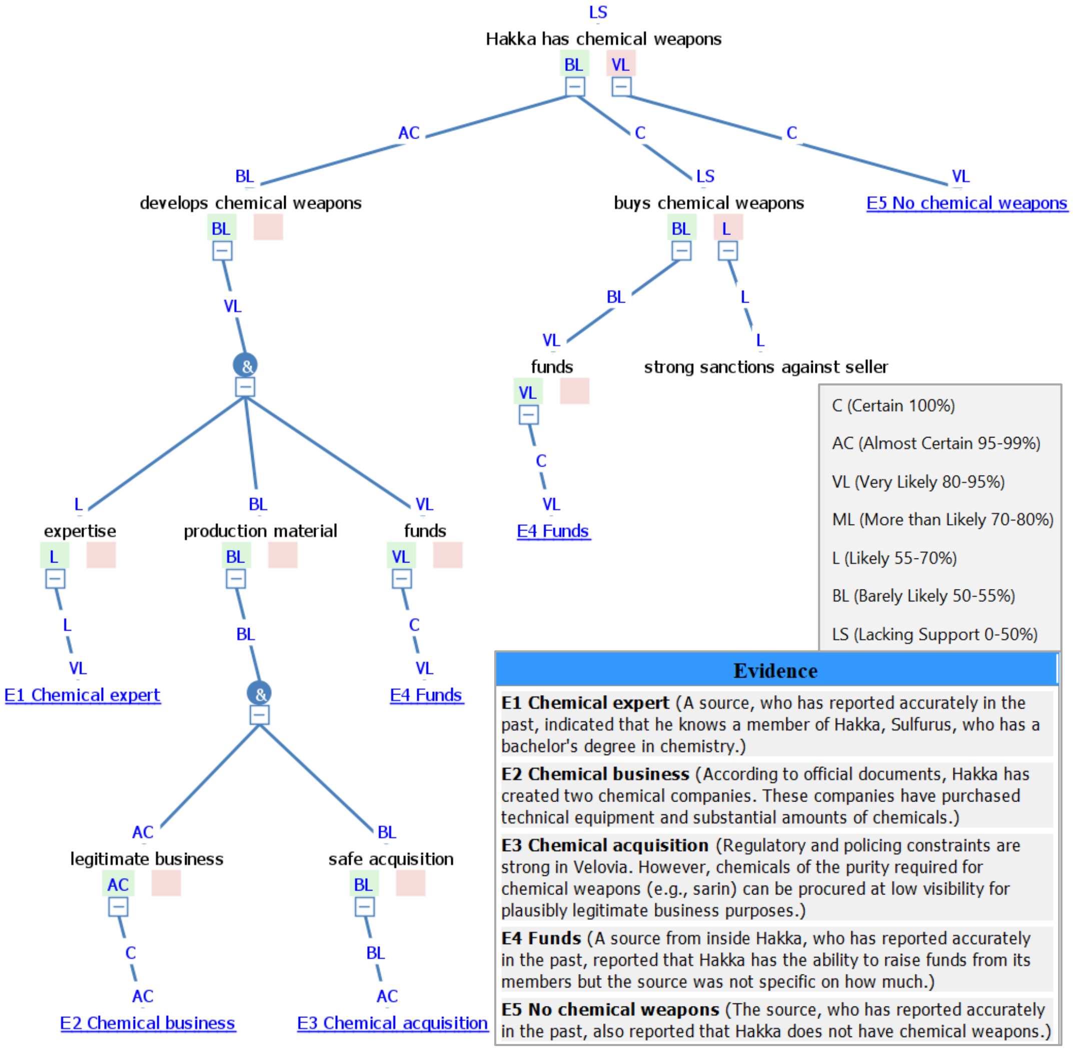

E1 Chemical expert (A source, who has reported accurately in the past, indicated that he knows a member of Hakka, Sulfurus, who has a bachelor's degree in chemistry.)

E2 Chemical business (According to official documents, Hakka has created two chemical companies. These companies have purchased technical equipment and substantial amounts of chemicals.)

E3 Chemical acquisition (Regulatory and policing constraints are strong in Velovia. However, chemicals of the purity required for chemical weapons (e.g., sarin) can be procured at low visibility for plausibly legitimate business purposes.)

E4 Funds (A source from inside Hakka, who has reported accurately in the past, reported that Hakka has the ability to raise funds from its members but the source was not specific on how much.)

E5 No chemical weapons (The source, who has reported accurately in the past, also reported that Hakka does not have chemical weapons.)

Figure 69. Augmented Wigmorean argumentation for the problem in Table 17.

6.3 Analytic Bias

The topic of bias comes up repeatedly in works on intelligence analysis. Much of the discussion is based on research performed by psychologists decades ago. As a result of this research, psychologists have made some depressing and quite unreasonable claims about the extent of our inferential rationality. Conclusions reached and reported in this research have formed an important basis for the very influential work of R. J. Heuer (1999).

In this section we will discuss different biases which have been identified in intelligence analysis and how employing Wigmorean argumentations can help recognize and partially counter them. We will begin by examining various meanings attached to the term bias and to its various origins and possible species and relations. Also important are possible value-related consequences for the persons whose views are labeled as being biased in various ways. Reading other existing works on bias in intelligence analysis, many persons are likely to conclude that the only origins of bias are the intelligence analysts themselves. But there are other important origins of possible bias including the sources of intelligence evidence (i.e., HUMINT), persons in chains of custody of intelligence evidence, and the policy-making customers of intelligence analyses. These additional classes of sources have their own species of bias.

6.3.1 Basic Interpretations of the Term "Bias"

The term *bias* arises in a variety of different contexts, some of which do not concern us such as in dressmaking and in the game of bowls. A dressmaker is said to make cuts along the bias, meaning that he/she makes oblique cuts across the warp of a fabric. In the game of bowls, the swerving course of a bowl when thrown or lagged, is termed bias. Bias does, of course, have some technical uses such as in statistics, machine learning, and engineering. In statistics and machine learning we speak of a biased result if it is distorted in some way and arises from a neglected factor or an approximate model learned. In electrical engineering, one form of bias refers to steady voltages applied to an electronic device to stabilize its operation, or to minimize distortions in recordings. *But our interests concern the use of the term bias with reference to people's views, beliefs, opinions, and related behaviors*; this use began in the mid 16th Century (Chantrell, 2004, p.52). The term bias comes from the French: *biais*. This word has its origin in the Greek: *epikarsios*, meaning "oblique."

The question of interest to us is: *What is meant by the term bias when it is applied to peoples' views, beliefs, opinions, and related behaviors?* One place to begin answering this question is by considering words that have been used as synonyms for the word bias. First, some meanings that have been commonly associated with the term bias are: prejudice, partiality, partisanship, favoritism, unfairness, one-sidedness, bigotry, intolerance, discrimination, leaning tendency, inclination, and predilection (Lindberg, 2004, p.82).

The term bias often occurs in the field of law. As happens on so many other occasions, the interpretations of evidential and inferential concepts in legal contexts are very useful in intelligence analysis. *Black's Law Dictionar*y provides several interpretations of the term bias (Black, 1968, p.205): inclination, bent, pre-conceived opinion, a predisposition to decide a cause or an issue in a certain way which does not leave the mind perfectly open to conviction, and the inability to judge a matter impartially in a particular case. As we expect, all involved in a legal dispute: the parties in the dispute, their advocates, judges, and fact-finders, have their own particular biases. The most well-known description of biases in intelligence analysis is that of Heuer (Heuer, 1999, pp.111-171) who defines them as consistent and predictable mental errors caused by our simplified information processing strategies (Heuer, 1999, p.111).

So, what can we do to reduce or eliminate recognized biases? We think that the best protection against biases in an intelligence analysis comes from the collaborative effort of teams of analysts, who become skilled in the evidential and argumentational elements of their tasks, and who are willing to share their insights with colleagues, who are also willing to listen. Employing a systematic approach to intelligence analysis which is based on scientific reasoning with evidence, which makes explicit all the reasoning steps, probabilistic assessments, and assumptions, so that they can be critically analyzed and debated, is the best protection against biases. That is why the use of Wigmorean argumentation and of an analytic tool like Cogent, which helps the analyst perform such an analysis, is one way to recognize and counter biases.

In the next section we will review the analysts' biases discussed by Heuer, as well as other widely known biases. After that we will present three other origins of bias that are rarely discussed, even though they may be at least as important on occasion as any analysts' biases. As we review various types of biases we also discuss how the use of Wigmorean argumentation helps identify and reduce them. But before we proceed, let us mention that it would be quite impossible for anyone to list all the biases that can occur since people will always find new ways to be prejudiced, one-sided, and partisan.

6.3.2 Biases of the Analyst

Analysts, like other persons, have preferences for certain kinds of evidence and these preferences can induce biases. In particular, analysts can have a distinct preference for vivid or concrete evidence when less vivid or concrete evidence may be more inferentially valuable. In addition, their personal observations may be over-valued.

First, the hypothesis in search of evidence phase of the analysis (see Section 1.5.2) helps identify a wide range of evidentiary needs. Second, performing a detailed and systematic evaluation of the relevance and credibility of each item of evidence, *regardless of its "vividness",* is helping us be more objective in the evaluation of the infrential force of evidence.

The absence of evidence bias concerns a failure to consider the degree of completeness of the available evidence. A Wigmorean argumentation would show very clearly that threre are sub-hypotheses for whch there is no relevant evidence. The oversensitivity to evidence consistency bias can easily manifest when using an analytic tool like Heuer's ACH (Heuer, 2008) where the analyst judges alternative hypotheses based on evidence, without building any argumentation. A Wigmorean argumentation would reveal if most of the evidence is only relevant to a small fraction of sub-hypotheses, while many other sub-hypotheses have no evidentiary support. If Heuer would had written his book in 2003, he might have used the case of Curveball as a very good example of the persistence of impressions based on discredited evidence bias (Drogin, 2007). In this case, Curveball's evidence was discredited on a number of grounds but was still believed and taken seriously by some analysts, as well as many others.

Wigmorean argumentations help countering this bias by incorporating in the argumentation an explicit analysis of the credibility of evidence, especially for key evidence that has a direct influence on the analytic conclusion. When such an evidence item is discredited, specific elements of its analysis are updated, and this leads to the automatic updating of the probability of each hypothesis to which it is relevant. For example, the credibility of the observations performed by a source (such as Curveball) depends on source's competence, vercity, objectivity and observational sensitivity under the conditions of observation. Moreover, competence depends on access and understandability. Thus, the bias that would result from the persistence of impressions based on discredited evidence is countered by a rigorous, detailed and explicit credibility analysis.

But there are additional biases in the evaluation of evidence that Heuer does not mention, particularly with respect to establishing the credentials of evidence: relevance, credibility, and inferential force or weight. An analyst may focus on the veracity of the source and ignore source's competence, objectivity and observational sensitivity. Analysts may fail to recognize possible synergisms in convergent evidence, as happened in the 9/11/2001 disaster. Analysts may even overlook evidence having significant inferential force.

Biases in the perception of cause and effect arise when analysts assign causal relations to the occurrence of events and phenomena that are actually accidental or random in nature. One related consequence is that analysts often overestimate their ability to predict future events from past events, because there is no causal association between them. One major reason for these biases is that analysts may not have the requisite level of understanding of the kinds and amount of information necessary to infer a dependable causal relationship. According to Heuer, when feasible, the "increased use of scientific procedures in political, economic, and strategic research is much to be encouraged", to counter these biases (Heuer, 1999, p.128), which is precisely what the computational model of intelligence analysis discussed in Section 1.5 does.

Now, here is something that can occur in any analysis concerning chains of reasoning. It is always possible that an analyst's judgment will be termed biased or fallacious, on structural grounds if it is observed that this analyst frequently leaves out important links in his/her chains of reasoning. This is actually a common occurrence since, in fact, there is no such thing as a uniquely correct or perfect argument. Someone can always find alternative arguments to the same hypothesis; what this says is that there may be entirely different inferential routes to the same hypothesis. Another possibility is that someone may find arguments based on the same evidence that lead to different hypotheses. This is precisely why there are trials at law; the prosecution and defense will find different arguments, and tell different stories, from the same body of evidence.

Biases in estimating probabilities (Heuer, 1999, p.122) are due to the difficulty of the human mind on coping with complicated probabilistic relationships. People tend to employ simple rules of thumb that reduce the burden of processing such information, but will also introduce errors. As discussed in Section 4.6, there are different views among probabilists on how to assess the force of evidence. The view of probability that Heuer assumes is the conventional or Kolmogorov view of probability discussed in Section 5.4. This is also the only view of probability considered by Heuer's sources of inspiration on biases: Daniel Kahneman, Amos Tversky, and their many colleagues in psychology (Kahneman and Tversky, 1974; Kahneman et al., 1982). In his writings, Kolmogorov makes it abundantly clear that his axioms apply only to instances in which we can determine probabilities by counting. So, also clearly, Kolmogorov probabilities apply only to replicable, repeatable, or enumerative phenomena, those we can observe over and over again. Since Heuer only considers numerical probabilities conforming to the Kolmogorov axioms, any biases associated with them (e.g., using the availability rule, the anchoring strategy, expressions of uncertainty, assessing the probability of a scenario) are either irrelevant or not directly applicable to a type of analysis that is based on different probability systems

But Heuer also notes that intelligence analysis usually deals with one-of-a-kind situations for which there are never any statistics. In such cases, analysts resort to subjective or personal numerical probability expressions. He discusses several reasons why verbal assessments of probability are frequently criticized for their ambiguity and misunderstanding. In his discussion he recalls Sherman Kent's advice that verbal assessments should always be accompanied by numerical probabilities (Kent, 1994). Heuer obviously agrees with Kent's advice. Then further Heuer notes (1999, p.123): "There seems to be little an analyst can do about this, short of breaking the analytical problem down in a way that permits assigning probabilities to individual items of information, and then using a mathematical formula to integrate these separate probability judgments." A Wigmorean argumentation achieves precisely what Heuer imagined that could be done for countering this bias. It breaks a hypothesis into simpler hypotheses and assesses the simpler hypotheses based on evidence. The probabilities are expressed in words rather than numbers, and are combinedusing simple min/max rules.

Analysts often overestimate the accuracy of their past judgments; customers often underestimate how much they have learned from an intelligence report; and persons who conduct post-mortem analysis of an intelligence failure will judge that events were more readily foreseeable than was in fact the case. They all have hindsight biases in evaluating intelligence reporting. As Heuer (1999, p.162) notes, they "take their current state of knowledge and compare it with what they or others did or could or should have known before the current knowledge was received. This is in sharp contrast with intelligence estimation, which is an exercise in foresight, and it is the difference between these two modes of thought—hindsight and foresight—that seems to be a source of bias. … After a view has been restructured to assimilate the new information, there is virtually no way to accurately reconstruct the pre-existing mental set."

Apparently Heuer did not envision the use of a Wigmorean argumentation that keeps track of the performed analysis, what evidence we had, what assumptions we made and what were their justifications, and what was the actual logic of our analytic conclusion. We can now add additional evidence and use our hindsight knowledge to restructure the argumentation and re-evaluate our hypotheses, and we can compare the hindsight analysis with the foresight one. But we will not confuse them. As indicated by Heuer (1999, pp.166-167): "A fundamental question posed in any postmortem investigation of intelligence failure is this: Given the information that was available at the time, should analysts have been able to foresee what was going to happen? Unbiased evaluation of intelligence performance depends upon the ability to provide an unbiased answer to this question.**"**

The confirmation bias**,** the most common bias according to Pherson and Boardman (2017), is the tendency to seek only that information that is consistent with the lead hypothesis, judgment, or conclusion; accept "confirming" evidence at face value while subjecting "disconfirming" evidence to critical scrutiny; and interpret ambiguous information to support the desired hypothesis. A Wigmorean argumentation would make it very clear that:

- The evidence is pertaining to only one aspect of the analysis i.e., there is only favoring evidence, and no (or very little) disfavoring evidence, or vice versa. This would suggest that the analyst is potentially trying to seek only information that favors their hypothesis.
- The analyst assigns higher probabilities to favoring evidence and lower probabilities to disfavoring evidence, or vice versa, suggesting that the analyst has a tendency of readily accepting supporting evidence, while being less receptive to and heavily scrutinizing the counter indicative evidence.
- The analyst is trying to interpret less relevant or ambiguous information as supporting his/her hypothesis. This implies the analyst sees all information as though it confirms

his/her preferred hypothesis, and does not take into account that the information might actually not be as relevant.

The satisficing bias (choosing the first hypothesis that appears good enough rather than carefully identifying all possible hypotheses and determining which one is the most consistent with the evidence) is signaled when the user has analyzed only one of the possible hypotheses, ignoring its alternatives. It is also signaled when several hypotheses are analyzed, but one of them has a significantly larger argumentation.

The conjunction fallacy (assuming that a conjunction of conditions is more probable than one of them) cannot be made in a Wigmorean argumentaton where the probability of a conjunction is computed by using the min rule: P(A and B) = minimum{P(A), P(B)}.

6.3.3 Some Frequently Overlooked Origins of Bias

So much of the discussion of bias in intelligence analysis is directed at intelligence analysts themselves. But we have identified three other origins of bias that are rarely discussed, even though they may be at least as important on occasion as any analysts' alleged biases. The three other origins of bias we will consider are:

- Persons who provide testimonial evidence about events of interest (i.e., HUMINT sources);
- Other intelligence professionals having varying capabilities who serve as links in what we term "chains of custody" linking the evidence itself, as well as its sources, with the users of evidence (i.e., the analysts);
- The "consumers" of intelligence analyses (government and military officials who make policy and decisions regarding national security).

HUMINT Sources

Our concern here is with persons who supply us with testimonial evidence consisting of reports of events about matters of interest to us. Heuer (1999, p.122) does mention the "bias on the part of the ultimate source", but he does not analyze it. In our work on evidence in a variety of contexts, we have always been concerned about establishing the credibility of its sources, particularly when they are human witnesses, sources, or informants (Schum, 1994/2001a). In doing so, we have made use of the 600 year-old legacy of experience and scholarship in the Anglo-American adversarial trial system concerning witness credibility assessments. As discussed in Section 4, we have identified the three major attributes of the credibility of ordinary witnesses: veracity, objectivity, and observational sensitivity. We will show how there are distinct and important possible biases associated with each such credibility attribute.

As discussed above, assessing the credibility of a human source $\mathcal{S}$ involves assessing $\mathcal{S}$'s veracity, objectivity, and observational sensitivity. We have to consider that source $\mathcal{S}$ can be biased concerning any of these attributes. On veracity, $\mathcal{S}$ might prefer to tell us that event **E** occurred, whether $\mathcal{S}$ believed **E** occurred or not. As an example, an analyst evaluating $\mathcal{S}$'s evidence **E*** might have evidence about $\mathcal{S}$ suggesting that $\mathcal{S}$ would tell us that **E** occurred because $\mathcal{S}$ wishes to be the bearer of what $\mathcal{S}$ believes we will regard as good news that event **E** occurred. On objectivity, $\mathcal{S}$ might choose to believe that **E** occurred because it would somehow be in $\mathcal{S}$'s best interests if **E** did occur. On observational sensitivity, there are various ways that $\mathcal{S}$'s senses could be biased in favor of recording event **E**; clever forms of deception supply examples.

These three species of bias possible for HUMINT sources must be considered by analysts attempting to assess the credibility of source $\mathcal{S}$ and how much weight or force $\mathcal{S}$'s evidence **E*** should have in the analyst's inference about whether or not event **E** did happen. The existence of any of these three biases would have an effect on an analyst's assessment of the weight or force of $\mathcal{S}$'s report **E***. As we know, all assessments of the credibility of evidence rest upon available evidence about its sources. In the case of HUMINT we need ancillary evidence about the veracity, objectivity, and observational sensitivity of its sources. In the process, we have to see whether any such evidence reveals any of the three biases just considered. Cogent supports the analyst in this determination by guiding her to answer specific questions based on ancillary evidence. The veracity questions to be considered are shown in Table 6 (p. 134), the objectivity questions are shown in Table 7 (p.136), and the observational sensitivity questions are shown in Table 8 (p.137) .

Persons in Chains of Custody of Evidence

Unfortunately, there are other persons, apart from HUMINT sources, whose possible biases need to be carefully considered. We know that analysts make use of an enormous variety of evidence that is not testimonial or HUMINT, but is tangible in nature. Examples include objects, images, sensor records of various sorts, documents, maps, diagrams, charts, and tabled information of various kinds.

But the intelligence analysts only rarely have immediate and first access to HUMINT assets or informants. They may only rarely be the first ones to encounter an item of tangible evidence. What happens is that there are several persons who have access to evidence between the times the evidence is first acquired and when the analysts first receive it. These persons may do a variety of different things to the initial evidence during the time they have access to it. In law, these persons constitute what is termed a "chain of custody" for evidence.

Heuer (1999, p.122) mentions the "distortion in the reporting chain from subsource through source, case officer, reports officer, to analyst" but he does not analyze it. In criminal cases in law, there are persons identified as "evidence custodians", who keep careful track of who discovered an item of evidence, who then had access to it and for how long, and what if anything they did to the evidence when they had access to it.

These chains of custody add three major additional sources of uncertainty for intelligence analysts to consider, that are associated with the persons in chains of custody whose credibility needs to be considered. The first and most important question involves *authenticity*: *Is the evidence received by an analyst exactly what the initial evidence said and is it complete?* The other questions involve assessing the *reliability* and *accuracy* of the processes used to produce the evidence if it is tangible in nature, or also used to take various actions on the evidence in a chain of custody, whether the evidence is tangible or testimonial. As an illustration, consider the chain of custody from 0 (p. 147), concerning an item of testimonial HUMINT coming from a foreign national whose code name is "Wallflower", who does not speak English. Wallflower gives his report to *case officer* Bob. This report is *recorded* by Bob and then *translated* by Husam. Then, Wallflower's translated report is *transmitted* to a *report's officer* Marsha who *edits* it and *transmits* it to the analyst Clyde who evaluates it and assesses its weight or force.

Now, here is where forms of bias can enter that can be associated with the persons involved in these chains of custody. The case officer Bob might have intentionally overlooked details in his recording of Wallflower's report. The translator Husam may have intentionally altered or deleted parts of this report. The report's officer Marsha might have altered or deleted parts of the translated report of Wallflower's testimony in her editing of it. The result of these actions is that the analyst Clyde receiving this evidence almost certainly did not receive an authentic and complete account of it, nor did he receive a good account of its reliability and accuracy. What he received was the transmitted, edited, translated, recorded testimony of Wallflower. Schum et al. (2009) show how Cogent may determine the credibility of the evidence received by the analyst. Although the information to make such an analysis may not be available, the analyst should adjust the confidence in his conclusion, in recognition of these biases.

Consumers of Intelligence Analyses

The policy-making consumers or customers of intelligence analysts are also subject to a variety of inferential and decisional biases that may influence the reported analytic conclusions. As is well known, the relationships between intelligence analysts and governmental policy makers are much discussed and involve considerable controversy (George and Bruce, 2008; Johnston, 2005). On the one hand we hear intelligence professionals say that they do not make policies but only try to help policy makers be as informed as they can be when they do form policies and make decisions in the nation's best interests. But we also learn facts about the intelligence process that complicate matters. An intelligence analysis is usually a hierarchical process involving many intelligence officers, at various grade levels, who become involved in producing an intelligence "product." At the most basic level of this hierarchy are the so-called "desk analysts" who are known and respected experts in the specific subject matter of the analysis at hand. An analysis produced by one or more desk analysts is then passed "upward" through many administrative levels, at each of which persons at these higher levels can comment on the desk analysts' report. It is often recognized that the higher an editor is in this hierarchy, the more political his/her views and actions become that may affect the content and conclusions of the analysis at hand. As this "upward" process continues, the analysis that results may be quite different from the one produced by the desk analysts, reflecting the biases of those who have successively edited it. In some cases, these editing biases are the direct result of the consumer's biases who may wish to receive a certain analytic conclusion. Using a system like Cogent that shows very clearly how the analytic conclusion is rooted in evidence would significantly help in reducing the above biases.

6.3.4 Biases and the Evaluation of Analysts

We must also be concerned about the consequences to a person, say an intelligence analyst, of being identified as displaying a bias of some sort. In the previous sections we considered an array of biases identified by psychologists as being ones to which none of us are allegedly immune to having on occasion. For example, some of these biases are said to involve the numerical probabilities we might use to hedge conclusions about the hypotheses we assert as a result of an analysis of evidence. Here comes analyst $\mathcal{A}$, whose assigned probability to hypothesis $\mathbf{H_K}$ is very high, say $P(\mathbf{H_K}) = 0.95$. $\mathcal{A}$ is now labeled biased by critics since they provide good arguments that $\mathcal{A}$ has been one-sided or narrow-minded in his present analysis of evidence concerning the hypotheses of interest. $\mathcal{A}$'s probabilistic assessment is not taken seriously and may even be the object of scorn among colleagues whose considerably smaller assigned probabilities to $\mathbf{H_K}$ are the ones reported to an interested customer who concludes that hypothesis $\mathbf{H_K}$ is not true. But time passes and it is discovered that hypothesis $\mathbf{H_K}$ true, much to the distress of the customer. This raises some interesting and difficult issues concerning the relation between bias and error.

The question is, *"Are all demonstrably biased judgments necessarily erroneous?"* In our example, Analyst $\mathcal{A}$'s biased judgment initially invited criticism. But because it was more correct than the judgments of other analysts, should it now invite praise? The issues raised in such instances are value-related. Suppose it is argued that, since $\mathcal{A}$'s judgments are frequently the result of one-sided or narrow-minded analyses, $\mathcal{A}$ was only lucky in the case of $\mathcal{A}$'s inferences regarding hypothesis $\mathbf{H_K}$, and therefore $\mathcal{A}$ deserves no praise. So, the answer to the above question seems to be: a biased judgment does not entail that it is necessarily erroneous. Whether a biased judgment that happens to be correct deserves praise involves some difficult choices.

There are some facts about the world that add great complexity to intelligence analysis and bear on the relations between bias and error. First, the world is not stationary and new things happen all the time. As a result, discovery in intelligence analysis is continuous and never ceases. We learn new things all the time. Beliefs about some hypothesis regarded as being very likely or unlikely a short time ago are overtaken by events that occurred just today. This means that intelligence analysis is a seamless process involving mixtures of three basic forms of reasoning: abductive (imaginative or insightful), deductive, and inductive. These mixtures of reasoning form one of the basic features of Cogent. Here is an example of how continuing discovery bears on bias and error.

There is a reason why Analyst $\mathcal{A}$'s high probability for $\mathbf{H_K}$ was criticized as being biased because $\mathcal{A}$'s analysis was one-sided or narrow-minded. Critics noted that $\mathcal{A}$ ignored even considering events **E**, **F**, and **G**, which, if they occurred, would be evidence against hypothesis $\mathbf{H_K}$. Perhaps $\mathcal{A}$ preferred not to take into account events that would make $\mathcal{A}$'s favored hypothesis less likely. These same critics either assumed or had evidence for some or all of events **E**, **F**, and **G**; this is why their assessments of the numerical probability of $\mathbf{H_K}$ were so much smaller than $\mathcal{A}$'s $P(\mathbf{H_K}) = 0.95$. But we have recently discovered evidence that events **E**, **F**, and **G** definitely did not occur; evidence of their occurrence was not credible for various reasons. And we also learned that $\mathbf{H_K}$ is true after all.

So, in light of these new discoveries, $\mathcal{A}$'s being one-sided or narrow-minded apparently worked to $\mathcal{A}$'s advantage in this instance. However, if these biases are routine characteristic of $\mathcal{A}$'s analyses, we would be entitled to be skeptical of $\mathcal{A}$'s probabilistic judgments. The critics' beliefs that $\mathcal{A}$ was simply lucky in the present analysis seem to have merit. Here is a view, quite reasonable but controversial, concerning how an intelligence analysis should be graded. *On this view, an analysis should be graded in terms of how well it was done and not whether it was correct or not.* This is precisely the view often taken in much of contemporary decision analysis (Clemen 1995, pp.3-4). According to this criterion, $\mathcal{A}$ should not be praised for the high probability $\mathcal{A}$ assigned to a true hypothesis, but criticized for the manner in which $\mathcal{A}$ inferred this high probability.

The above discussion of biases adds a strong argument in favor of using structured analytic methods, in the debate on how to significantly improve intelligence analysis (Marrin and Clemente, 2005; Marrin, 2011).

6.4 COGENT: Cognitive Agent for Cogent Analysis

Cogent is an intelligent system that assists an intelligence analyst solve typical intelligence analysis problems. It implements the computational theory of intelligence analysis described in this book. It builds on a series of analytical tools that includes Disciple-CD (Tecuci et al., 2016a), TIACRITIS (Tecuci et al., 2010; 2011a; 2011b), and Disciple-LTA (Tecuci et al., 2008a).Cogent enables a synergistic integration of an analyst's imagination and expertise with the computer's knowledge and critical reasoning.

6.4.1 Obtaining Cogent

Cogent is a research prototype implemented in JAVA and tested on PC and MAC. It is a stand-alone system that needs to be installed on the user's computer. For installation requirements and to download the system please go to http://lac.gmu.edu/Cogent/Download.html

6.4.2 Getting Started with Cogent

Figure 70 shows some of the panes from the interface of Cogent. The *Whiteboard area* displays the current analysis. The *Assistants area* shows several assistants, each helping in performing a group of related operations, such as building the argumentation or defining evidence. The bottom panes provide help in using the system.

6.4.3 Cogent Operations Help

This annotated Cogent interface in Figure 71 explains how to use the Operations Help. The top-left pane shows part of the current argumentation. When you click on a node, the "Operations Help" tab shows all the operations that can be performed at that node.

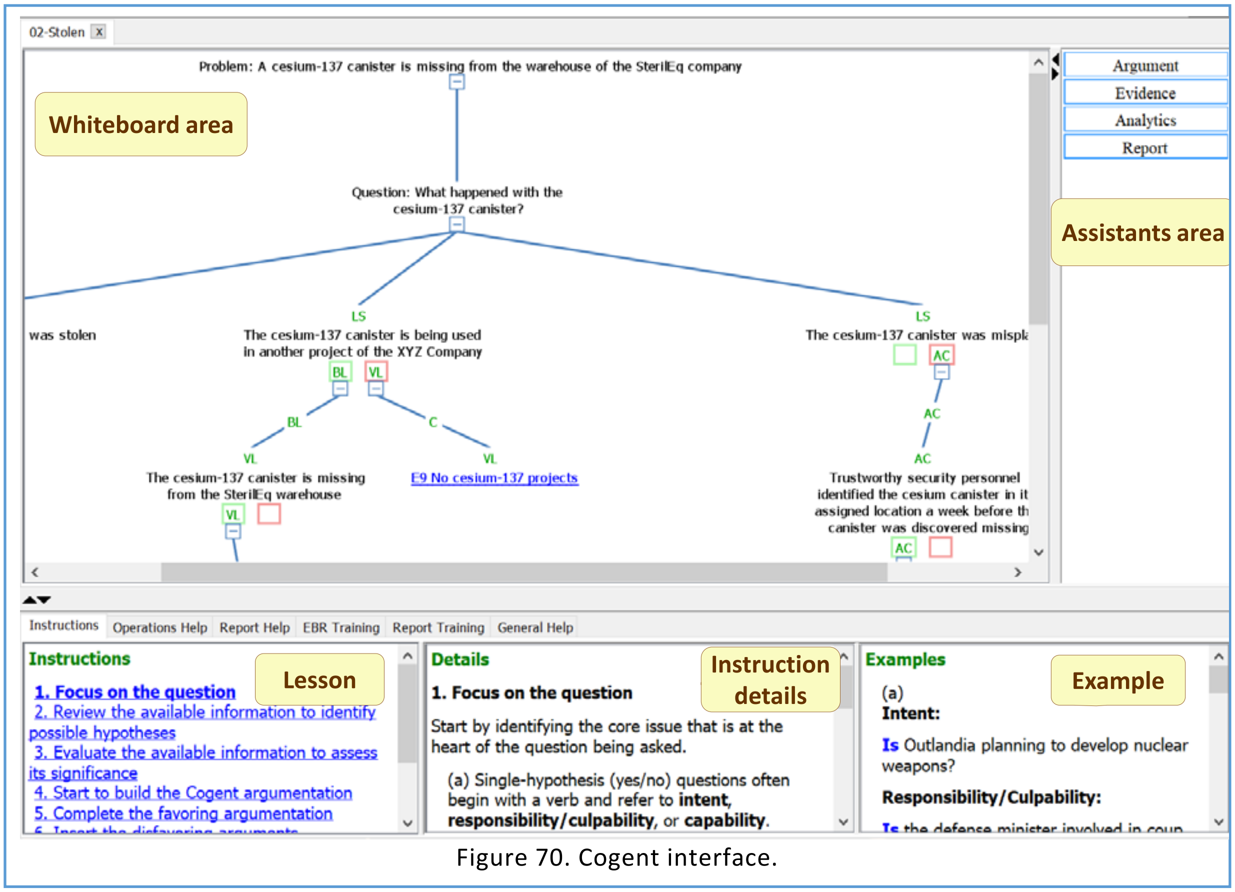

Problem: A cesium-137 canister is missing from the warehouse of the SterilEq company

Question: What happened with the cesium-137 canister?

was stolen

The cesium-137 canister is being used in another project of the XYZ Company

The cesium-137 canister was mispl

The cesium-137 canister is missing from the SterilEq warehouse

E9 No cesium-137 projects

Trustworthy security personnel identified the cesium canister in it assigned location a week before th canister was discovered missing

Instructions

1. Focus on the question
2. Review the available information to identify possible hypotheses
3. Evaluate the available information to assess its significance
4. Start to build the Cogent argumentation
5. Complete the favoring argumentation
6. Insert the disfavoring arguments

Details

1. Focus on the question

Start by identifying the core issue that is at the heart of the question being asked.

(a) Single-hypothesis (yes/no) questions often begin with a verb and refer to **intent**, **responsibility/culpability**, or **capability**.

Examples

(a)
Intent:

Is Outlandia planning to develop nuclear weapons?

Responsibility/Culpability:

Is the defense minister involved in coup

Figure 70. Cogent interface.

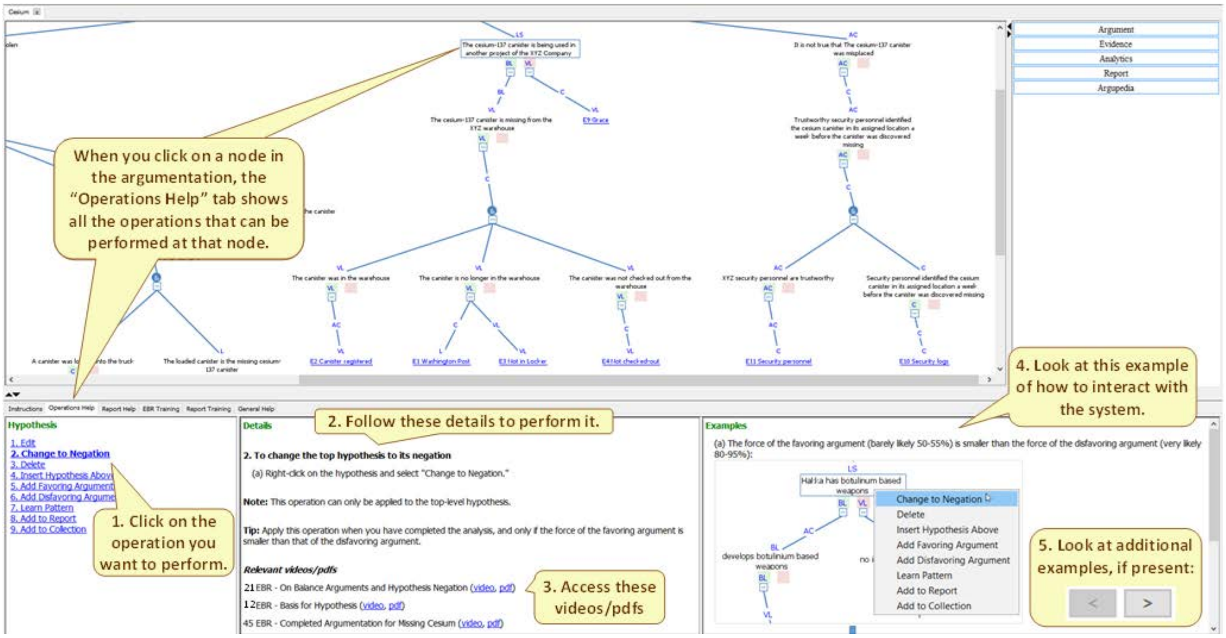
Cesium
Argument
Evidence
Analytics
Report
Argupedia
LS
The cesium-137 canister is being used in another project of the XYZ Company
AC
It is not true that The cesium-137 canister was misplaced
The cesium-137 canister is missing from the XYZ warehouse
E9 Grace
Trustworthy security personnel identified the cesium canister in its assigned location a week before the canister was discovered missing
The canister was in the warehouse
The canister is no longer in the warehouse
The canister was not checked out from the warehouse
XYZ security personnel are trustworthy
Security personnel identified the cesium canister in its assigned location a week before the canister was discovered missing
The loaded canister is the missing cesium-137 canister
E2 Canister registered
E1 Washington Post
E3 Not in Locker
E4 Not checked-out
E11 Security personnel
E10 Security logs
Instructions
Operations Help
Report Help
EBR Training
Report Training
General Help
Hal-ka has botulinum based weapons
develops botulinum based weapons
Change to Negation
Delete
Insert Hypothesis Above
Add Favoring Argument
Add Disfavoring Argument
Learn Pattern
Add to Report
Add to Collection
<
>

When you click on a node in the argumentation, the "Operations Help" tab shows all the operations that can be performed at that node.

Hypothesis

1. Edit
2. Change to Negation
3. Delete
4. Insert Hypothesis Above
5. Add Favoring Argument
6. Add Disfavoring Argume
7. Learn Pattern
8. Add to Report
9. Add to Collection

1. Click on the operation you want to perform.

Details

2. Follow these details to perform it.

2. To change the top hypothesis to its negation

(a) Right-click on the hypothesis and select "Change to Negation."

Note: This operation can only be applied to the top-level hypothesis.

Tip: Apply this operation when you have completed the analysis, and only if the force of the favoring argument is smaller than that of the disfavoring argument.

Relevant videos/pdfs

21 EBR - On Balance Arguments and Hypothesis Negation (video, pdf)
12 EBR - Basis for Hypothesis (video, pdf)
45 EBR - Completed Argumentation for Missing Cesium (video, pdf)

3. Access these videos/pdfs

Examples

4. Look at this example of how to interact with the system.

(a) The force of the favoring argument (barely likely 50-55%) is smaller than the force of the disfavoring argument (very likely 80-95%):

5. Look at additional examples, if present:

Figure 71. Operation help interface.

6.4.4 Cogent Practice

Power Plant Security

The United States has passed information to the country Upland that terrorists are planning to attack Upland's nuclear power plants in order to seize radioactive material for a "dirty" bomb. The United States is concerned that Upland is not taking this information seriously and that some plants are still vulnerable to attack.

Question: Has Upland increased security at all of its nuclear power plants?

Available Information:

- In an intercepted message, the Power Minister told the Prime Minister that additional fences and security cameras were installed at all nuclear power plants in June.
- A newspaper report said that security has been upgraded at all nuclear power plants.
- Imagery shows that new fences were installed at 3 nuclear power plants. Imagery of the other plant was not available.

Explained Solution:

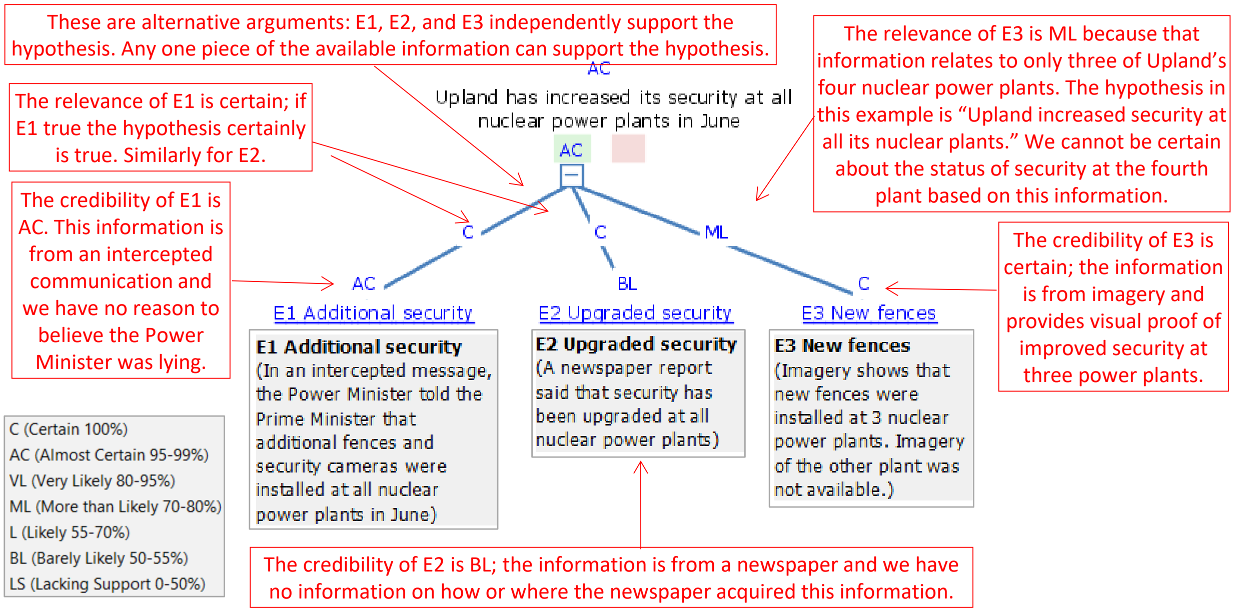
AC
Upland has increased its security at all nuclear power plants in June
AC
C
C
ML
AC
BL
C
E1 Additional security
E2 Upgraded security
E3 New fences

These are alternative arguments: E1, E2, and E3 independently support the hypothesis. Any one piece of the available information can support the hypothesis.

The relevance of E1 is certain; if E1 true the hypothesis certainly is true. Similarly for E2.

The credibility of E1 is AC. This information is from an intercepted communication and we have no reason to believe the Power Minister was lying.

The relevance of E3 is ML because that information relates to only three of Upland's four nuclear power plants. The hypothesis in this example is "Upland increased security at all its nuclear plants." We cannot be certain about the status of security at the fourth plant based on this information.

The credibility of E3 is certain; the information is from imagery and provides visual proof of improved security at three power plants.

E1 Additional security (In an intercepted message, the Power Minister told the Prime Minister that additional fences and security cameras were installed at all nuclear power plants in June)

E2 Upgraded security (A newspaper report said that security has been upgraded at all nuclear power plants)

E3 New fences (Imagery shows that new fences were installed at 3 nuclear power plants. Imagery of the other plant was not available.)

The credibility of E2 is BL; the information is from a newspaper and we have no information on how or where the newspaper acquired this information.

C (Certain 100%)
AC (Almost Certain 95-99%)
VL (Very Likely 80-95%)
ML (More than Likely 70-80%)
L (Likely 55-70%)
BL (Barely Likely 50-55%)
LS (Lacking Support 0-50%)

Eonomics Minister Fired

The Prime Minister has fired the Economics Minister.

Question: Why did the Prime Minister fire the Economics Minister?

Available Information:

- A reliable source who is a longtime aide to the Prime Minister said that the Prime Minister values complete loyalty above all else.
- The Economics Minister told the President that the Prime Minister was incompetent and lazy, according to a reliable source on the Prime Minister's staff.

Explained Solution:

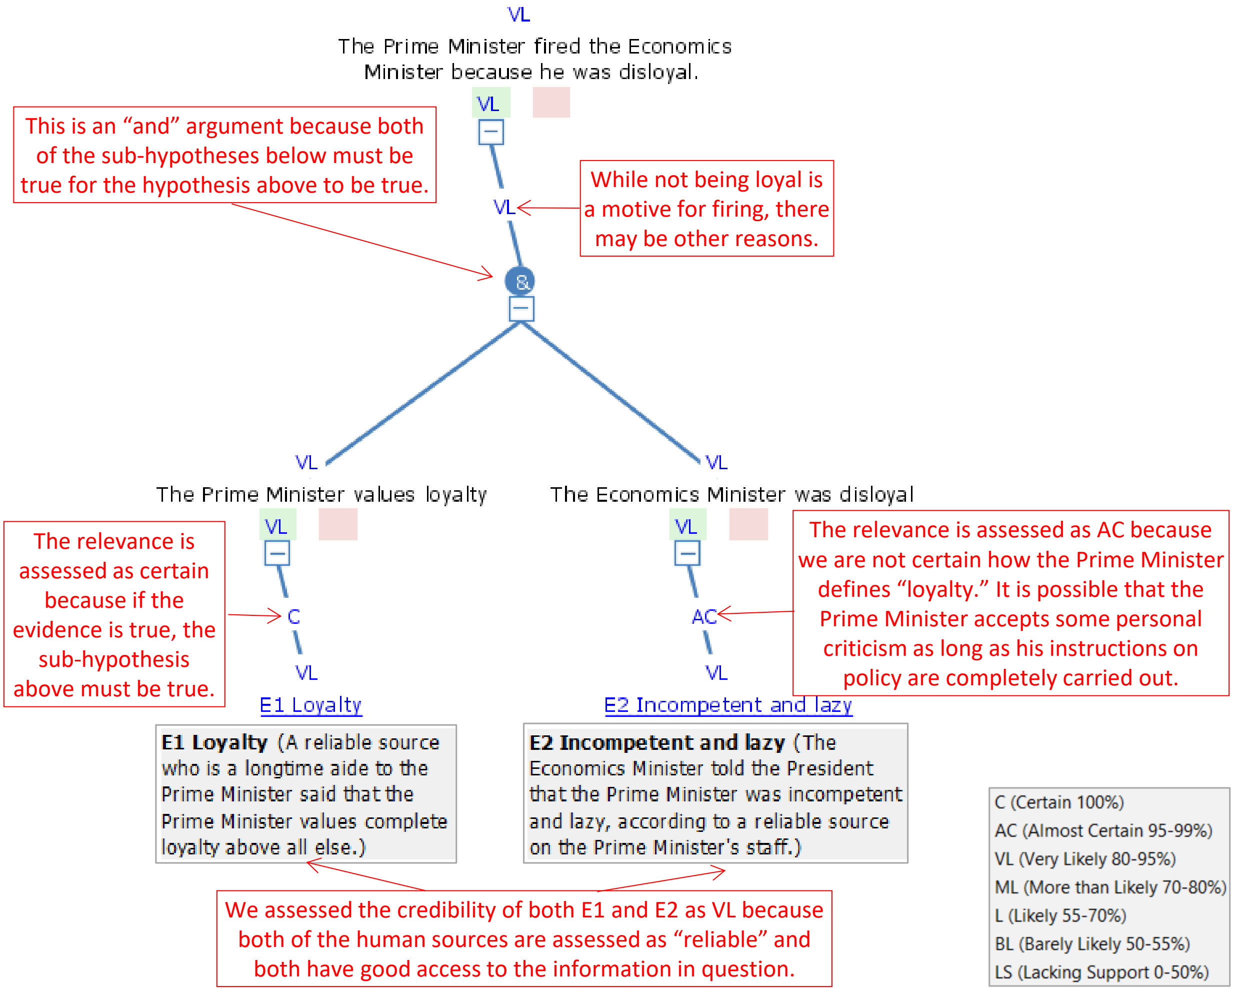
VL
The Prime Minister fired the Economics Minister because he was disloyal.
VL
VL
&
VL
The Prime Minister values loyalty
VL
C
VL
E1 Loyalty
VL
The Economics Minister was disloyal
VL
AC
VL
E2 Incompetent and lazy
C (Certain 100%)
AC (Almost Certain 95-99%)
VL (Very Likely 80-95%)
ML (More than Likely 70-80%)
L (Likely 55-70%)
BL (Barely Likely 50-55%)
LS (Lacking Support 0-50%)

This is an "and" argument because both of the sub-hypotheses below must be true for the hypothesis above to be true.

While not being loyal is a motive for firing, there may be other reasons.

The relevance is assessed as certain because if the evidence is true, the sub-hypothesis above must be true.

The relevance is assessed as AC because we are not certain how the Prime Minister defines "loyalty." It is possible that the Prime Minister accepts some personal criticism as long as his instructions on policy are completely carried out.

E1 Loyalty (A reliable source who is a longtime aide to the Prime Minister said that the Prime Minister values complete loyalty above all else.)

E2 Incompetent and lazy (The Economics Minister told the President that the Prime Minister was incompetent and lazy, according to a reliable source on the Prime Minister's staff.)

We assessed the credibility of both E1 and E2 as VL because both of the human sources are assessed as "reliable" and both have good access to the information in question.

Law School

John is graduating from college and his parents are wondering what John will do next.

Question: Will John apply to law school?

Available Information:

- John told his girlfriend he has the legal mind of a Supreme Court justice.
- John, according to numerous comments in his year book, is the only student at his college that wears a tie and suit to class.
- John told his brother that he wants to be making at least $150,000 a year by the time he is 30.

Explained Solution

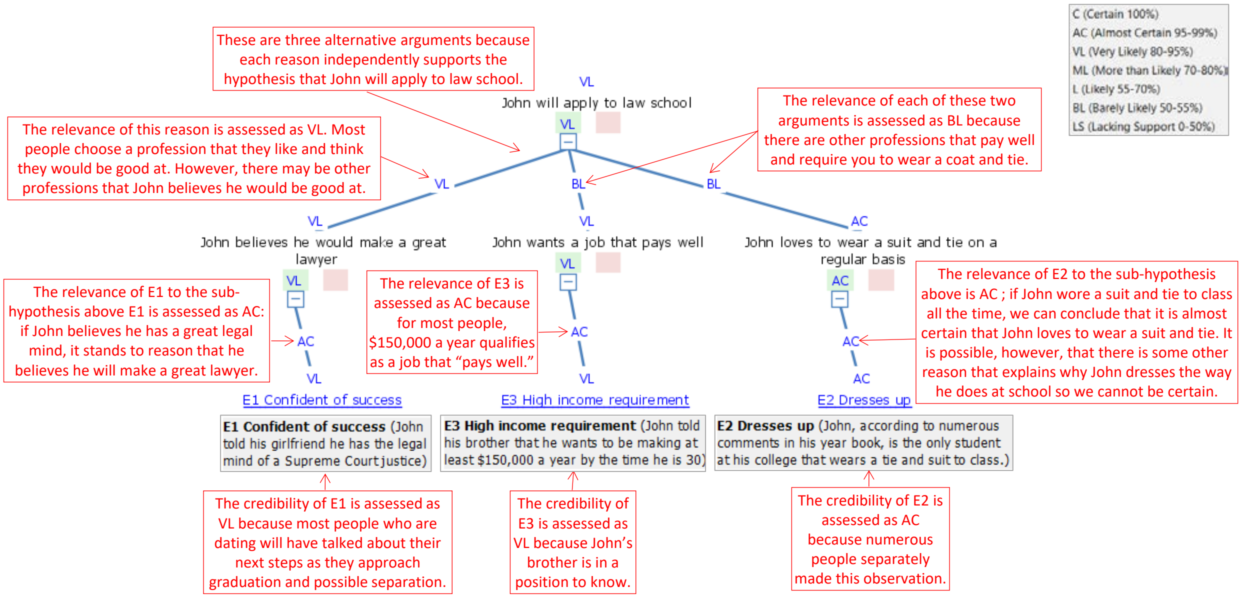
C (Certain 100%)
AC (Almost Certain 95-99%)
VL (Very Likely 80-95%)
ML (More than Likely 70-80%)
L (Likely 55-70%)
BL (Barely Likely 50-55%)
LS (Lacking Support 0-50%)
VL
John will apply to law school
VL
VL
BL
BL
VL
John believes he would make a great lawyer
VL
AC
VL
E1 Confident of success
VL
John wants a job that pays well
VL
AC
VL
E3 High income requirement
AC
John loves to wear a suit and tie on a regular basis
AC
AC
AC
E2 Dresses up

These are three alternative arguments because each reason independently supports the hypothesis that John will apply to law school.

The relevance of this reason is assessed as VL. Most people choose a profession that they like and think they would be good at. However, there may be other professions that John believes he would be good at.

The relevance of each of these two arguments is assessed as BL because there are other professions that pay well and require you to wear a coat and tie.

The relevance of E1 to the sub-hypothesis above E1 is assessed as AC: if John believes he has a great legal mind, it stands to reason that he believes he will make a great lawyer.

The relevance of E3 is assessed as AC because for most people, $150,000 a year qualifies as a job that "pays well."

The relevance of E2 to the sub-hypothesis above is AC ; if John wore a suit and tie to class all the time, we can conclude that it is almost certain that John loves to wear a suit and tie. It is possible, however, that there is some other reason that explains why John dresses the way he does at school so we cannot be certain.

E1 Confident of success (John told his girlfriend he has the legal mind of a Supreme Court justice)

E3 High income requirement (John told his brother that he wants to be making at least $150,000 a year by the time he is 30)

E2 Dresses up (John, according to numerous comments in his year book, is the only student at his college that wears a tie and suit to class.)

The credibility of E1 is assessed as VL because most people who are dating will have talked about their next steps as they approach graduation and possible separation.

The credibility of E3 is assessed as VL because John's brother is in a position to know.

The credibility of E2 is assessed as AC because numerous people separately made this observation.

Entrance Exam

Nick wants to go to law school. He will have to do well on his law school entrance exams because he barely graduated as an undergraduate with a 2.02 GPA.

Question: How will Nick do on the law school entrance exam?

Available Information:

- His IQ was measured at 145.
- Nick has been working at the car wash 60 hours a week for the past three months, according to his employer's records.
- He scored a 1475 on the SATS and 1490 on the GRE.
- Nick's dad said that Nick, who now lives at home, is partying all the time when he isn't working.

Explained Solution:

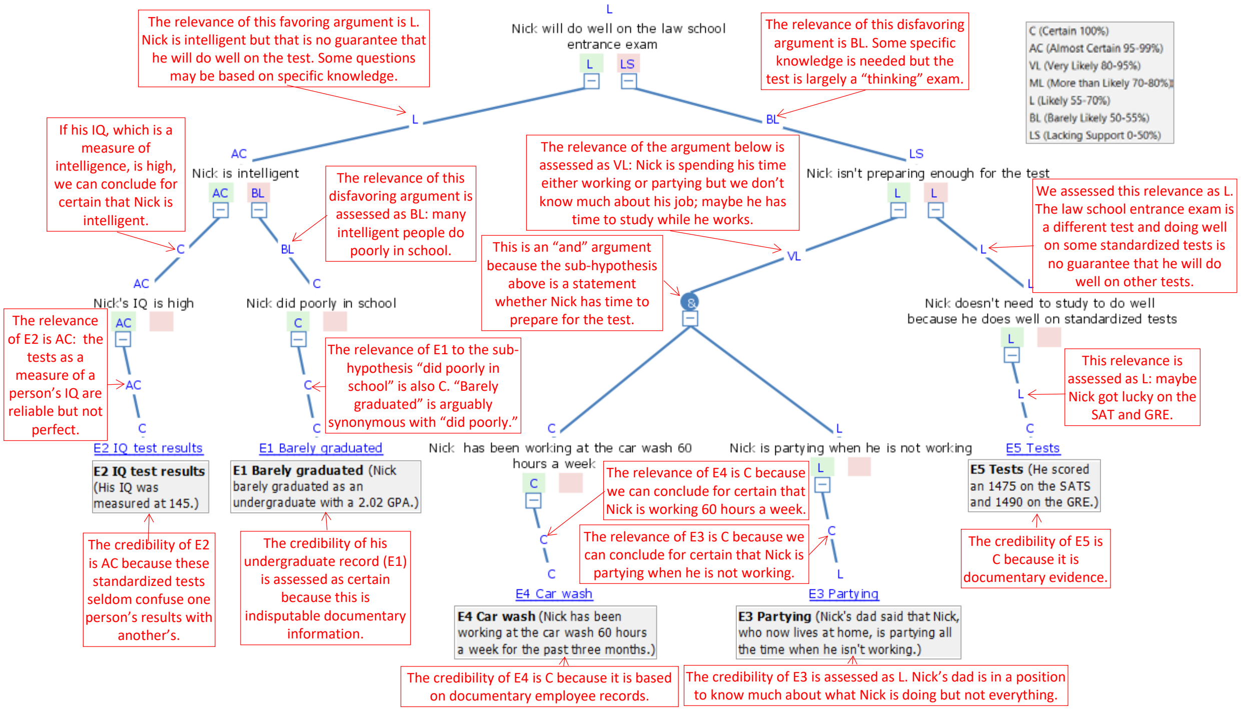
L
Nick will do well on the law school entrance exam
L
LS
C (Certain 100%)
AC (Almost Certain 95-99%)
VL (Very Likely 80-95%)
ML (More than Likely 70-80%)
L (Likely 55-70%)
BL (Barely Likely 50-55%)
LS (Lacking Support 0-50%)
L
BL
AC
Nick is intelligent
AC
BL
LS
Nick isn't preparing enough for the test
L
L
C
BL
VL
L
AC
Nick's IQ is high
C
Nick did poorly in school
C
L
Nick doesn't need to study to do well because he does well on standardized tests
AC
C
L
AC
C
L
C
C
C
L
C
E2 IQ test results
E1 Barely graduated
Nick has been working at the car wash 60 hours a week
Nick is partying when he is not working
E5 Tests
E2 IQ test results (His IQ was measured at 145.)
E1 Barely graduated (Nick barely graduated as an undergraduate with a 2.02 GPA.)
C
L
E5 Tests (He scored an 1475 on the SATS and 1490 on the GRE.)
C
C
C
L
E4 Car wash
E3 Partying
E4 Car wash (Nick has been working at the car wash 60 hours a week for the past three months.)
E3 Partying (Nick's dad said that Nick, who now lives at home, is partying all the time when he isn't working.)

The relevance of this favoring argument is L. Nick is intelligent but that is no guarantee that he will do well on the test. Some questions may be based on specific knowledge.

The relevance of this disfavoring argument is BL. Some specific knowledge is needed but the test is largely a "thinking" exam.

If his IQ, which is a measure of intelligence, is high, we can conclude for certain that Nick is intelligent.

The relevance of this disfavoring argument is assessed as BL: many intelligent people do poorly in school.

The relevance of the argument below is assessed as VL: Nick is spending his time either working or partying but we don't know much about his job; maybe he has time to study while he works.

We assessed this relevance as L. The law school entrance exam is a different test and doing well on some standardized tests is no guarantee that he will do well on other tests.

This is an "and" argument because the sub-hypothesis above is a statement whether Nick has time to prepare for the test.

The relevance of E2 is AC: the tests as a measure of a person's IQ are reliable but not perfect.

The relevance of E1 to the sub-hypothesis "did poorly in school" is also C. "Barely graduated" is arguably synonymous with "did poorly."

This relevance is assessed as L: maybe Nick got lucky on the SAT and GRE.

The relevance of E4 is C because we can conclude for certain that Nick is working 60 hours a week.

The relevance of E3 is C because we can conclude for certain that Nick is partying when he is not working.

The credibility of E2 is AC because these standardized tests seldom confuse one person's results with another's.

The credibility of his undergraduate record (E1) is assessed as certain because this is indisputable documentary information.

The credibility of E5 is C because it is documentary evidence.

The credibility of E4 is C because it is based on documentary employee records.

The credibility of E3 is assessed as L. Nick's dad is in a position to know much about what Nick is doing but not everything.

6.5 Anticipatory Intelligence with Cogent

Let us continue the anticipatory analysis started in Section 1.5, where we have already established that the cesium-137 canister is missing.

6.5.1 What Happened to the Cesium-137 Canister?

Then, if H_1 is true, Mavis infers H_2: The cesium-137 canister was stolen. So, now we have the next abductive step shown in Figure 72.

***Why** am I inferring $\boldsymbol{H_2}$ from $\boldsymbol{H_1}$?*

Things of value that are missing might have been stolen.

***Why** do I infer that the cesium-137 canister was stolen?*

My standpoint objective makes me naturally suspicious when radioactive materials go missing. Theft is a plausible explanation.

The records show that the XYZ Company has never lost any cesium-137 in the past.

There has been no denial of this report from the XYZ Company.

***What** other explanations other than theft am I considering?*

The cesium-137 canister was misplaced.

The cesium-137 canister is used in a project without being checked-out from the XYZ warehouse.

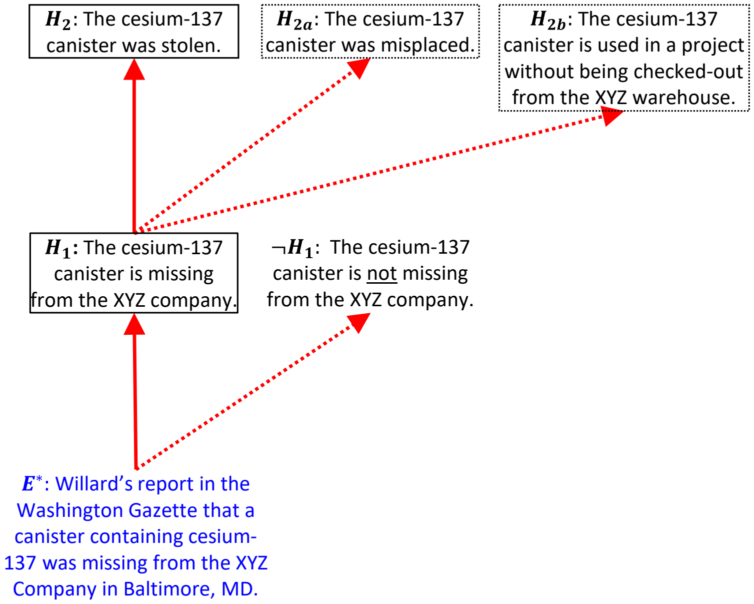
$\boldsymbol{H_2}$: The cesium-137 canister was stolen.
$\boldsymbol{H_{2a}}$: The cesium-137 canister was misplaced.
$\boldsymbol{H_{2b}}$: The cesium-137 canister is used in a project without being checked-out from the XYZ warehouse.
$\boldsymbol{H_1}$: The cesium-137 canister is missing from the XYZ company.
$\neg\boldsymbol{H_1}$: The cesium-137 canister is not missing from the XYZ company.
$\boldsymbol{E^*}$: Willard's report in the Washington Gazette that a canister containing cesium-137 was missing from the XYZ Company in Baltimore, MD.

Figure 72. Abducing that the cesium-137 canister was stolen.

Was the Cesium-137 Canister Stolen?

We have to put the hypothesis $\boldsymbol{H_2}$ to work to guide the collection of relevant evidence. One strategy is to search for indicators that the cesium-137 canister was stolen, such as those from Figure 73.

In searching for evidence that the hazardous material locker storing the cesium-137 canister was breached, we talked to a professional locksmith named Clyde who said that the lock had been forced, but that it was a clumsy job (see I3 Clyde inTable 20).

In searching for evidence of suspicious activity, we investigated the security camera of the warehouse and discovered a video segment showing a person loading a container into a U-Haul panel truck (see I4 Security Camera inTable 20).

In searching for evidence that the entrance of the STEMQ warehouse was breached, we investigated its security procedures and obtained the information labeled I5 Guard in Table 20 concerning the panel truck having Maryland license plate MDC-578 which was inside the STEMQ warehouse on the day before Willard's discovery of the missing cesium-137 canister. This led us to the identification of the renting company owing the truck, as well as the name and the address of the person who rented the truck (see I6 TRUXINC). Further investigation of the person and the truck revealed the information in I7 Silver Spring and I8 Investigative Record.

Table 20. Additional information on the missing of the cesium-137 canister.

I3 Clyde: We talked to a professional locksmith named Clyde, who said that the lock had been forced, but it was a clumsy job.

I4 Security Camera: The security camera of the STEMQ warehouse contains a video segment showing a person loading a container into a U-Haul panel truck.

I5 Guard: There is a security perimeter around the STEMQ warehouse and employee parking area having just one gate that is controlled by a guard. On the day before the missing canister was observed, the security guard, Sam, recorded that a panel truck having Maryland license plate MDC-578 was granted entry at 4:45PM just before the STEMQ closing hour at 5:00PM. The driver of this vehicle showed the guard a manifest containing items being delivered to the STEMQ warehouse. This manifest contained a list of packing materials allegedly ordered by the STEMQ Company. The vehicle was allowed to enter the parking area. At 8:30PM this same vehicle was allowed to exit the parking area. A different guard was on duty in the evenings and noticed that his records showed that this vehicle had been permitted entry and so he allowed the vehicle to exit the parking area.

I6 TRUXINC: Maryland DOT's record indicates that the panel truck carrying the license plate number MD-578 is registered in the name of a truck-rental company called TRUXINC, located in Silver Spring, MD. The manager of this agency showed records indicating that this truck was rented to a person who gave his name as Omer Riley, having as his listed address: 6176 Williams Ave. in Silver Spring. The truck was rented on the day before Willard's discovery of the missing cesium-137, and it was returned the day after he made the discovery.

I7 Silver Spring: Silver Spring city record according to which there is no residence at 6176 Williams Ave. in Silver Spring, MD.

I8 Investigative Record: An examination of the panel truck rented by Omer Riley, using a Geiger counter, revealed minute traces of cesium-137.

I9 Grace: Grace, the Vice President for Operations at STEMQ, tells us that no one at the STEMQ Company had checked out the canister for work on any project the STEMQ Company was working on at the time. She says that the STEMQ Company had other projects involving hazardous materials, but none that involved the use of cesium-137.

We have to identify the "dots" in the text from Table 20 which are fragments representing relevant evidence items for the leaf hypotheses inFigure 73. These dots are presented in Table 21.

The collected information from Figure 73 suggests the following scenario: A truck entered the company, the canister was stolen from the locker, the stolen canister was loaded into the truck, and the truck left with the canister. This leads to the development of the argumentation inFigure 74. The scenario is obviously a sufficient condition for the stolen hypothesis and therefore its relevance is certain.

Let's now consider each component action, starting with: A truck entered the company. We have E8 Guard Report favoring this hypothesis. Its *relevance* is certain because it asserts the event (that a truck entered the company). Its *credibility* is assessed as almost certain (95-99%) because this is the actual report and Sam is the security guard (with no reason to mislead).

Table 21. Dots from Table 20

E6 Clyde: Locksmith Clyde's report that the lock was forced.

E7 Security Camera: Video segment on the security camera of the STEMQ warehouse showing a person loading a container into a U-Haul panel truck.

E8 Guard Report: The record, made by Sam, security guard at the STEMQ Company, that a panel truck bearing Maryland license plate number MDC-578 was in the STEMQ parking area on the day before Willard's discovery of the missing cesium-137 canister.

E10 Security logs: Every week security personnel conduct an inventory on all canisters containing hazardous material. The week before the canister was discovered missing, the cesium canister was noted in company records as being in its assigned location.

E11 Security personnel: Security personnel at STEMQ are rigorously vetted for honesty and trustworthiness.

E9 MDDOT Record: Maryland DOT's record that the truck bearing license plate number MDC-578 is registered in the name of the TRUXINC Company in Silver Spring, MD.

E10 TRUXINC Record1: TRUXINC's record that the truck bearing MD license plate number MDC-578 was rented to a man who gave his name as Omer Riley on the day before Willard's discovery of the missing cesium-137 canister.

E11 TRUXINC Record2: TRUXINC's record that Omer Riley gave his address as 6176 Williams Ave.

E12 Silver Spring Record: Silver Spring city record according to which there is no residence at 6176 Williams Ave. in Silver Spring, MD.

E13 Investigative Record: Investigative record that traces of cesium-137 were found in the truck bearing license plate number MDC-578.

E14 Grace: Grace, the Vice President for Operations at STEMQ, tells us that no one at the STEMQ Company had checked out the canister for work on any project.

Let's now consider the hypothesis: The canister was stolen from locker. We have shown that the cesium-137 canister is missing from the warehouse. And we have evidence that the warehouse locker was forced open. This suggests an AND argument whose relevance we assess as being almost certain (95-99%).

One item of evidence favoring the hypothesis that the warehouse locker was forced open is Ralph's statement that the lock appears to have been forced open. We assess its *relevance* as likely (55-70%) because Ralph himself is not so sure, using the term "appears."

The second item of evidence favoring this hypothesis is Clyde's conclusion that the lock has been forced open. Its *relevance* is certain because the evidence asserts the event (that the locker was forced open). *Credibility* in this case is assessed as almost certain (95-99%) because Clyde is an independent expert and we know of no reason why he would not tell the truth.

Let's now consider the hypothesis: The stolen canister was loaded into the truck. We have E7 Security Camera favoring this hypothesis. We assess its *relevance* as likely (55-70%) because the loaded canister may not be the missing canister. We assess its *credibility* as *certain* because the video segment shows very clearly the loading of the canister into the truck, and we have no indication that it has been tampered with.

Finally, for the hypothesis "The truck left with the canister", we have no evidence. However, in the context of the other actions of the considered scenario, one may assume to be certain that this event happened. We discuss the assumptions in more details in the next section.

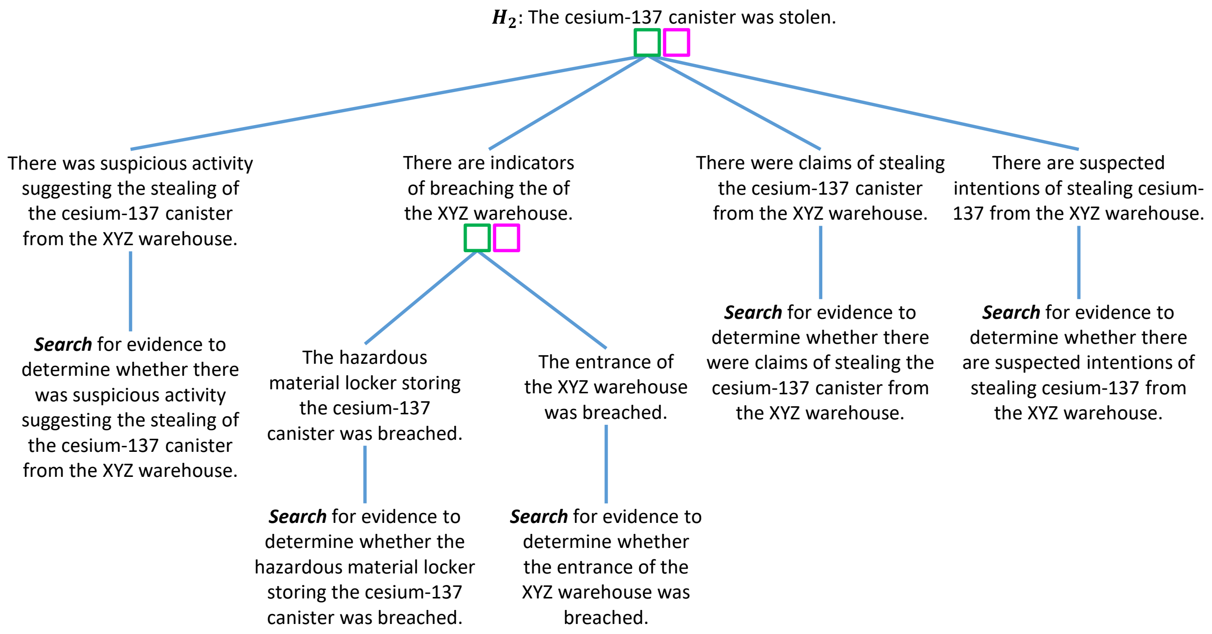
H_2: The cesium-137 canister was stolen.
There was suspicious activity suggesting the stealing of the cesium-137 canister from the XYZ warehouse.
There are indicators of breaching the of the XYZ warehouse.
There were claims of stealing the cesium-137 canister from the XYZ warehouse.
There are suspected intentions of stealing cesium-137 from the XYZ warehouse.
Search for evidence to determine whether there was suspicious activity suggesting the stealing of the cesium-137 canister from the XYZ warehouse.
The hazardous material locker storing the cesium-137 canister was breached.
The entrance of the XYZ warehouse was breached.
Search for evidence to determine whether there were claims of stealing the cesium-137 canister from the XYZ warehouse.
Search for evidence to determine whether there are suspected intentions of stealing cesium-137 from the XYZ warehouse.
Search for evidence to determine whether the hazardous material locker storing the cesium-137 canister was breached.
Search for evidence to determine whether the entrance of the XYZ warehouse was breached.

Figure 73. Evidence search guided by indicators of stealing.

We now have all the elements for a system, such as Cogent, to infer the probability of the top hypothesis. Based on the evidence, we have determined the probabilities of the simplest hypotheses. We have also estimated the probabilities of the two assumptions. The probability of the hypothesis "The canister was stolen from locker", very likely, is obtained as the smallest between the relevance of its argument and the probabilities of the sub-hypotheses. Similarly, the probability of the hypothesis "The stolen canister was loaded into the truck", likely, is obtained as the smallest between the relevance of its argument and the probabilities of the sub-hypotheses. Further up, the probability of the top hypothesis "The cesium-137 canister was stolen", likely, is the smallest between the relevance of its argument and the probability of the sub-hypotheses.

Use of Assumptions

Statements taken to be likely true, but without supporting evidence, are called ***assumptions****.* Analysts use assumptions to fill in gaps or address conflicting information, to support their analytic conclusions. Assumptions might be based on past behavior or comparable but unrelated situations, or “commonsense” reasoning about cause and effect. In general, for every assumption, one has to assess the probability that the assumption is true under the current situation, and to justify this probability.

When should we make an assumption? Consider the argument for E7 Security Camera fromFigure 74, reproduced in the left-hand side of Figure 75. Its *relevance* was assessed as only likely (55-70%) because the canister loaded into the truck may not be the stolen canister.

A relevance which is less than certain is an indication that there is a missing sub-hypothesis in the argument. We may make it explicit as shown in the right-hand side of Figure 75. First, above E7 Security Camera, we have introduced the corresponding event: A canister was loaded into the truck. If we now make the *assumption* that the loaded canister is the precisely the missing cesium-137 canister, then we create an argument with relevance certain. If the probability of this assumption is likely, then the probability that the stolen canister was loaded into the truck is also likely. The advantage, however, is that now we have made *explicit* what before was a *hidden assumption*. This increases the persuasiveness of our argumentation.

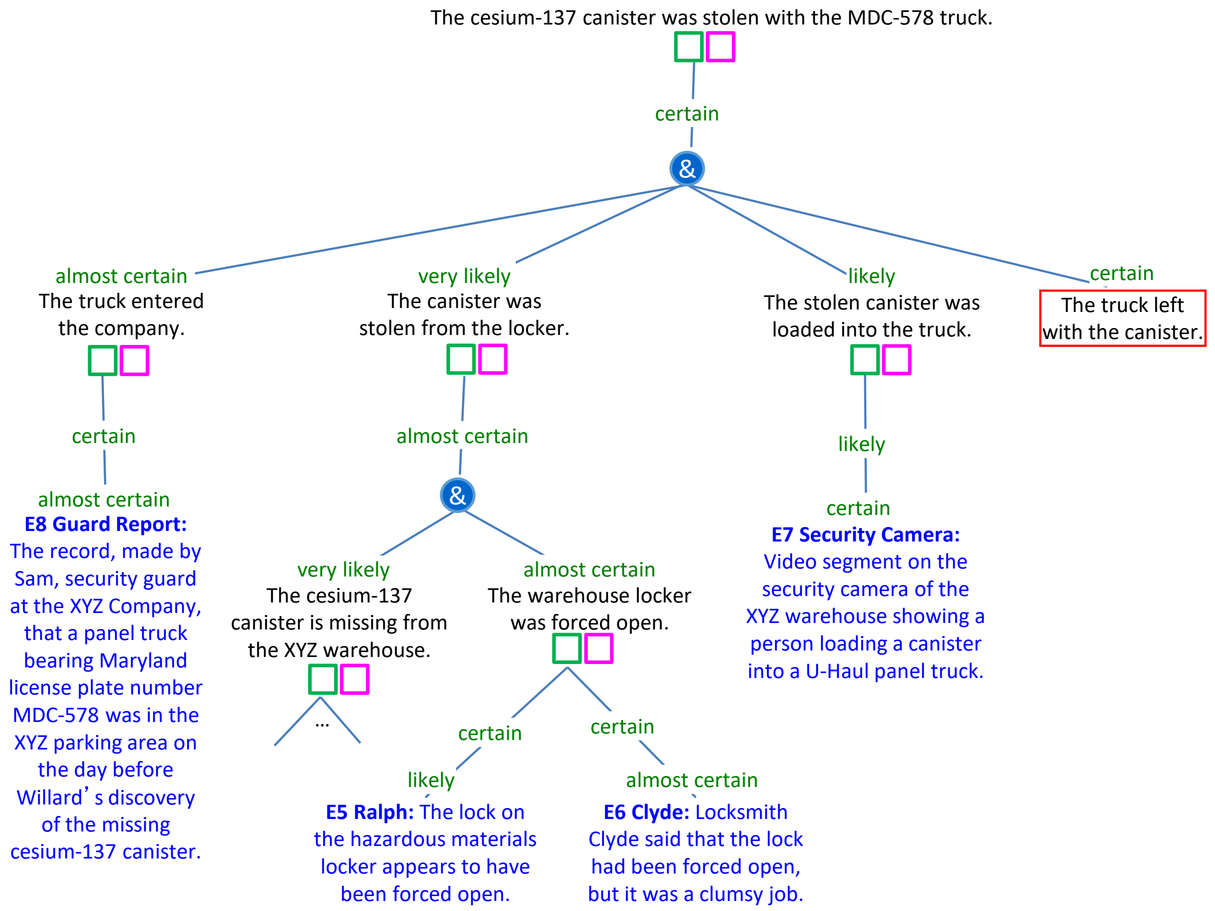
The cesium-137 canister was stolen with the MDC-578 truck.
certain
&
almost certain
The truck entered the company.
very likely
The canister was stolen from the locker.
likely
The stolen canister was loaded into the truck.
certain
The truck left with the canister.
certain
almost certain
E8 Guard Report: The record, made by Sam, security guard at the XYZ Company, that a panel truck bearing Maryland license plate number MDC-578 was in the XYZ parking area on the day before Willard's discovery of the missing cesium-137 canister.
almost certain
&
very likely
The cesium-137 canister is missing from the XYZ warehouse.
...
almost certain
The warehouse locker was forced open.
certain
certain
likely
E5 Ralph: The lock on the hazardous materials locker appears to have been forced open.
almost certain
E6 Clyde: Locksmith Clyde said that the lock had been forced open, but it was a clumsy job.
likely
certain
E7 Security Camera: Video segment on the security camera of the XYZ warehouse showing a person loading a canister into a U-Haul panel truck.

Figure 74. Initial argumentation for the hypothesis that the cesium canister was stolen.

As done above, it is recommended that, above each item of evidence, we insert the fact or the event asserted by that item, as was illustrated in this case with E7. This not only improves the clarity of the argumentation but also facilitates adding other items of evidence, either favoring or disfavoring that fact or event.

We may not always be able to apply this technique. There will be cases where there will be several possible alternative assumptions, as well as cases where we may not know what additional assumption to make. In such cases we will leave the initial argument, without adding any assumption, but should also use Cogent's capability to add an explanation to capture the reasoning supporting the relevance assessment.

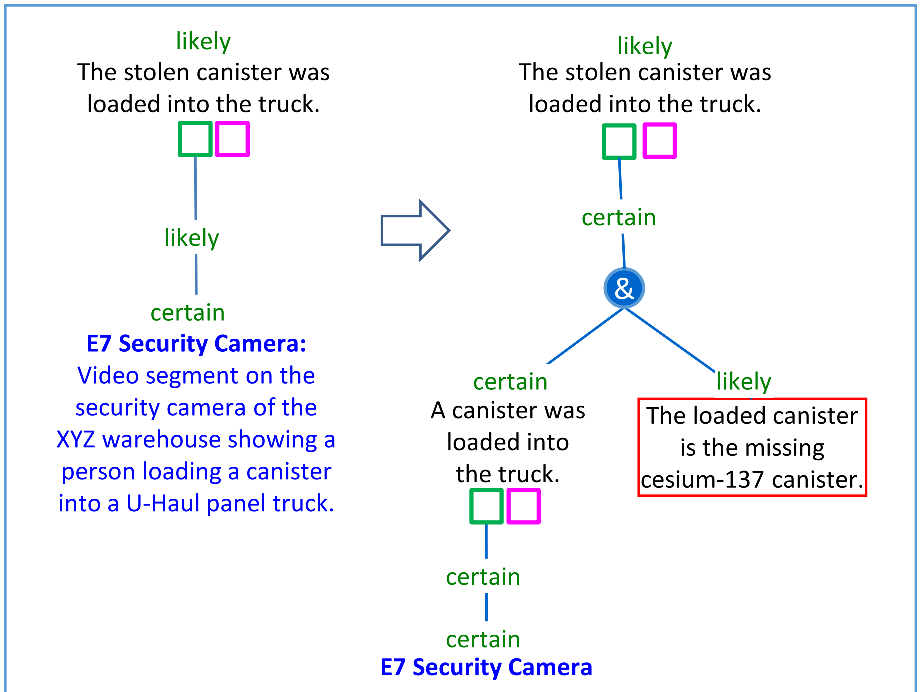
likely
The stolen canister was loaded into the truck.
likely
certain
E7 Security Camera:
Video segment on the security camera of the XYZ warehouse showing a person loading a canister into a U-Haul panel truck.
likely
The stolen canister was loaded into the truck.
certain
&
certain
A canister was loaded into the truck.
likely
The loaded canister is the missing cesium-137 canister.
certain
certain
E7 Security Camera

Figure 75. Transforming an argument to make an implicit assumption explicit.

Was the Cesiu-137 Canister Misplaced?

Let's now analyze the hypothesis H_{2a}: The cesium-137 canister was misplaced. What evidence is relevant to this hypothesis?

E10 Security logs: Every week security personnel conduct an inventory on all canisters containing hazardous material. The week before the canister was discovered missing, the cesium canister was noted in company records as being in its assigned location.

E11 Security personnel: Security personnel at STEMQ are rigorously vetted for honesty and trustworthiness.

Together they suggest the argumentation from Figure 76.

Trustworthy security personnel identified the cesium canister in its assigned location a week before the canister was discovered missing. This is disfavoring argument to the hypothesis that the cesium-137 canister was misplaced. Its inferential force is almost certain. Therefore, the evidence does not support the hypothesis that the cesium-137 canister was misplaced.

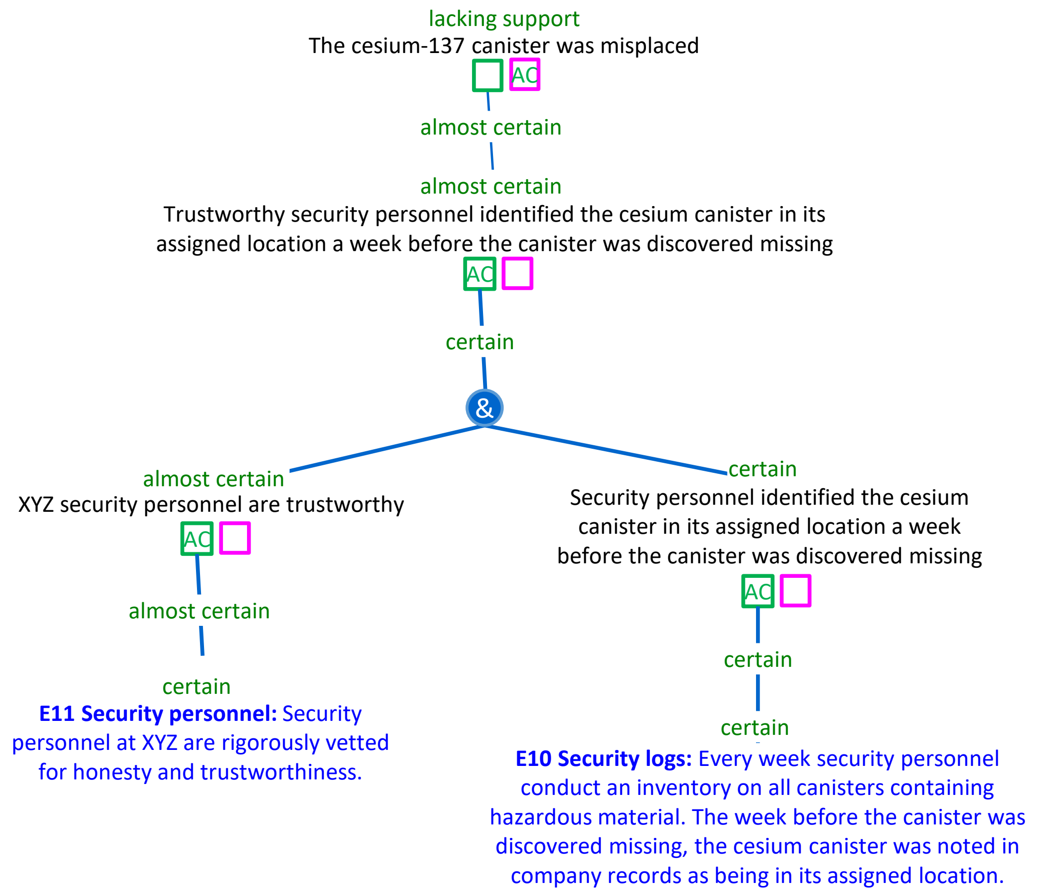
lacking support
The cesium-137 canister was misplaced
AC
almost certain
almost certain
Trustworthy security personnel identified the cesium canister in its assigned location a week before the canister was discovered missing
AC
certain
&
almost certain
XYZ security personnel are trustworthy
AC
almost certain
certain
E11 Security personnel: Security personnel at XYZ are rigorously vetted for honesty and trustworthiness.
certain
Security personnel identified the cesium canister in its assigned location a week before the canister was discovered missing
AC
certain
certain
E10 Security logs: Every week security personnel conduct an inventory on all canisters containing hazardous material. The week before the canister was discovered missing, the cesium canister was noted in company records as being in its assigned location.

Figure 76. Argumentation for the c misplaced esium hypothesis.

Is the Cesium-13 Canister Being Used in a Project?

We are finally analyzing the third hypothesis: H_{2b}: The cesium-137 canister is used in a project without being checked-out from the STEMQ warehouse. In searching for evidence we have contacted Grace, the Vice President for Operations at STEMQ. She tells us that no one at the STEMQ Company had checked out the canister for work on any project the STEMQ Company was working on at the time. She says that the STEMQ Company had other projects involving hazardous materials but none that involved the use of cesium-137 (see I9 Grace in Table 20).

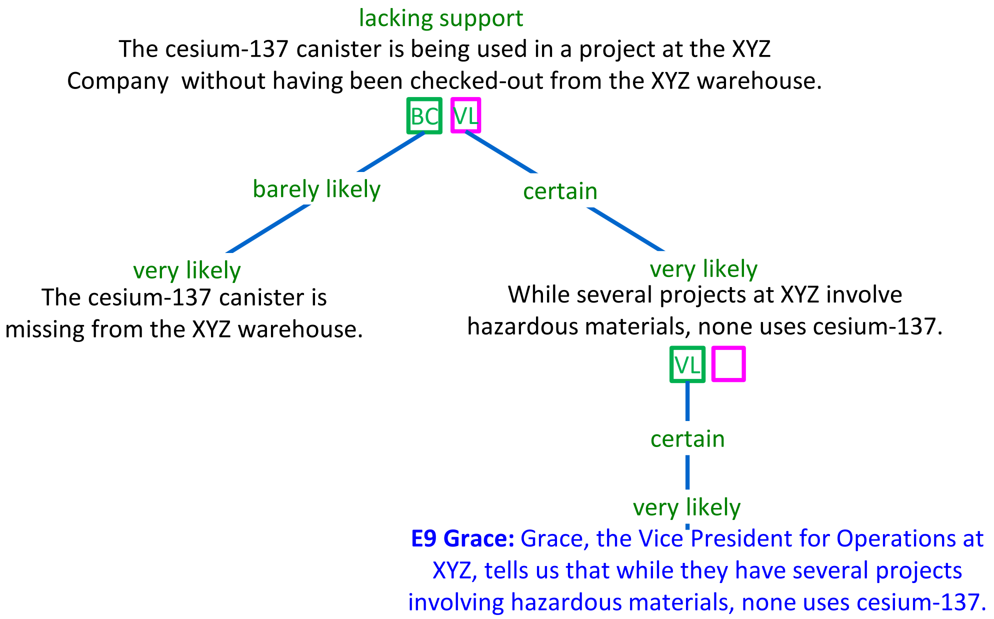
lacking support
The cesium-137 canister is being used in a project at the XYZ Company without having been checked-out from the XYZ warehouse.
BC
VL
barely likely
certain
very likely
The cesium-137 canister is missing from the XYZ warehouse.
very likely
While several projects at XYZ involve hazardous materials, none uses cesium-137.
VL
certain
very likely
E9 Grace: Grace, the Vice President for Operations at XYZ, tells us that while they have several projects involving hazardous materials, none uses cesium-137.

Figure 77. Argumnetation for the cesium used in a project hypothesis.

We have shown that it is very likely that the cesium canister is missing from the STEMQ warehouse. This is a barely likely indicator that the cesium canister is being used in a project at the STEMQ Company. So we have a favoring argument with inferential force barely likely, the smaller between the probability of the “missing” hypothesis and its relevance (seeFigure 77).

On the other hand, we have evidence from Grace, the Vice President for Operations at STEMQ, who tells us that while they have several projects involving hazardous materials, none uses cesium-137. This is disfavoring evidence for the top hypothesis. The relevance of this evidence item is certain because it directly contradicts the top hypothesis. We have no reason not to believe Grace, so let’s assess the credibility of this evidence item as very likely (80-95%). Thus we also have a disfavoring argument with inferential force very likely, the smaller between the relevance and the credibility of the evidence item E9 Grace.

The inferential force of the favoring argument is barely likely, but the inferential force of the disfavoring argument is much higher: very likely. Therefore, the evidence does not support the conclusion that the hypothesis is true, but it supports the conclusion that it is not true.

Therefore, the most likely hypothesis is $\boldsymbol{H_2}$: The cesium-137 canister was stolen.

6.5.2 The Cesium-137 Canister was Stolen by Omar al-Massari

Having been established that the available evidence favors the hypothesis that the cesium-137 canister was stolen from the STEMQ warehouse with the MDC-578 truck, the next step is to identify who has actually stolen it. A natural suspect is Omer Riley who has rented the MDC-578 truck. As indicated by I10 Santa in Table 22, our asset, Santa, has given us significant leads for getting additional information about our suspect. This has led to our obtaining of the items of information I11 Test and I12 Walsh. The information from Table 22 and the corresponding dots in Table 23 suggest the following scenario: Omar al-Massari rented the MDC-578 truck, giving his alias, Omer Riley, and a false address, and then used it to steal the cesium-137 canister, which in turn caused it to become contaminated because cesium-137 is a radioactive material. This leads to the development of the hypothesis analysis tree fromFigure 79.

Table 22. Information on the presumed stealing of the cesium-137 canister.

I10 Santa: An asset code-named "Santa" tells us that the name Omer Riley is one of the aliases used by a person named Omar al-Massari, who came to the USA in 2000, apparently from Saudi Arabia, on an extended work permit. Omar al-Massari is a physicist employed for the past two years by the Ultratech company in Silver Spring. He lives with two other males at 403 Winston Road in Silver Spring. Santa also tells us that Omar al-Massari (alias Omer Riley) is intimately associated with an unnamed jihadist organization in the Washington, DC area.

I11 Test: Omar al-Massari (alias Omer Riley) was apprehended at his place of work at the Ultratech company in Silver Spring. During his questioning, he was given a test called "whole body counting" with a Geiger counter that can detect the gamma radiation emitted by cesium-137. This test indicated the presence of traces of cesium-137 on his skin and hair.

I12 Walsh: The president of Ultratech company, a Mr. John Walsh, reported that Omar al-Massari's (alias Omer Riley) work at Ultratech does not involve his handling any radioactive substances.

Table 23. Dots from Table 22.

E15 Santa Alias: Santa's testimony that Omer Riley is an alias used by Omar al-Massari.

E16 Santa Work: Santa's testimony that Omar al-Massari is a physicist employed for the past two years by the Ultratech company in Silver Spring.

E17 Santa Adr: Santa's testimony that Omer Riley lives with two other males at 403 Winston Road in Silver Spring.

E18 Santa Ter Org: Santa's testimony that Omar al-Massari (alias Omer Riley) is intimately associated with an unnamed jihadist organization in the Washington, DC area.

E19 Omar Test: "Whole body counting" test result on Omar al-Massari (alias Omer Riley) with a Geiger counter indicating traces of cesium-137 on his skin and hair.

E20 Walsh: Walsh's testimony that Omar al-Massari's work does not involve handling of radioactive substances.

6.5.3 Who is Omar al-Massari?

We have established that the cesium-137 canister was stolen by Omar al-Massari. Now the question is: Who is Omar al-Massari? Is he someone working for a competitor, someone hoping to sell this valuable material, or someone having terrorist connections? Figure 78shows the indicators that guide us in collecting evidence to prove that Omar al-Massari has ties to terrorist organizations. As a result, we collect the information in Table 24.

6.5.4 Extracting Evidence from Information

Up to now, most of the case studies included both the items of information collected and the items of evidence extracted from them. In this case study you will have to extract the items of evidence from the collected information yourself. As discussed before, this involves the necessity for parsing incoming information to see what evidential dots or trifles this information reveals. Testimonial information or descriptions of tangible items might contain very many details, dots, or trifles. Some of the details might be interesting and relevant evidence, and others not. What we always have to do is to parse the information to extract the information that we believe is relevant in the inference task at hand.

In so many instances we have seen persons taking a lump of information containing many details, some interesting and some not, and treating it as a single item of potential evidence. There are two problems here. The first is that the relevant individual details in this lump might bear on different inferential issues; they will rarely all bear on the same issue. The second is that the irrelevant details only act to confuse the inferential bearings of the relevant details. Now, the problem is that determining what the relevant and irrelevant details are is a subjective matter. We might not all agree that a particular detail is relevant or irrelevant.

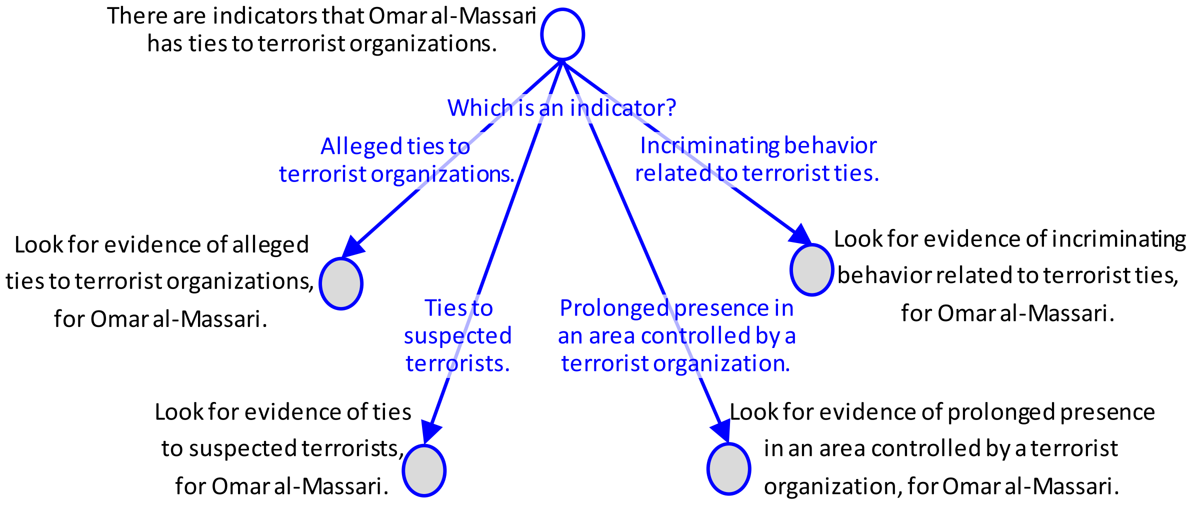
There are indicators that Omar al-Massari has ties to terrorist organizations.
Which is an indicator?
Alleged ties to terrorist organizations.
Incriminating behavior related to terrorist ties.
Look for evidence of alleged ties to terrorist organizations, for Omar al-Massari.
Look for evidence of incriminating behavior related to terrorist ties, for Omar al-Massari.
Ties to suspected terrorists.
Prolonged presence in an area controlled by a terrorist organization.
Look for evidence of ties to suspected terrorists, for Omar al-Massari.
Look for evidence of prolonged presence in an area controlled by a terrorist organization, for Omar al-Massari.

Figure 78. Evidence collection guided by indicators of terrorist ties.

Here comes an example involving I14 Clark from Table 24.

Table 24. Information on the presumed terrorist ties of Omar al-Massari.

I13 Laptop: In further investigation of Omar al-Massari, we tell him that we wish to see his laptop computer. We are, of course, interested in what it might reveal about the terrorists he may be associating with. He refuses to tell us where it is. We referred to this as the non-production of evidence.

I14 Clark: As we have learned, Omar al-Massari lives with two other males at 403 Winston Road in Silver Spring, Richard Clark and Fahd al-Quso. We were only able to interview Richard Clark. Clark is an American citizen of Anglo-Saxon origin who was born in 1973 in Lanham, MD, and is the owner of the residence at 403 Winston Road in Silver Spring. He has lived there since 2005, when he purchased the house. He says he had trouble making his house payments and was forced to take in renters who could contribute to his house payment. Clark went on to say that he rented rooms to al-Massari and al-Quso (who came together to look at the rooms) because they were professionals who could pay their rents on time, which Clark says they always did. We told Clark about al-Massari's detention and asked for further information about him and al-Quso. Clark responded that Fahd al-Quso had suddenly moved out two days ago, and that he had not seen al-Massari for the past two days. However, Clark says that al-Massari had driven up two days ago in a U-Haul panel truck, but had only stayed for a minute before he left. Clark says he wondered why al-Massari was driving around in a U-Haul truck. Finally, Clark says that his two renters kept to themselves most of the time they were home and that he never looked inside their rooms (which would have been an invasion of their privacy), but that he frequently overheard their conversations concerning a place called Allied Import, that Clark guessed was a business of some kind.

I15 Quso: The other housemate (besides Richard Clark) of Omar al-Massari is Fahd al-Quso, a Yemeni who, like Omar al-Massari, has been here on an extended work permit. Fahd al-Quso is also a physicist who has been employed for the past three years by a company called Physicom in Laurel, MD. We have not been able to interview him. Just yesterday, Fahd al-Quso boarded Emirates Flight #207 bound for Dubai. He purchased a one-way ticket using a credit card. He was not on any no-fly lists. We have just learned from a trusted source that Fahd al-Quso has been detained by United Arab Emirate authorities for questioning about his association with terrorist incidents in the UAE.

I16 FBI: Allied Import is a business at 2121 M Street East in Washington, DC, that deals with a variety of items from various places in the Middle East. FBI contacts tell us that Allied Import has been under surveillance for several months in connection with possible narcotics trafficking. We asked the FBI for any surveillance records they might have about trucks entering and leaving Allied Import. We were shown a surveillance video of a single man arriving two days ago in a U-Haul panel truck with MD license number MDC-578. The driver was positively identified as Omar al-Massari. The video shows al-Massari handing off a single canister-like container to a man. So, Allied Import may be in the terrorist business in addition to being in the narcotics business, since we know from experience that these two activities often go hand in hand. The man in the video accepting the container from al-Massari was identified by the FBI as a Maryland resident named Kenny Derwish. Derwish lists his residence as 113 4th St. in Oxon Hill, MD. Derwish has been with Allied Import for four years.

I17 Yasmin: A source code-named "Yasmin" tells us that she knew a man in Saudi Arabia named Omar al-Massari. Yasmin says she is "quite sure" that Omar spent two years "somewhere" in Afghanistan "sometime" in the years 1998-2000. Yasmin also tells us that she once met Omar al-Massari at a large gathering last August in Bethesda, MD, held (allegedly) to support charities in the Middle East. She said that Omar al-Massari attended this gathering with a person he identified as his roommate and that they were both physicists. Yasmin says that funds collected at this gathering were never intended to be used for charitable purposes, but to support terrorist activities both here and around the globe. Additionally, Yasmin tells us that Kenny Derwish is an alias used by Saeed al-Nami. She says that Saeed was associated with a now-disbanded Islamic Jihad cell in Herndon, VA. She further tells us that Saeed is a now a principal member of an active terrorist cell in Washington, DC, called "Jihad Bis Sayf" (Striving Through the Sword).

Table 25. New information related to Saeed al-Nami.

I18 Miller: A loading dock worker at Allied Import named Rocky Miller says that Derwish put the object he received from the U-Haul guy in the trunk of his car.

I19 Garcia: To learn more about al-Nami (alias Derwish) we interview the management of the Allied Import Company. The first thing we are told by Jose Garcia, a vice president of Allied Import, was that they only knew Kenny Derwish by this name; Garcia says he was very surprised to learn that this name was an alias. Garcia also said that Derwish had worked for Allied Import for five years and is an expert in the evaluation of firearms and explosives that Allied Import purchases from foreign suppliers. Garcia said that Derwish knew more about these items than anyone he had ever known. We asked Garcia what kind of explosives they import. He said that they import only plastic explosives like Semtex, RDX, and C-4 that are very stable and can be shipped safely by ground and sea transport. Many American demolition companies use these explosives but can get them cheaper from China and some European companies. We asked Garcia if any of their imported explosives had gone missing. He said that this has rarely happened, but about two weeks ago, a small amount of RDX, about two pounds, went missing from a storage facility to which al-Nami (Derwish) had access. We then asked Garcia if we could talk with Derwish, Garcia said that Derwish had gone on vacation two days ago.

I20 Yasmin: We contacted our source, Yasmin, again. All she has told us so far was that Saeed al-Nami used an alias (Kenny Derwish) and that he was associated with the militant jihadist group Jihad Bis Sayf. We now asked her for more information about al-Nami. She said that the name Derwish was a Yemeni name but al-Nami was a Saudi name. Yasmin said that he took this alias, Derwish, for two reasons. First, because it would sound a much less Muslim name to Americans, and second, because he admired the Yemenis and had spent a year in 2003 in an al Qaeda training camp in Yemen, where he received training in the use and construction of explosive devices. We asked Yasmin if any of this training might have included the assembly of dirty bombs. She says she would not rule this out because al Qaeda wished to have these weapons for many years, and had assembled some stocks of radiological materials.

In this case Clark tells us a variety of things, some of which seem potentially relevant and others not. Looking carefully at this testimony from Clark, we can first identify details that seem irrelevant as far as al-Massari's terrorist activities are concerned. Here are some of them:

- Clark had trouble making payments on the house he purchased in 2005. So he had to take in renters.
- The renters, al-Massari and al-Quso, always paid their rent on time.
- The renters, al-Massari and al-Quso, kept to themselves most of the time when they were home.

But here are some apparently relevant details, dot, or trifles that bear on different matters:

Clark is of Anglo-Saxon origin and was born in 1973 in Lanham, MD. These details are potentially relevant concerning Clark's credibility. That he is of Anglo-Saxon origin says that he is probably a kufr, a Muslim word for an infidel. That he is 37 years old might arise in assessments of his competence.

- Al-Massari and al-Quso roomed together. This bears on al-Massari's association with a suspected terrorist.
- Al-Massari and al-Quso frequently discussed the Allied Import Co. This bears on where al-Massari might have taken the cesium-137.
- Two days ago al-Massari drove up to their residence in a U-Haul truck. This also bears on where al Massari might have taken the cesium-137.
- Al-Massari only stayed for a minute two days ago when he drove up in the U-Haul truck. This may be relevant on when al-Massari took the cesium to Allied Import.
- Al-Quso suddenly moved out two days ago. This bears on al Quso's behavior as a possible terrorist.

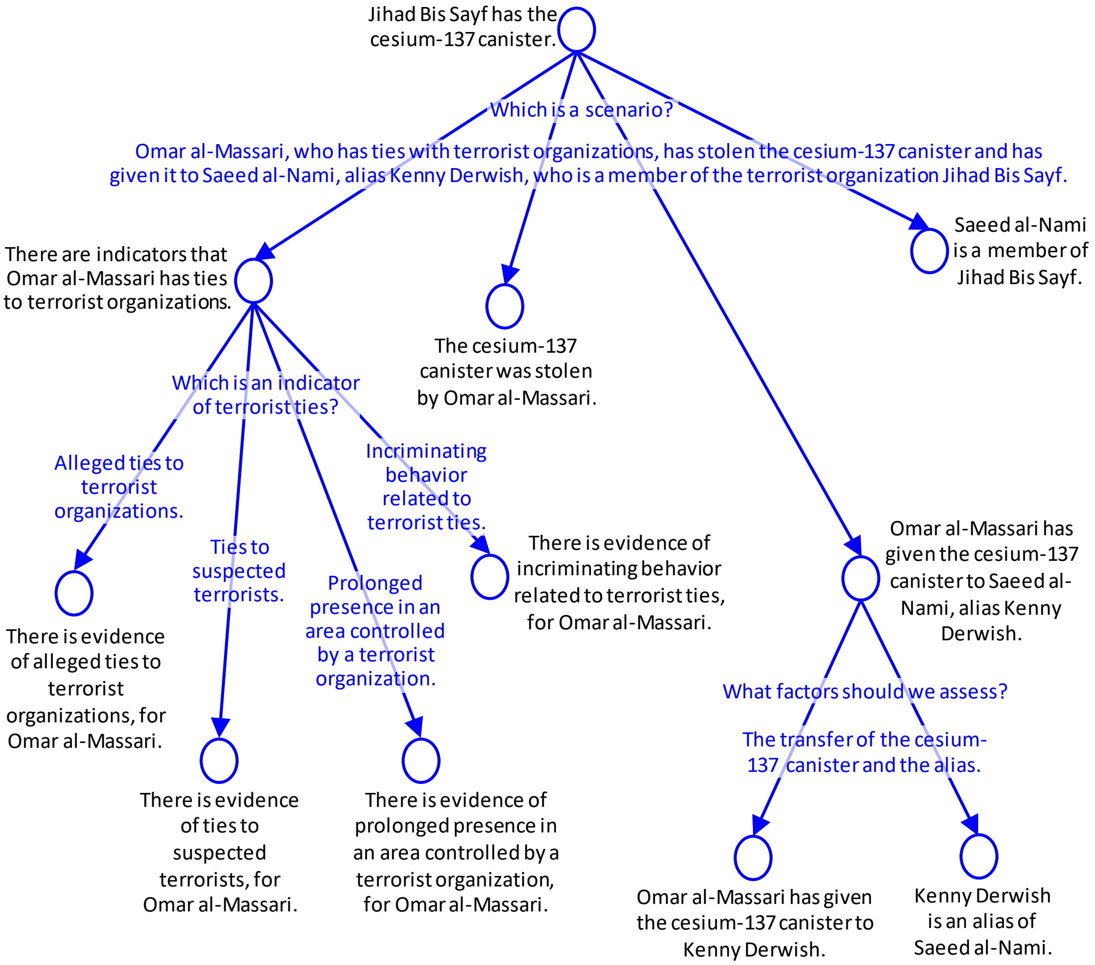
Jihad Bis Sayf has the cesium-137 canister.
Which is a scenario?
Omar al-Massari, who has ties with terrorist organizations, has stolen the cesium-137 canister and has given it to Saeed al-Nami, alias Kenny Derwish, who is a member of the terrorist organization Jihad Bis Sayf.
There are indicators that Omar al-Massari has ties to terrorist organizations.
The cesium-137 canister was stolen by Omar al-Massari.
Saeed al-Nami is a member of Jihad Bis Sayf.
Which is an indicator of terrorist ties?
Alleged ties to terrorist organizations.
Ties to suspected terrorists.
Prolonged presence in an area controlled by a terrorist organization.
Incriminating behavior related to terrorist ties.
There is evidence of alleged ties to terrorist organizations, for Omar al-Massari.
There is evidence of ties to suspected terrorists, for Omar al-Massari.
There is evidence of prolonged presence in an area controlled by a terrorist organization, for Omar al-Massari.
There is evidence of incriminating behavior related to terrorist ties, for Omar al-Massari.
Omar al-Massari has given the cesium-137 canister to Saeed al-Nami, alias Kenny Derwish.
What factors should we assess?
The transfer of the cesium-137 canister and the alias.
Omar al-Massari has given the cesium-137 canister to Kenny Derwish.
Kenny Derwish is an alias of Saeed al-Nami.

Figure 79. Analysis of the hypothesis that a terrorist organization has the stolen cesium-137.

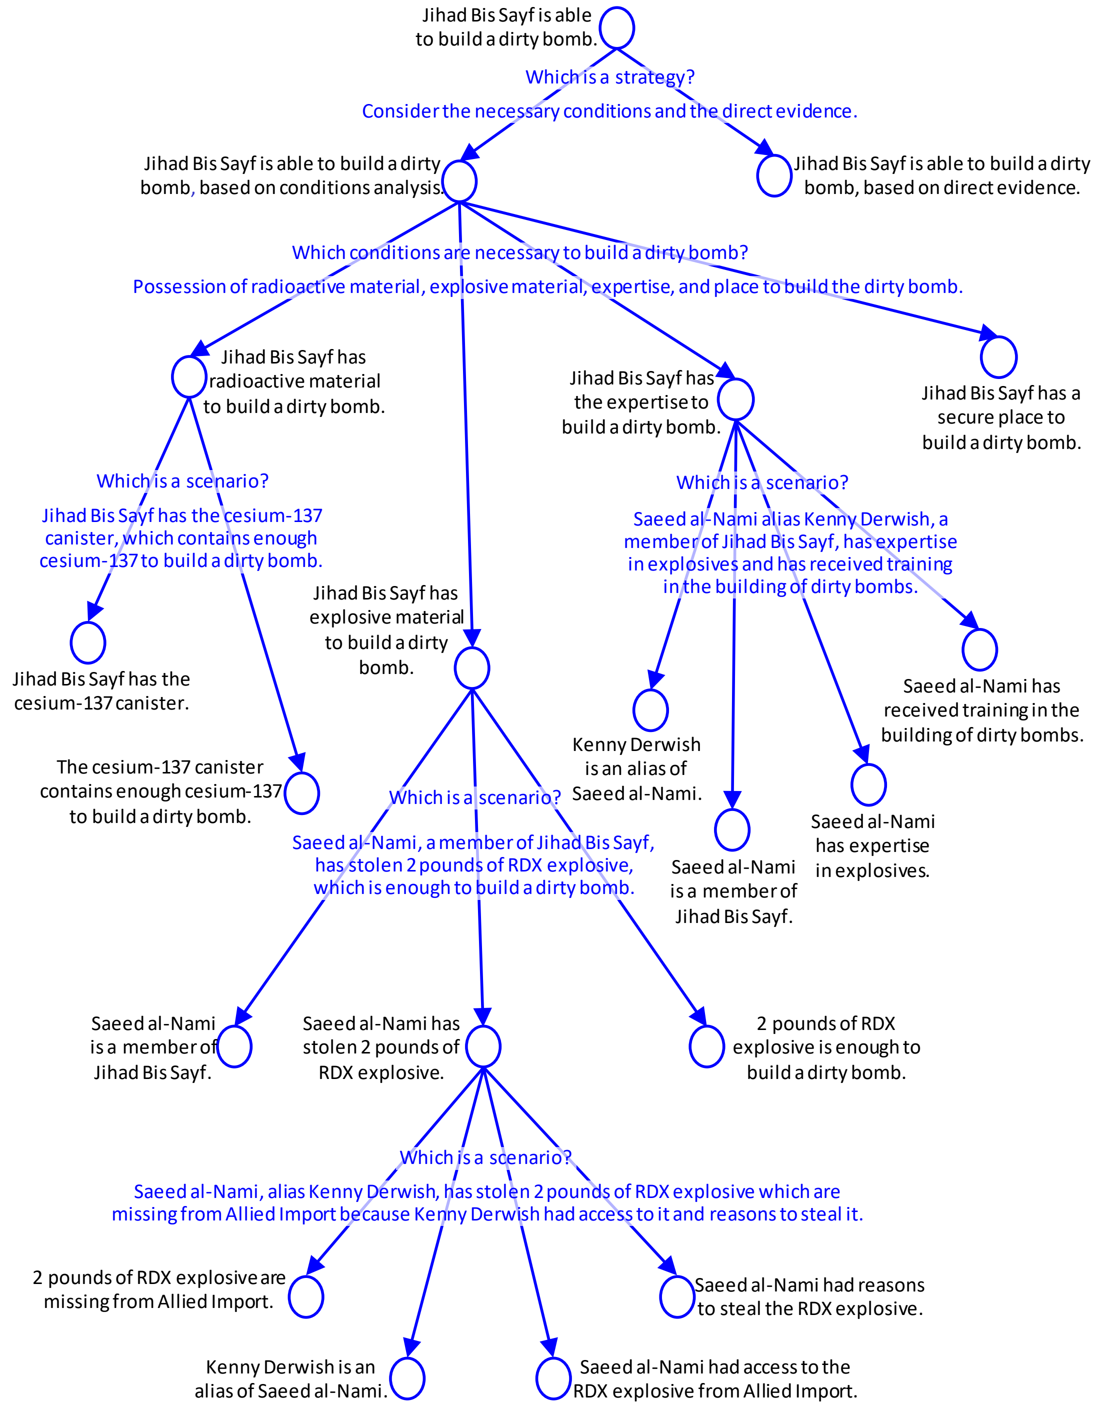
Jihad Bis Sayf is able to build a dirty bomb.
Which is a strategy?
Consider the necessary conditions and the direct evidence.
Jihad Bis Sayf is able to build a dirty bomb, based on conditions analysis.
Jihad Bis Sayf is able to build a dirty bomb, based on direct evidence.
Which conditions are necessary to build a dirty bomb?
Possession of radioactive material, explosive material, expertise, and place to build the dirty bomb.
Jihad Bis Sayf has radioactive material to build a dirty bomb.
Jihad Bis Sayf has the expertise to build a dirty bomb.
Jihad Bis Sayf has a secure place to build a dirty bomb.
Which is a scenario?
Jihad Bis Sayf has the cesium-137 canister, which contains enough cesium-137 to build a dirty bomb.
Which is a scenario?
Saeed al-Nami alias Kenny Derwish, a member of Jihad Bis Sayf, has expertise in explosives and has received training in the building of dirty bombs.
Jihad Bis Sayf has explosive material to build a dirty bomb.
Jihad Bis Sayf has the cesium-137 canister.
Saeed al-Nami has received training in the building of dirty bombs.
Kenny Derwish is an alias of Saeed al-Nami.
The cesium-137 canister contains enough cesium-137 to build a dirty bomb.
Which is a scenario?
Saeed al-Nami has expertise in explosives.
Saeed al-Nami, a member of Jihad Bis Sayf, has stolen 2 pounds of RDX explosive, which is enough to build a dirty bomb.
Saeed al-Nami is a member of Jihad Bis Sayf.
Saeed al-Nami is a member of Jihad Bis Sayf.
Saeed al-Nami has stolen 2 pounds of RDX explosive.
2 pounds of RDX explosive is enough to build a dirty bomb.
Which is a scenario?
Saeed al-Nami, alias Kenny Derwish, has stolen 2 pounds of RDX explosive which are missing from Allied Import because Kenny Derwish had access to it and reasons to steal it.
2 pounds of RDX explosive are missing from Allied Import.
Saeed al-Nami had reasons to steal the RDX explosive.
Kenny Derwish is an alias of Saeed al-Nami.
Saeed al-Nami had access to the RDX explosive from Allied Import.

Figure 80. Analysis of the hypothesis that Jihad Bis Sayf is able to build a dirty bomb.

This analysis is an example of the necessity for parsing lumps of information to identify what different specific details, dot, or trifles are relevant, and how they are relevant on different issues in an analysis.

6.5.5 The Cesium-137 Canister was Stolen by Someone Aassociated with Terrorists

The information from Table 24 suggests a scenario where Omar al-Massari, who has ties with terrorist organizations, has stolen the cesium-137 canister and has given it to Saeed al-Nami, alias Kenny Derwish, who is a member of Jihad Bis Sayf. Consequently, we develop the analysis tree in Figure 79 to assess the hypothesis that the terrorist organization Jihad Bis Sayf has the cesium-137 canister. First you will define evidence based on the available information. Then you will associate it with elementary hypotheses and evaluate them. As a result, you will assess the probability of the top-level hypothesis. This analysis will be used in the next case study.

Will a Bomb Be Set Off in Washington DC?

We have concluded that the terrorist organization Jihad Bis Sayf, though its member Saeed al-Nami, has taken possession of the cesium-137 canister. The next hypothesis to evaluate is whether this organization is able to construct a dirty bomb. We direct our intelligence collection efforts on Saeed al-Nami and obtain the items of information from Table 25, which will be used to assess this hypothesis, as indicated in

Finally, we investigate the top-level hypothesis, that Jihad Bis Sayf will set-off a dirty bomb in the Washington, DC area. This reduces to assessing its reasons, desires, and capabilities, as shown in Figure 81. The case study in the next section allows you to assess this top-level hypothesis.

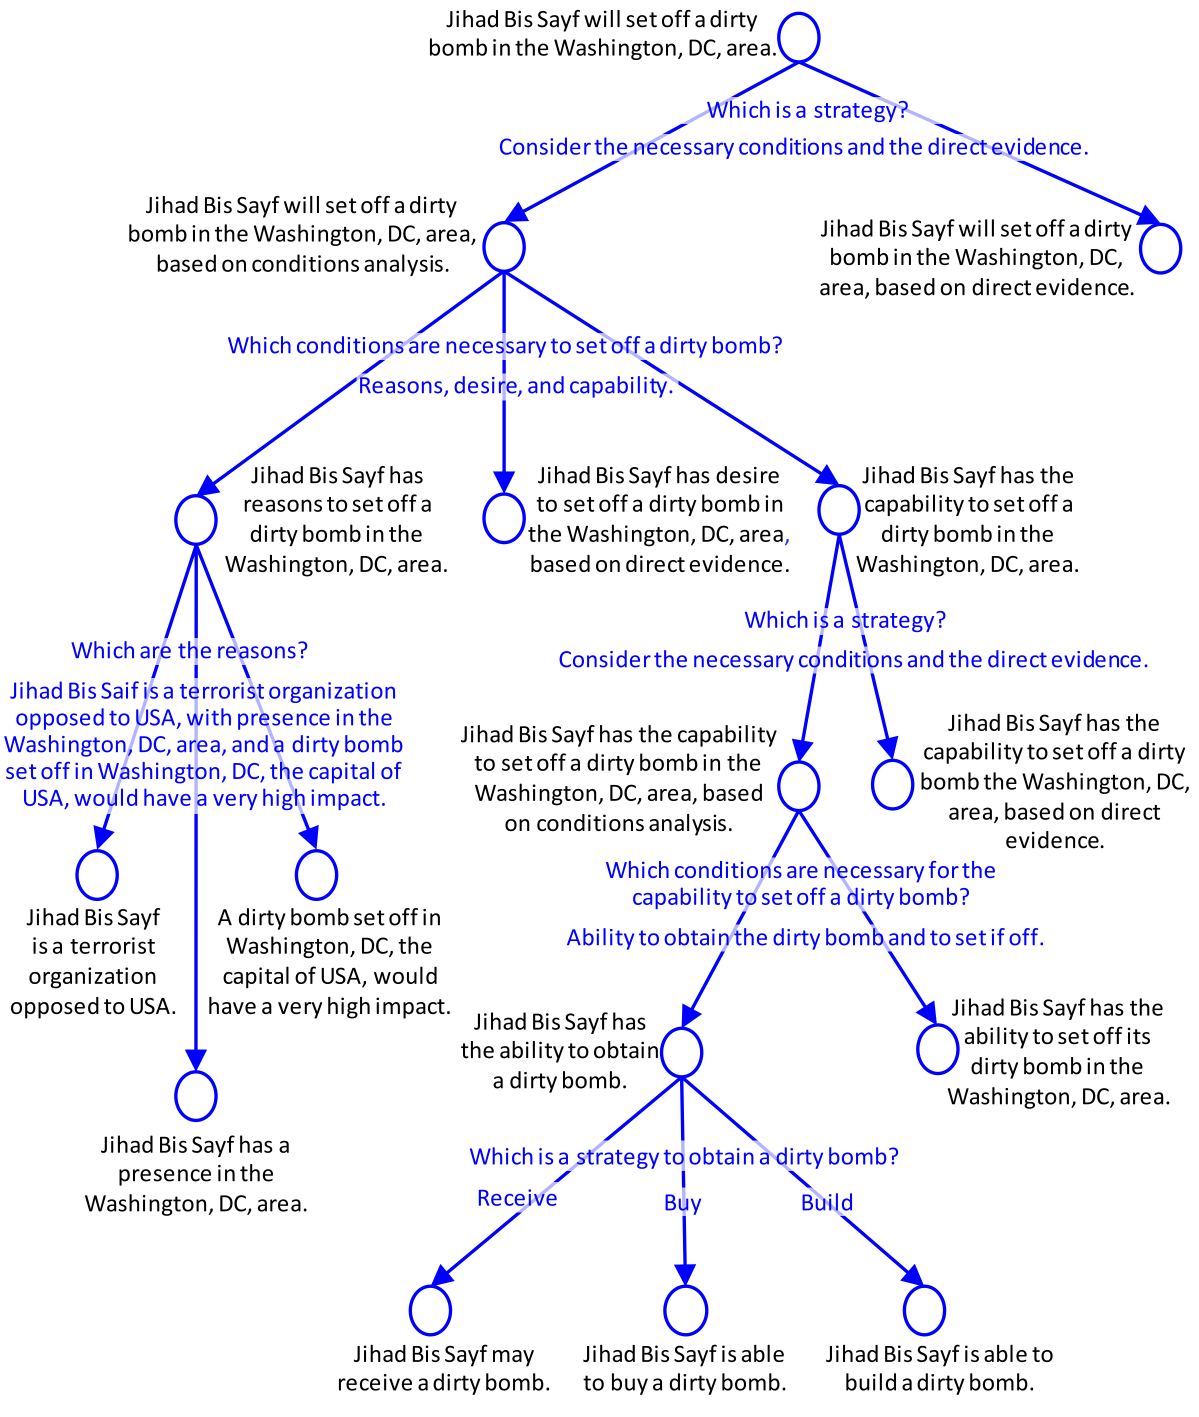
Jihad Bis Sayf will set off a dirty bomb in the Washington, DC, area.
Which is a strategy?
Consider the necessary conditions and the direct evidence.
Jihad Bis Sayf will set off a dirty bomb in the Washington, DC, area, based on conditions analysis.
Jihad Bis Sayf will set off a dirty bomb in the Washington, DC, area, based on direct evidence.
Which conditions are necessary to set off a dirty bomb?
Reasons, desire, and capability.
Jihad Bis Sayf has reasons to set off a dirty bomb in the Washington, DC, area.
Jihad Bis Sayf has desire to set off a dirty bomb in the Washington, DC, area, based on direct evidence.
Jihad Bis Sayf has the capability to set off a dirty bomb in the Washington, DC, area.
Which are the reasons?
Jihad Bis Saif is a terrorist organization opposed to USA, with presence in the Washington, DC, area, and a dirty bomb set off in Washington, DC, the capital of USA, would have a very high impact.
Jihad Bis Sayf is a terrorist organization opposed to USA.
A dirty bomb set off in Washington, DC, the capital of USA, would have a very high impact.
Jihad Bis Sayf has a presence in the Washington, DC, area.
Which is a strategy?
Consider the necessary conditions and the direct evidence.
Jihad Bis Sayf has the capability to set off a dirty bomb in the Washington, DC, area, based on conditions analysis.
Jihad Bis Sayf has the capability to set off a dirty bomb the Washington, DC, area, based on direct evidence.
Which conditions are necessary for the capability to set off a dirty bomb?
Ability to obtain the dirty bomb and to set if off.
Jihad Bis Sayf has the ability to obtain a dirty bomb.
Jihad Bis Sayf has the ability to set off its dirty bomb in the Washington, DC, area.
Which is a strategy to obtain a dirty bomb?
Receive
Buy
Build
Jihad Bis Sayf may receive a dirty bomb.
Jihad Bis Sayf is able to buy a dirty bomb.
Jihad Bis Sayf is able to build a dirty bomb.

Figure 81. Jihad Bis Sayf will set-off a dirty bomb in the Washington, DC area.

6.6 Improving the State of the Art in Critical Thinking

6.6.1 Current State of the Art

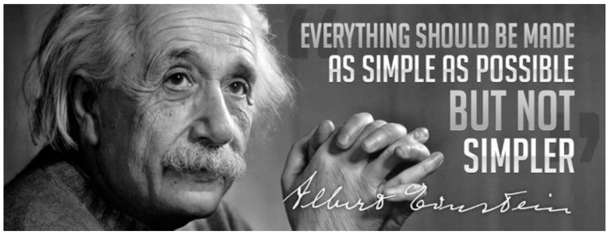

The prevailing approach to critical thinking in intelligence analysis and in many other domains is a *holistic* one where the analyst, after reviewing large amounts of information and mentally processing the data, reaches a conclusion. Consequently, there is very limited *transparency* on how exactly the conclusion has been reached from evidence, what assumptions have been made, and how exactly the probability of the conclusion and the confidence in this probability have been assessed (Marrin, 2011).

Complementary to the holistic approach is the *structured analysis* approach, but the current practice relies on very simple analytic techniques, such as those described in (Heuer and Pherson, 2015). By far the most popular analysis method remains Heuer's *Analysis of Competing Hypotheses* (Heuer, 1999; 2008).

However, in his laudable attempt to simplify the complex analysis task, Heuer made several over-simplifications that affect the result of the analysis process, such as:

- The *credibility* of evidence is assessed holistically, without considering that different types of evidence have different credentials. For example, the credibility of a drone image depends on its *accuracy* and the *reliability* of the camera, while the credibility of an intelligence report depends on the source's *competence*, *veracity*, *objectivity*, and *observational sensitivity,* as discussed in
- The *relevance* of evidence is also assessed holistically, without constructing relevance arguments necessary for more accurate assessments.
- The selection of the most likely hypothesis is ad-hoc. For example, a variant of the ACH method assigns numeric values to credibility and relevance and computes sums of "credibility * relevance" to select the hypothesis with the highest value.
- There is no assessment of the probability of a hypothesis, nor of the confidence in that probability.

The highly-acclaimed book "Critical Thinking for Strategic Intelligence" (Pherson and Pherson, 2021) presents general guidelines for good analysis, but no method for assessing the probability of a hypothesis or the confidence in that probability.

Much more advanced methods use Bayesian probabilistic inference networks but, despite their implementation in advanced analytical tools, such as Netica (https://www.norsys.com/), developing a Bayesian network is a difficult task for an intelligence analyst.

6.6.2 Improving the Analysis of Competing Hypotheses

We think that the simple structured analytic techniques in wide use today could be improved through the methods discussed in this volume, as illustrated below with the *Analysis of Competing Hypotheses (ACH).* Our present comments are based upon the account of a system being developed to implement the ACH approach (Heuer, 2008).

Using the Substance-blind Classification of Evidence

The basis of ACH consists of a matrix in which various items of interest in an intelligence analysis are recorded, as illustrated in the abstract example from .

Table 26. An illustration of Heuer's Analysis of Competing Hypotheses.

Evidence	Source Type	Credibility	Relevance	H_1	H_2	H_3
E_1	Inference	medium	high	C	C	I
E_2	Assumption	high	low	C	I	C
E_3	Intel reporting	low	high	I	C	I
E_4	HUMINT	medium	medium	C	C	C
E_5	Liaison	high	low	C	C	I
E_6	Lack of intel reporting despite vigorous search	low	medium	I	C	C
E_7	Contrarian hypothesis	high	high	C	I	I

In this two-dimensional matrix, analysts first list the substance or content of the evidence in the first column. Then, in the second column, analysts list what Heuer calls "source type", which should guide them in evaluating the credibility and relevance of evidence (columns 3 and 4).

Here are the actual examples Heuer provides of source types: Inference, Assumption, Intel reporting, HUMINT, Liaison, Lack of intel reporting despite vigorous search, and Contrarian hypothesis. One problem with this classification is that the credibility of evidence in the same category (e.g., Liaison) is evaluated based on certain credentials, if it is tangible evidence (e.g., authenticity), and on other credentials, if it is testimonial evidence (e.g., veracity, objectivity, etc.). Thus, this classification does not help with this evaluation.

As discussed in Section 3.3.1, there is a "substance-blind" classification of evidence that emerges precisely from the fact that entirely different credibility questions must be asked of tangible and testimonial evidence. Therefore, an improvement of the ACH method is to use the forms of evidence shown in Figure 38, which will guide the analyst in assessing its believability. In fact, several of Heuer's types can easily be mapped to these forms. For example, HUMINT is a species of testimonial evidence; Intel Reporting may either involve testimonial or tangible evidence; Liaison evidence (obtained from contacts with representatives of friendly or neutral governments) may be either tangible or testimonial in nature. Heuer's "Lack of intel reporting despite vigorous search" qualifies as "missing evidence" having potential inferential value, as discussed in Section 3.3.5.

Heuer uses a very broad interpretation of evidence as "all the factors that influence an analyst's judgment about the relative likelihood of the hypotheses" (Heuer, 2008, p. 253). However, according to the Science of Evidence (Schum, 2009) and as discussed in Section 3.3, all evidence, regardless of its substance or content, has three credentials that must be established by defensible arguments: relevance, credibility, and inferential force or weight. From this point of view, three of the examples provided by Heuer (inference, assumption, and contrarian hypothesis) do not qualify as evidence. We agree with Heuer that they play an important role in evidential reasoning, but they should be accounted for not as evidence (how do we ever establish the credibility of an assumption or a hypothesis?), but as components of arguments. For example, assumptions could be used to assess the relevance or the credibility of evidence, as well as discussed below.

Assessing the Credibility of Evidence

In the third column, ACH requires the analyst to rate the credibility of the "source type" of an item of intelligence evidence as high, medium, or low. First, as discussed in Section 4, we think that it is better to talk about the "credibility" of evidence which may also include "competence" considerations in addition to "credibility" ones.

As discussed in Section 5, creddibility assessments for some items of evidence may be very complex, especially if these items have been obtained through chains of custody (Schum et al., 2009). Cogent has a lot of knowledge about the credibility of evidence and its constituents, and supports the analyst in making these assessments. For example, it knows about the necessity for determining the authenticity, accuracy, and reliability of the demonstrative tangible evidence. It knows that it has to establish both the competence and the credibility of the human sources of testimony. As discussed in Section 4, source credibility and source competence are entirely different characteristics, each with its own ingredients. For example, in order to determine the a source's credibility one has to determine the its veracity, objectivity, and observational sensitivity. On the other hand, in order to determine a source's competence, one would need to determine the its access and understandability. Each of these assessments may be a very complex. It is therefore important to assist the analysts in performing them, for instance, by incorporating into ACH the Cogent procedures for evaluating the credibility of evidence. In particular, the arguments developed with Cogent for establishing the credibility of evidence may include the use of assumptions.

Assessing the Relevance of Evidence

In the fourth column of the ACH table the analyst has to rate the relevance of an item of evidence as high, medium, or low. However, if the relevance arguments are not specifically constructed they can never be subjected to any form of critical reasoning. Cogent can help with this issue because it involves both the top-down and bottom-up argument-structuring methods, drawing upon and Wigmore's concern and methods for assessing the relevance of evidence (Wigmore, 1937).

Assessing the Probability of Hypotheses

The last columns in the ACH table correspond to the hypotheses being considered in the analysis at hand. A significant advancement of ACH over the conventional intuitive analysis approach is precisely the requirement to look at several competing hypotheses. In contrast, conventional intuitive analysis focuses on what is suspected to be the most likely hypothesis and then assesses whether or not the available evidence supports it. This may lead to wrong conclusions because the same evidence may also support other hypotheses.

In the column corresponding to a hypothesis, the analyst grades the bearing of an item of evidence on that hypothesis as either consistent (**C**) or inconsistent (**I**). Then the most likely hypothesis is the one with the least evidence against it, that is, the hypothesis with the least number of **I**s. But there is no indication of how relatively strong any of the **I**s are. Suppose we have ten items of evidence for which H_1 and H_2 have the same number of **I**s. How do we decide which hypothesis to accept, given the fact that the evidence items assessed as **I** under H_1 might be different from the evidence items assessed as **I** under H_2? In their extension of the ACH method, Good and his colleagues attempted to address this issue by associating numbers to the high, medium, and low gradations of credibility and relevance, and scorings the competing hypotheses (Good et al., 2001). The problem with this approach is that numbers applied to hypotheses will have little meaning in the absence of any specific relevance arguments, considerations of credibility and competence attributes for different sources of evidence, and characteristics of the evidence itself. This also applies to any ordinary probability assessments under alternative hypotheses that will have little meaning either in the absence of specific arguments justifying them. In that sense, Good's extension of ACH may do more harm than help because it may provide the analysts with a false sense of confidence rather than encouraging them to give more careful attention to the arguments necessary to justify their conclusions regarding the competing hypotheses.

An additional difficulty with the ACH method is that it requires us to begin with what Heuer calls a full set of hypotheses (Heuer, 2008, p. 256); presumably this means that the hypotheses are mutually exclusive and exhaustive. In some cases, such as in the example Heuer provides, we may consider a set of hypotheses that occur in response to a specific question we have been asked. The analysis example Heuer provides is in answer to the question: *What is the status of Iraq's nuclear weapons program?* The three hypotheses he lists as being a full set are:

H_1: "Dormant or shut down."
H_2: "Has been started up again."
H_3: "Weapon available within this decade."

It could, of course, be argued about whether the hypotheses on this list are in fact either exhaustive or mutually exclusive. For example, H_3 and H_2 are not mutually exclusive. If the weapons program has been started up again (H_2) then we might infer that there might be *at least* one weapon available within this decade (H_3). Conversely, for a weapon to be available, Iraq must have started-up its weapons program. What this shows is that it may be difficult to assure that we have a complete set of mutually exclusive hypotheses. However, if the set of hypotheses is not complete, it may just be the case that the most likely hypothesis is among the missing ones. Cogent may help with this issue by estimating the probabiliyu of each of the competing hypotheses considered or, at least, the one selected through the ACH method. If the ACH-selected hypothesis does not have a high enough probability, then this is an indication that additional hypotheses should be considered.

A simplification made by the ACH method is to consider that both the credibility and the relevance of an item of evidence are independent of the particular hypothesis being considered. Let us consider, for example, an item of evidence revealing the number of years needed by North Korea to develop its nuclear program. This item of evidence is relevant to H_3 but it is not at all relevant to the other two hypotheses. One way to address this issue is to simply estimate a different believability and relevance for each hypothesis.

James Bruce, who is well-known for his valuable work on the importance of epistemology in intelligence analysis, discusses reasons why the ACH method does represent a significant advance over analytic methods that are entirely unsystematic and have so often resulted in a favored hypothesis being uncritically endorsed on a very shaky evidential foundation (Bruce, 2008). He also mentions various reasons why the ACH method enjoys current popularity among many intelligence analysts. However, the example he provides illustrating the virtues of ACH also illustrates one of its most severe limitations. He mentions the unjustified conclusions reached about Saddam's alleged possession and development of WMDs based on the reports provided by "Curveball." Bruce argues that had these reports been subjected to analysis using ACH, a possibly different conclusion would have been reached, especially regarding bioweapons. There are, however, some good reasons why ACH might not have helped regarding this conclusion. The trouble here is that the ACH method says nothing about the attributes of the competence and credibility of HUMINT or the attributes of the credibility of various forms of tangible evidence, such as the diagrams of bioweapons facilities that Curveball provided. We are just as concerned as James Bruce about the epistemology of intelligence analysis but we are especially concerned that intelligence analysts be provided with appropriate background knowledge regarding such tasks as assessing the credibility of sources of evidence and establishing the relevance of evidence on alternative hypotheses. Cogent has significant knowledge about the properties, uses, discovery, and marshaling of evidence that it can share with the intelligence analysts who use it. It also knows about the necessary credibility-related questions. This knowledge can be integrated into the ACH method, as suggested above.

There is a problem that seems endemic in intelligence analysis that the ACH method does not address. The problem is that, in so many situations of interest to the Intelligence Community, we have a seamless activity in which we have evidence in search of hypotheses *at the same time* as hypotheses in search of evidence. Suppose we wish to consider hypothesis H_2, that Iraq's weapons program has been started up again. There is no mechanism in ACH for putting this hypothesis to use in generating new lines of evidence and inquiry. This mechanism should address the question: *What things need to be tested by what evidence in order to sustain this hypothesis?* What this amounts to is generating main lines of argument under H_2, showing what evidence would be necessary to prove or disprove the hypothesis that the Iraqis have started up their weapons program. Many possibilities come to mind, such as the acquisition of necessary materials, the bringing together of necessary talented scientific and technical people, and the development of facilities necessary to make weapons of various sorts. You recognize here that this is what we described in Section 1.5.2 as hypotheses in search of evidence. To put some hypothesis to use requires us to generate arguments from it that will eventually identify classes of observable evidence necessary to sustain this hypothesis. But the world continues to change as we are attempting to understand events in it. The result is that we must continually generate new hypotheses or revise the ones we have constructed. Thus, a major item left out in ACH is the crucial importance of the discovery process in which we have evidence in search of hypotheses at the same time with hypotheses in search of evidence. As discussed in Section 2, Cogent promotes a systematic approach to this complex issue.

Heuer has conceived ACH as a manual method that can be easily used by the analysts and has therefore made many simplifications. The Cogent-inspired improvements suggested above will complicate the original ACH method, but the added complexity will not create any problem if one can use the corresponding components of Cogent. For example, assessing the credibility of evidence could easily be done with Cogent.

Finally, let us notice that many of the improvements suggested above for ACH may be applicable to other evidence-based analytic methods. This suggests that Cogent is an excellent tool for teaching intelligence analysts because the concepts and method for evidence-based reasoning that would be learned with it would help the analysts no matter what specific evidence-based analytic methods they would use.

6.7 Review Questions

119. Which of the following is an *assumption* in the argument that John has a higher IQ than Mark:
    a) John scored higher than Mark on the SAT.
    b) Individuals that score higher on the SAT have a higher IQ.
    c) John scored higher than Mark on an IQ test.

120. Possible answers to a question about a situation are considered:
    a) assumptions
    b) hypotheses
    c) items of evidence

121. Consider the hypothesis "Mark's grades have improved" and the sub-hypotheses:

    Last year Mark maintained a C average.

    Mark is maintaining a B average this year.

    How should you represent them in Cogent?

    a) As two alternative (separate) arguments
    b) As a single AND argument

122. True or false:
    If "Last year Mark maintained a C average" and "Mark is maintaining a B average this year", then it is <u>likely</u> that "Mark's grades have improved."

123. Consider the hypothesis
    Material on the web site of terrorist group X persuaded John to become a terrorist.
    and the sub-hypotheses
    John did not harbor any pro-terrorist views prior to March.
    John visited the terrorist web site 22 times in March.
    John offered his services to terrorist group X in April.

    How should you represent them in Cogent?
    a) As three alternative (separate) arguments.
    b) As a single AND argument.

124. Assuming that "John did not harbor any pro-terrorist views prior to March" and "John visited

the terrorist web site 22 times in March" and "John offered his services to terrorist group X in April", how certain are you that "Material on the web site of terrorist group X persuaded John to become a terrorist"?

a) certain
b) almost certain

125. Consider the hypothesis

President Doe's intelligence service assassinated dissident James Fairley at the behest of Doe

and the sub-hypotheses

The intelligence service of Doe assassinated Fairley
The intelligence service would only assassinate Fairley if Doe gave the order

How should you represent them in Cogent?

a) As two alternative (separate) arguments
b) As a single AND argument

126. True or false:

If "The intelligence service of Doe assassinated Fairley" and "The intelligence service would only assassinate Fairley if Doe gave the order" then it is <u>certain</u> that "The intelligence service assassinated dissident James Fairley at the behest of President Doe."

127. True or false:

If "The intelligence service of Doe assassinated Fairley" and "Doe maintains very strict control over his intelligence service" then it is certain that "President Doe gave the order to his intelligence service to assassinate dissident James Fairley."

128. Consider the hypothesis

John conducted the terrorist attack against the government on 1 May

and the sub-hypotheses

John was planning to conduct a terrorist attack against the government on 1 May
John was involved in several attacks against the government last year

How should you represent them in Cogent?

a) As two alternative (separate) arguments
b) As a single AND argument

129. True or false:

Given the hypothesis "John conducted the terrorist attack against the government on 1 May", the sub-hypothesis "John was planning to conduct a terrorist attack against the government on 1 May" should be assessed as having <u>lower relevance</u> than the sub-hypothesis "John was involved in several attacks against the government last year."

130. Consider the hypothesis

John, who is in his early 40's, is capable of embezzling money from his firm

and the sub-hypotheses

John was convicted of shoplifting as a teenager

John has a reputation of having little or no integrity among most of his current co-workers

How should you represent them in Cogent?

a) As two alternative (separate) arguments

b) As a single AND argument

131. True or false:

Given the hypothesis "John is capable of embezzling money", the sub-hypothesis "John was convicted of shoplifting as a teenager" should be assessed as having lower relevance than the sub-hypothesis "John has a reputation of having little or no integrity among most of his current co-workers."

132. A small bomb was exploded in the city's center near a major university. No one was hurt but a local monument sustained some damage. Initial police investigation indicates that the bomb had been hidden in a black backpack. Person X, a male student at the university and a member of an anti-government, anti-business anarchist group has been identified as a person of interest. One witness has claimed to have seen Person X near the bombing site just before the explosion while another witness claims that Person X was elsewhere at the time.

133. *Hypothesis:* Person X was in the vicinity at the time of the bombing.

Evidence 1 (Eyewitness 1 claims to have seen Person X near the bombing site shortly before the bombing occurred. This eyewitness had never met or seen Person X before.)

a) Favors the truthfulness of the hypothesis and thus is a favoring argument in Cogent.

b) Disfavors the truthfulness of the hypothesis and thus is a disfavoring argument in Cogent.

c) It is not relevant to the hypothesis.

134. Which of the following values is a more accurate assessment of the relevance of Evidence 1:

a) certain (100%)

b) very likely (80-95%)

c) barely likely (50-55%)

135. Which of the following values is a more accurate assessment of the credibility of Evidence 1:

a) almost certain (95-99%)
b) likely (55-70%)

136. Evidence 2 (Eyewitness 2 claims that Person X was at another location at the time of the bombing. Eyewitness 2 is a member of the same anarchist group.)
a) Favors the truthfulness of the hypothesis and is a favoring argument in Cogent.
b) Disfavors the truthfulness of the hypothesis and is a disfavoring argument in Cogent.
c) It is not relevant to the hypothesis.

137. Which of the following values is a more accurate assessment of the relevance of Evidence 2:
a) certain (100%)
b) very likely (80-95%)
c) barely likely (50-55%)

138. Which of the following values is a more accurate assessment of the credibility of Evidence 2:
a) almost certain (95-99%)
b) barely likely (50-55%)

139. Do the fhe five evidence items E1 Washington Gazette, E4 Not checked-out, E5 Forced lock, E14 Grace, and E6 Clyde (from Table 2 and Table 21) conclusively show that the proposition "The canister containing cesium-137 was stolen" is true?

140. What type of evidence item is E9 MDDOT Record in Table 21 on page 226?

141. What type of evidence item is E8 GuardReport in Table 21?

142. In new evidence regarding the dirty bomb example, suppose we have a source code-named "Yasmin." She tells us that she knew a man in Saudi Arabia named Omar al-Massari. Yasmin says she is "quite sure" that Omar spent two years "somewhere" in Afghanistan "sometime" in the years 1998-2000. What type of evidence is this?

143. We return to our asset "Yasmin" who has given us further evidence about Omar al-Massari in our cesium-137 example. Suppose we have a tangible document recording Yasmin's account of her past experience with Omar al-Massari. In this document Yasmin tells us about having seen a document detailing plans for constructing weapons of various sorts that was in Omar al-Massari's possession. What kind of evidence is this and how should it be analyzed?

144. Consider our discussion on the cesium-137 canister. Upon further investigation we identify the person who rented the truck as Omar al-Massari, alias Omer Riley. We tell him that we wish to see his laptop computer. We are, of course, interested in what it might reveal about the terrorists he may be associating with. He refuses to tell us where it is. What kind of

evidence is this?

145. What type of evidence item is E6 Clyde in Table 21?

146. What type of evidence item is E14 Grace in Table 21?

147. Convergent evidence involves evidence about different events, all of which point to the same conclusion. Look at evidence items from Table 21 and identify convergent evidence.

148. One of the unrealistic features about our cesium-137 example is that all the evidence we have so far is harmonious in pointing toward the hypothesis that a dirty bomb containing cesium-137 will be set off somewhere in Washington, DC. In short, we have no contradictory or divergent evidence so far. Could you imagine what some items of dissonant evidence might be.

149. An intelligence analysis has miscarried on an important matter concerning national security and a post mortem hearing is now in progress to determine what went wrong. Attention is focused on the work of analyst $\mathcal{A}$, who provided key judgments during the analysis process. At the hearing a critic notes, "Our main trouble was that we paid too much attention to analyst $\mathcal{A}$ who gave us a biased assessment of the force of evidence E*. A said this evidence very strongly favored hypothesis H_2, which we now know did not occur. H_4 really happened and we have all been embarrassed since we reported that H_2 was true." What could have happened that led this critic to say that $\mathcal{A}$ was biased? Who or what determines analytic bias? And, can analytic bias be prevented?

150. Are there sources of bias that cannot be linked to individual analysts or teams of analysts?

151. An episode of intelligence analysis can go wrong for many reasons. On many accounts we have read, assorted alleged analytic biases are the major reasons why an analysis has gone wrong. In some cases it seems that it is argued that analytic bias is the only reason why an intelligence analysis can go wrong. However, an analysis may go wrong for other reasons not involving bias but rather for an assortment of analytic errors that might be made. What is the distinction between bias and error in intelligence analysis and why is this distinction so important to recognize and discuss?

152. In discussions of bias, so much attention has been based on numerical assessments of the probability of hypotheses considered in intelligence analysis. What other properties of intelligence analysis represent a much more important emphasis in assessing the quality of an analysis?

153. Can analysts ever be criticized for having drawn incorrect conclusions? Or, are some alleged "intelligence failures" actually failures after all?

ANSWERS TO QUESTIONS

Introduction

What Ingredients of Analysis are to be Generated by Imaginative Thought?

1. Characterize each of the questions below with respect to the number of answers it can have.

   Q_1: Will the president select General Martin to be the country's next defense minister?

   Q_2: Which of the country's four-star generals is the president likely to nominate as the country's next defense minister?

   Q_3: Why did the president select General Martin to be the next defense minister?

Answer

Q_1 is a binary question because it has only two possible answers, Yes or No.

Q_2 has a limited number of answers given by the number of the four-start genrals?

Q_3 may have any number of answers, each corresponding to a reason for selecting General Martin to be the next defense minister. The available information will limit this number.

2. A terrorist incident occurred two weeks ago in an American city involving considerable destruction and some loss of lives. After an investigation, two foreign terrorist groups have been identified as possible initiators of this terrorist action: an Al Qaeda-affiliated Group A from Yemen, and a Taliban Group B from Pakistan. Which are some hypotheses we could entertain about this event?

Answer

Some hypotheses we could entertain are the following ones:

H_1: Group A was the one involved in this incident.

H_2: Group B was the one involved in this incident.

H_3: Both Groups A and B were involved in this incident.

H_4: Neither Group A nor B were involved in this incident.

3. You might have reason to suspect that Iran is now supplying improvised explosive devices (IEDs) to a Taliban group in Afghanistan. Since there are other possible sources for these weapons you will have more than one major hypothesis about possible suppliers of these IEDs. What are some of these other hypotheses?

Answer

- Iran is now supplying improvised explosive devices to a Taliban group in Afghanistan.
- Hezbollah is now supplying improvised explosive devices to a Taliban group in

Afghanistan.
- Pakistan is now supplying improvised explosive devices to a Taliban group in Afghanistan.

4. Consider the hypothesis that Iran is now supplying IEDs to a Taliban group in Afghanistan. What evidence we might find concerning this hypothesis?

Answer
We might find evidence showing that the detonators of IEDs we have examined are similar to those previously employed in Iraq that are known to be of Iranian origin.

5. Consider the hypothesis that Al Qaeda-affiliated Group A from Yemen was the one involved in the terrorist incident. What evidence we might find concerning this hypothesis?

Answer
We might find evidence showing that Person X, seen at the location of the incident, traveled to Yemen six months ago.

A Computational Approach to Intelligence Analysis

6. Sometimes we have evidence in search of hypotheses or possible explanations. For example, consider the dog-tag containing the name of one of our soldiers who has been missing since the end of our conflict with Country Z. This tag was allegedly given to a recent visitor in Country Z who then gave it to us. One possibility is that this soldier is still being held as a prisoner in Country Z. What are some other possibilities?

Answer

Some plausible alternative hypotheses are the following ones:
- The dog-tag was taken from the body of this soldier who died during the conflict.
- The soldier lost this tag while evading capture and is still in hiding somewhere.
- The soldier lost this tag or threw it away. He was captured but chose to remain in Country Z after the conflict was over.

7. Sometimes we have hypotheses in search of evidence. Suppose our hypothesis is that Person X was involved in the terrorist incident. So far, all we have is evidence that he was at the scene of the incident an hour before it happened. If this hypothesis were true, what other kinds of evidence might we be able to observe about X?

Answer
Some plausible answers are:
- Evidence of X's association with known terrorist groups.
- Evidence that X has skills in using whatever weapons or materials were used in the

terrorist incident.

- Evidence that X made prior threats to persons against whom the terrorist act was committed.
- Evidence that X has the weapons or remains of materials used in the terrorist incident.

8. True or false: Source A provides information on subject B. If source A is a longstanding enemy of subject B, the credibility of this information, all other things being equal, should be increased.

Answer

False: Because A is a longstanding enemy of B, A is biased against B and the credibility of A should be decreased, not increased.

9. True or false: The relevance of evidence is an assessment of the extent to which the evidence may be believed.

Answer

True: If the evidence is true, then the hypothesis is true. Therefore the relevance of this item of evidence is certain. It is its credibility that is not certain because the source is unverified.

10. Inferential force is an assessment that takes into account:
    a) the credibility of evidence
    b) the relevance of evidence
    c) both the credibility and relevance of evidence

Answer

Inferential force is an assessment that takes into account both the credibility and relevance of evidence.

Anticipatory Intelligence

11. Consider the hypothesis that Countries A and B are about to engage in armed conflict. Here is a report you have just obtained; it says that there has just been an attempt on the life of the president of Country B by an unknown assailant. Why is this report, if credible, relevant evidence on the hypothesis that Countries A and B are about to engage in armed conflict?

Answer

A possible argument showing the relevance of the report is the one fromFigure 82. Notice that each link in this argument represents a source of doubt or uncertainty. Each such source of doubt opens up a new potentially valuable line of evidence you might gather. This shows how argument construction and discovery are linked together.

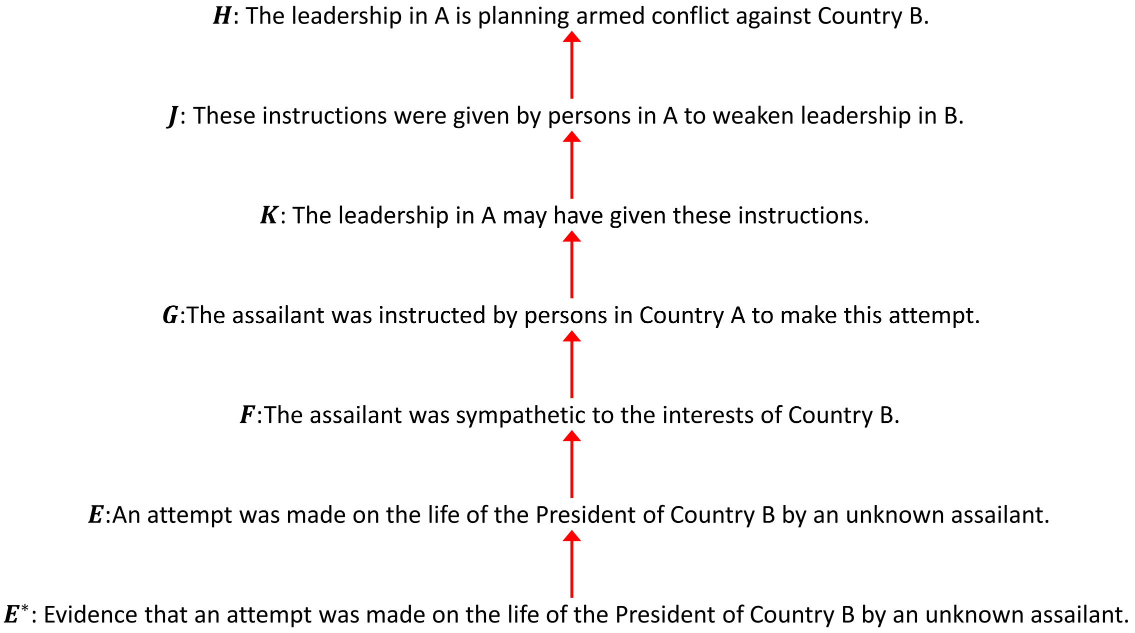

$\boldsymbol{H}$: The leadership in A is planning armed conflict against Country B.

$\boldsymbol{J}$: These instructions were given by persons in A to weaken leadership in B.

$\boldsymbol{K}$: The leadership in A may have given these instructions.

$\boldsymbol{G}$:The assailant was instructed by persons in Country A to make this attempt.

$\boldsymbol{F}$:The assailant was sympathetic to the interests of Country B.

$\boldsymbol{E}$:An attempt was made on the life of the President of Country B by an unknown assailant.

$\boldsymbol{E}^*$: Evidence that an attempt was made on the life of the President of Country B by an unknown assailant.

Figure 82. Possible argument showing the relevance of $\boldsymbol{E}^*$ to a $\boldsymbol{H}$.

12. A car bomb was set off in front of a power sub-station in Washington DC on 25 November. The building was damaged but, fortunately, no one was injured. From the car's identification plate, which survived, it was learned that the car belonged to Quick Car Rental Agency. From information provided by Quick, it was learned that the car was last rented on 24 November by a man named M.
    - Construct an argument from this evidence (E*) to the hypothesis that person M was involved in this car-bombing incident.
    - Suppose that we have determined that evidence E* is believable and therefore we think that M indeed rented a car on November 24. We need additional evidence to assess F, which states that M drove the car on November 25. As discussed in Section 1.5 and illustrated in Section 1.5.3, we can use this hypothesis to guide us in collecting new evidence. Employ this approach to find the needed evidence.

Answer

a) A possible argument from evidence E* to the hypothesis H is shown in Figure 83. Between evidence E* and hypothesis H we show three *interim sources of doubt*. The first E concerns the necessary believability-related foundation step. Just because Quick Car Rental Agency says that M rented the car on 24 November does not entail that this is so. Perhaps Quick made a mistake in their records or in reporting them to us. In other words, someone else may have rented the car on 24 November. Suppose that E is true and that M did rent the car from Quick on 24 November. From E we infer F that M drove the rented car the next day, 25 November. From F, that M drove the rented car on 25 November, we infer G that M parked the car in front of the power sub-station on 25 November. Finally, from G we infer hypothesis H, that M was involved in the car-bombing incident.

b) A possible decomposition of the hypothesis that M drove the rented car on 25 November is shown in Figure 83. This guides us to look for evidence that M did not return the car to Quick and the car was not stolen. We find such evidence from Quick which indicates that M did not return the car he rented, and that M did not report that the car he had rented was stolen.

13. Defendant Dave is accused of shooting a victim Vic. When Dave was arrested sometime after the shooting, he was carrying a 32 cal. Colt automatic pistol. Let H be the hypothesis that it was Dave who shot Vic. A witness named Frank appears and says he saw Dave fire a pistol at the scene of the crime when it occurred; that's all Frank can tell us.

e) Construct a simple chain of reasoning that connects Frank's report to the hypothesis H that it was Dave who shot Vic.

f) The chain of reasoning that connects Frank's report to the hypothesis that it was Dave who shot Vic shows only the possibility of this hypothesis being true. What are some alternative hypotheses?

g) In order to prove the hypothesis that it was Dave who shot Vic, we need additional evidence. As discussed in Section 1.5 and illustrated in Section 0, we need to use this hypothesis to guide us in collecting new evidence. Employ this approach to find the needed evidence.

h) Our investigation has led to the discovery of additional evidence. By itself, each evidence item is hardly conclusive that Dave was the one who shot Vic. Someone else might have been using Dave's Colt automatic. But Frank's testimony along with the fact that he was carrying his weapon, and with the ballistics evidence puts additional heat on Dave. Assess the

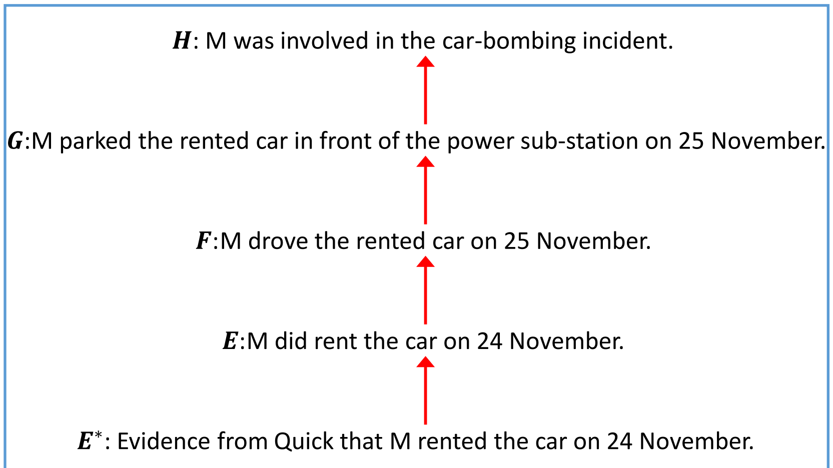
H: M was involved in the car-bombing incident.
G:M parked the rented car in front of the power sub-station on 25 November.
F:M drove the rented car on 25 November.
E:M did rent the car on 24 November.
E^*: Evidence from Quick that M rented the car on 24 November.

Figure 83. Argument from evidence to hypothesis.

probability of the hypothesis that Dave was the one who shot Vic.

Answer

A simple chain of reasoning that connects Frank's report to the hypothesis that it was Dave who shot Vic is shown in Figure 84.

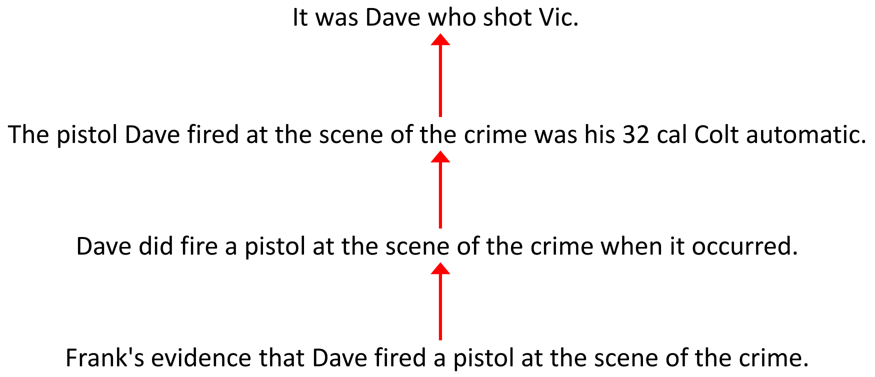
It was Dave who shot Vic.
The pistol Dave fired at the scene of the crime was his 32 cal Colt automatic.
Dave did fire a pistol at the scene of the crime when it occurred.
Frank's evidence that Dave fired a pistol at the scene of the crime.

Figure 84. Possible argument showing the relevance of an item of evidence to a hypothesis.

a) Frank's evidence, indeed, does not tell us very much since, even if Dave did fire his 32 cal. Colt during the crime, this does not prove that Dave shot Vic. Dave might have missed Vic with his shot. It was someone else who actually shot Vic. Remember that all Frank said

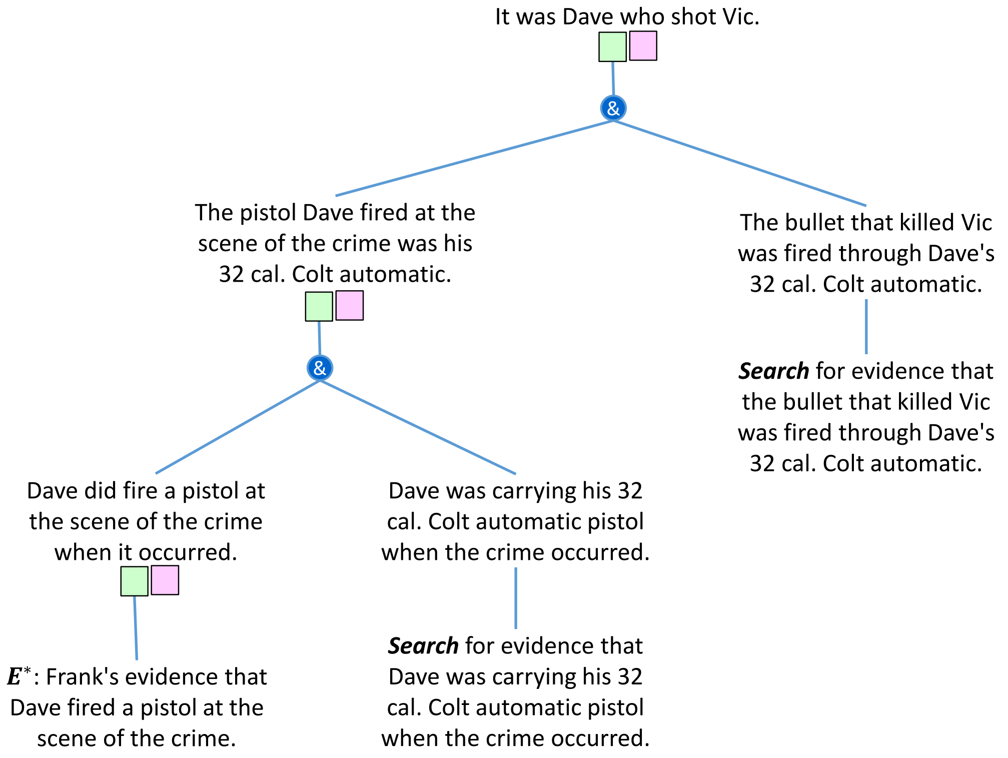
It was Dave who shot Vic.
&
The pistol Dave fired at the scene of the crime was his 32 cal. Colt automatic.
The bullet that killed Vic was fired through Dave's 32 cal. Colt automatic.
&
Search for evidence that the bullet that killed Vic was fired through Dave's 32 cal. Colt automatic.
Dave did fire a pistol at the scene of the crime when it occurred.
Dave was carrying his 32 cal. Colt automatic pistol when the crime occurred.
Search for evidence that Dave was carrying his 32 cal. Colt automatic pistol when the crime occurred.
E^*: Frank's evidence that Dave fired a pistol at the scene of the crime.

Figure 85. Hypothesis in search of evidence.

was that Dave fired a pistol at the scene of the crime when it happened. We do not know which pistol he fired and can only infer that he fired the 32 cal. Colt he had on him when he was arrested.

b) A possible decomposition of the hypothesis that it was Dave who shot Vic is shown in Figure 85This guides us to search for evidence that Dave was carrying his 32 cal. Colt automatic pistol when the crime occurred. We know that, when Dave was arrested, sometime after the shooting, he was carrying his 32 cal. Thus, this is one item of evidence.

c) The reasoning tree in Figure 85 also guides us to look for evidence that the bullet that killed Vic was fired through Dave's 32 cal. Colt automatic. And indeed, the report of a ballistics test shows that the bullet that killed Vic was a 32 cal. bullet that was fired through the 32 cal. Colt automatic that Dave had on him when he was arrested.

d) An analysis tree is shown in Figure 86. What we have here is an example of the <u>evidential synergism</u>. This is a most important evidential subtlety. Two or more evidence items, each considered separately, may say very little. But, when they are considered jointly, they may say a

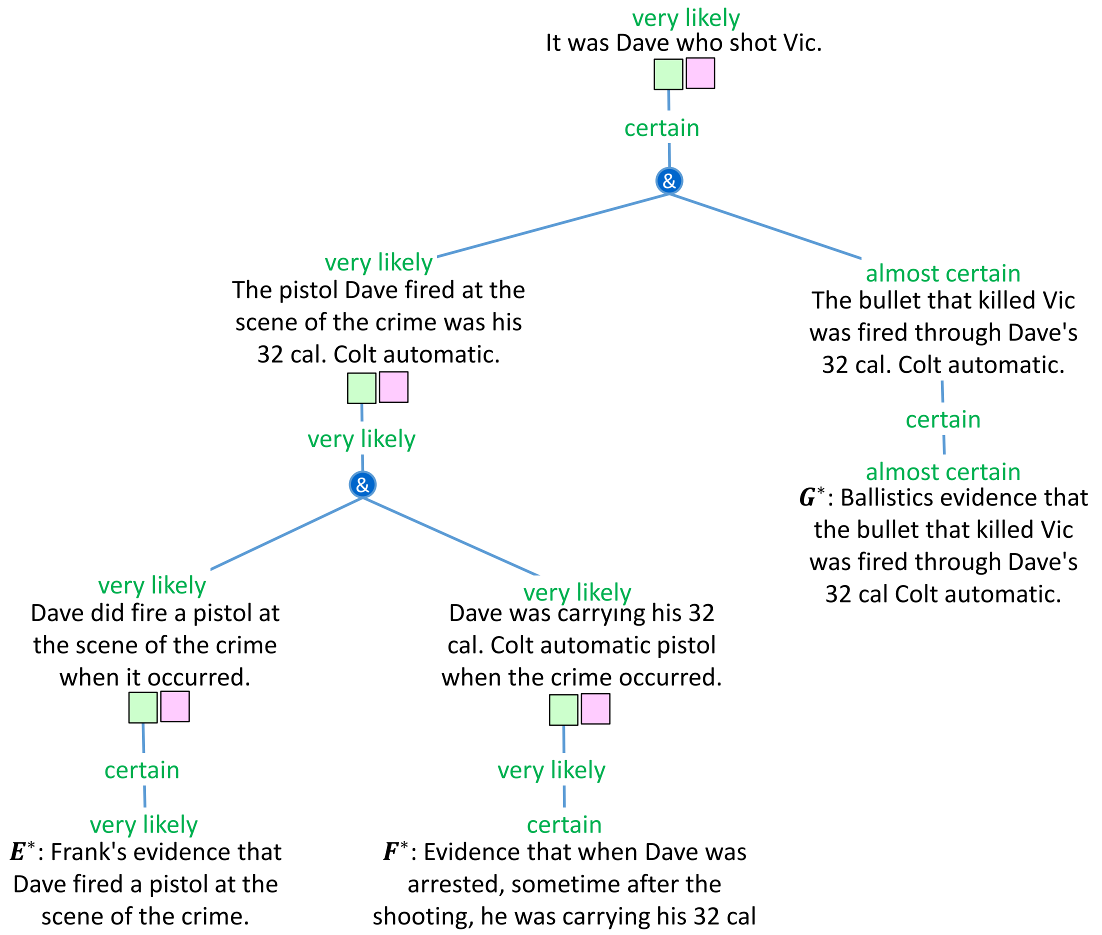
very likely
It was Dave who shot Vic.
certain
&
very likely
The pistol Dave fired at the scene of the crime was his 32 cal. Colt automatic.
very likely
&
almost certain
The bullet that killed Vic was fired through Dave's 32 cal. Colt automatic.
certain
almost certain
G^*: Ballistics evidence that the bullet that killed Vic was fired through Dave's 32 cal Colt automatic.
very likely
Dave did fire a pistol at the scene of the crime when it occurred.
certain
very likely
E^*: Frank's evidence that Dave fired a pistol at the scene of the crime.
very likely
Dave was carrying his 32 cal. Colt automatic pistol when the crime occurred.
very likely
certain
F^*: Evidence that when Dave was arrested, sometime after the shooting, he was carrying his 32 cal

Figure 86. An analysis of the hypothesis that Dave shot Vic.

great deal.

14. Justify the assessments of relevance of E1 Washington Gazette, E3 Not in Locker, and E4 Not checked-out shown in Table 2.

Answer

The relevance of E1 Washington Gazette: If the Washington Gazette report is true, then it is certain that the canister is no longer in the warehouse. For E3 Not in Locker there is the following justification of the relevance: If the canister is not located anywhere in the hazardous material locker, then it is very likely that it is no longer in the warehouse because it is not allowed to be kept anywhere else. Finally, for E4 Not checked-out, if no one at the XYZ Company has checked the canister out, then it is certain that it was not checked out because only XYZ personnel can check it out.

15. Justify the assessments of credibility of E1 Washington Gazette, E3 Not in Locker, and E4 Not checked-out shown inTable 2.

Answer

In general, the credibility of an item of evidence depends on the credibility of its source. The source of E1 is the Washington Gazette. What is, in general, the probability that an item of information from the Washington Gazette is true? Washington Gazette is a very reputable publication, but this article does not mention the source of information, so we will assess its credibility only as likely. The source of both E3 Not in Locker and E4 Not checked-out is Ralph. Ralph is the supervisor of the warehouse and has a reputation for honesty. We can therefore assess his credibility as very likely.

16. True or false: The relevance of the evidence "Willard, who is an unverified source, said that a canister containing cesium-137 is missing from the XYZ warehouse" to the hypothesis "The canister is no longer in the warehouse" is certain.

Answer

True. If the evidence is true, then the hypothesis is true. Therefore, the relevance of this item of evidence is certain. It is its credibility that is not certain because the source is unverified

17. If the credibility of an item of evidence is low and the relevance of this evidence to a hypothesis is high, the inferential force of this evidence is:

    d) high
    e) medium
    f) low

Answer

The inferential force is the smallest between the credibility (low) and the relevance (high), and is therefore low.

18. True or false: In the problem about the missing cesium canister, Ralph's reputation for honesty must be taken into account when assessing the <u>relevance</u> of the information provided by Ralph.

Answer

False. Ralph's reputation for honesty speaks to the accuracy of the information and is a factor in assessing the information's credibility, not relevance. Relevance, which measures how strongly an item of evidence supports a specific hypothesis or sub-hypothesis, is an assessment of the probability that the hypothesis or sub-hypothesis is true, assuming the information is true.

19. In the problem about the missing cesium canister, the credibility of information provided by Ralph (very likely) was assessed at a higher level than the information in Willard's article in the Washington Gazette (likely) because:
    d) Ralph has a reputation for honesty. His position at the XYZ warehouse is not pertinent to the information's credibility.
    e) Ralph has first-hand access to the information. His reputation for honesty is not pertinent to the information's credibility.
    f) Ralph has a reputation for honesty and has first-hand access to the information.

Answer

Two of the most important considerations in assessing the credibility of information are the source's honesty and the source's access to the information. Ralph is honest <u>and</u> has first-hand knowledge that the canister is not in the hazardous materials locker. Willard works for a reputable newspaper but he is getting information on the missing canister from an unknown source, whose honesty and access cannot be determined.

20. True or false: If Willard had a brother who was fired by the XYZ warehouse four years ago, the credibility of the information in Willard's report would increase.

Answer

False. If Willard had a brother who was fired by the XYZ warehouse four years ago, the credibility of the information would, if anything, be lowered. In this situation, Willard might be nursing a grudge against the XYZ warehouse that is prejudicing his views."

21. Consider the hypothesis ”The canister is no longer in the warehouse” and the following items of evidence:

    E1: Willard's report in the Washington Gazette that a canister containing cesium-137 was missing from the STEMQ warehouse in Baltimore, MD.

E2: Ralph, who has a reputation for honesty, reports that the cesium-137 canister […] is not located anywhere in the hazardous materials locker.

True or false: *E1* should have lower relevance than *E2*.

Answer

Relevance measures how strongly an item of evidence (assumed to be true) supports a specific hypothesis. E1 specifies that the cesium was missing from the warehouse, from which we can conclude for certain that *"The canister is no longer in the warehouse."* E2 notes only that the cesium is not in the hazardous materials locker, leaving open the possibility that it is in the warehouse.

22. The argument that "The cesium-137 canister is missing from the XYZ warehouse" needs to establish:
    - only that the canister is not in the warehouse
    - only that the canister was in the warehouse but is no longer there
    - that the canister was in the warehouse but is no longer there and was not checked out from the warehouse

Answer

In the context of this problem, "missing" means the canister is not accounted for. To be "missing" from the warehouse, the canister had to have been there previously, is no longer there, and was not checked out. If the canister was never there, it can't be "missing." Similarly, if the location of the canister is known, it also can't be "missing."

Marshaling Thoughts and Evidence for Imaginative Analysis

Marshaling "Magnets" or Attractors

23. From any collection of information arranged in chronological order, a virtual infinity of different stories might be told. Suppose we have the following three items of evidence whose temporal ordering we have reason to believe:

    E1: Person Y agreed on March 4 to supply us with information about the military in his country.

    E2: On August 1, Source Y supplied us with a HUMINT report saying that the commanding general of the military was planning to launch a coup attempt against the elected leadership in his country on August 15.

E3: On August 18, the leadership in this country announced that the commanding general of its military, along with several members of his staff, were being held in prison.

Think of all the different stories that might be told about why the event predicted in Y's HUMINT did not come to pass.

Answer

Here's just one possible story. The events described in Y's HUMINT did not happen because the evidence about the general and his staff was not believable. This evidence was provided by an opponent of the general who wished to cause trouble for the general. Also of course, the general and his staff now being in jail would reduce the probability of his promoting a coup.

24. The use of index cards and shoeboxes to organize incoming intelligence information is old hat. Are there any computer-based methods you have tried? Have they been helpful in allowing you to generate hypotheses for any analysis you have been working on?

Answer

There is a variety of software systems that are useful for <u>archiving</u> your information. These systems might have been quite useful for the convenience of finding what you are looking for. However, they may be useless for *heuristic* purposes in generating imaginative and productive hypotheses and new lines of inquiry and evidence.

25. Here is a HUMINT source S who tells us that a person P has been assembling explosive devices in his garage. What kinds of questions should you be asking about and of source S? Have another look at the examples shown in Table 3 concerning questions of and about our sources.

Answer

The major questions necessary to ask <u>about</u> source S involve S's competence and credibility. For competence, we should ask, "Did S really have recent access to P's garage and could have seen the explosive materials?" and "Could S have known that these materials were explosives?" For S's credibility we should ask questions about S's veracity, objectivity, and observational sensitivity or accuracy. Questions <u>of</u> S's report concern a variety of matters concerning the meaning of this report. For a start, we would obviously be interested in finding more about person P and his present situation. Is P a member of any group associated with terrorist activities of any kind? And, of course, we should ask questions about the explosive devices S said he saw. Perhaps they were simply fireworks.

26. In Section 2.4.2 we presented an example of the importance of event ordering to Person P who we assume does not wish to have a certain event ordering happen as he left work today;

the event of concern involves his having consumed three double martinis. Can you think of other cases in which event ordering is so important?

Answer

Here comes another example concerning the importance of the Chronology Magnet. Quite often the simple ordering of events can be an important source of ideas having heuristic value. Here is an example concerning a couple of persons. Consider the three events: A = "They met"; B = "They married"; and C = "They fell in love". Now, there are 3! = 6 possible orders of these three events. Suppose all we know about these two persons is the ordering of these three events. From this ordering we can determine a considerable amount of knowledge about these persons. First, suppose the ordering (B→A→C) applies to them: they married, they met, and then they fell in love. This suggests that this couple is a member of a culture in which arranged marriages are common. Now consider the ordering (C→A→-B): they fell in love, they met, and then they married. This suggests that they formed an attachment while they were communicating by phone or computer, they met, and then married. Finally, consider the order (A→C→B): they met, they fell in love, and then they married. This is the ordering most common in our present society.

27. What other questions seem natural to ask about the terrorist organization described in Section 2.4.3?

Answer

Among the other questions to ask are:

- Why have the alleged terrorists chosen this particular target and how difficult would be this target be to defend?
- What other targets would be attractive for this organization to strike?
- If this organization is domestic, how much and what kinds of assistance have they received from which foreign organizations?

28. Hypotheses become most useful "marshaling magnets" for attracting productive combinations of evidence to consider. Here we consider instances of hypotheses in search of evidence we mentioned earlier. As an example, suppose we form the hypothesis that S is a credible source of information about an important event E. This source might either tell us that event E occurred or it did not occur. What evidence would we need to justify our hypothesis that S is credible?

Answer

To decide whether S is a credible source, we will need to answer three questions concerning S's veracity, objectivity, and observational sensitivity or accuracy. The veracity question is: Does S really believe what he/she is telling us? We would not say that S has veracity or is truthful if we believed that S did not believe what he/she just told us. Second, the objectivity question asks:

Did S form a belief about the event reported on the basis of S's sensory evidence, or was it formed on the basis of what S either expected or wished to have observed? We would only say that S was an objective observer if S formed a belief on sensory evidence S obtained. The third question is: If S believed the evidence of S's senses, how good was this evidence. Here we would expose a variety of questions involving S's relevant sensory capabilities in addition to the ambient conditions in which S made the observations.

29. Consider the argument magnet discussed in Section 2.4.5. Here we must consider arguments favoring or disfavoring of sub-hypotheses we are considering. Consider again testing your hypothesis that S is a credible source. What arguments should you be prepared to offer in support of this hypothesis?

Answer

If you considered the veracity, objectivity, and observational sensitivity arguments described in Answer 19 above, you might well encounter both favoring and disfavoring evidence for each of these three arguments.

30. Consider the situation in Section 2.4.5 in which we are concerned with the leakage of classified information from an intelligence office. You have been charged with investigating the activities of a person Y who is suspected of being the leaker. As a result of your investigation you report that person Y can be eliminated from consideration. At some time later, the classified documents are found on a laptop belonging to Y, and Y admits to having been the leaker. You are then confronted by your boss who says, "You managed to seize defeat from the jaws of victory, how could you have been so foolish? You had Y and you let him go. You made all of us look bad and I am considering demoting you." What defense can you offer your boss and perhaps preserve your position?

Answer

Suppose that you have been careful to preserve the evidence you had gathered that led you to eliminate Y from consideration as the source of the leaks. You review with your boss the array of evidence you gathered showing that Y had no access to the leaked documents, that Y never attempted to gather information about matters in these documents from someone else, and that Y's record shows him to have been a stellar member of your office with no apparent reason for ever considering leaking classified information. You then ask your boss what she would have concluded based on the evidence about Y you have so carefully complied. You hope she will reconsider the threat to demote you.

31. We hope you will appreciate the many heuristic virtues of telling yourself stories or constructing scenarios based on evidence you have gathered. From the same collection of available evidence, you may be able to tell an array of different stories depending on the

"gap-fillers" or hypothetical events you include. Every different story you can tell suggests different hypotheses and new lines of evidence you might consider. What is another different story you could tell?

Answer

Here are two gap-fillers that lead to a quite different story about the killing of Premier X of Country A. Between Evidence 1 and Evidence 2, you insert a gap-filler saying that the leadership in Country B regards Premier X of Country A as a fool whose repeated threats against Country B are not taken seriously. The break of diplomatic relations between B and A is due to unfavorable economic pressures by A against B. The new gap-filler between Evidence 2 and Evidence 3, is that the Premier X of Country A was involved in repeated harassment of the wives of X's own governmental officials. The theme of this story is that Premier X not only says reckless things that Country B does not worry about, but X is also unpopular among members of his own government who are given reasons for getting rid of X.

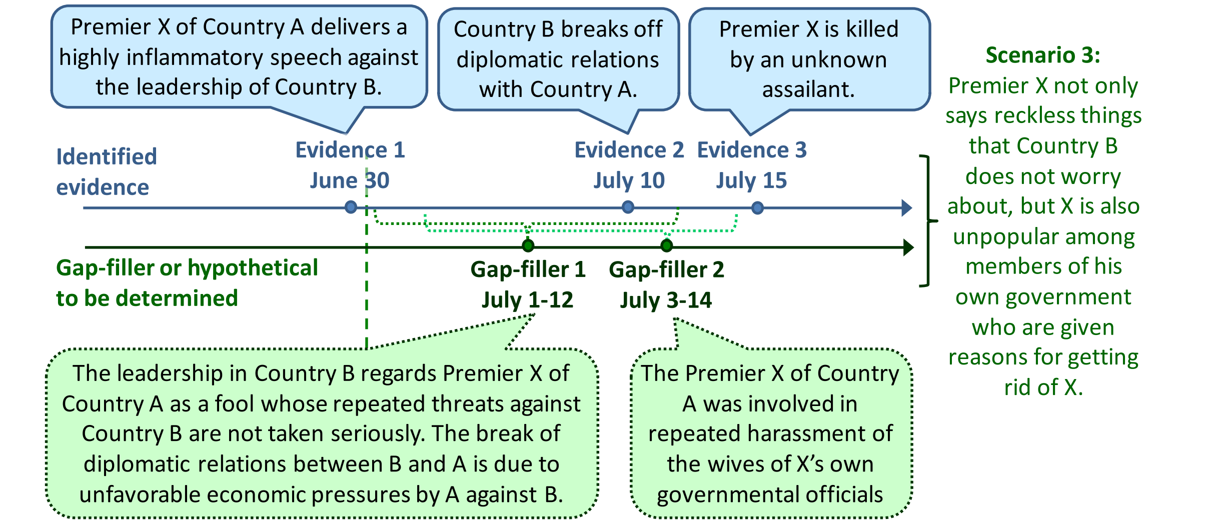
Premier X of Country A delivers a highly inflammatory speech against the leadership of Country B.
Country B breaks off diplomatic relations with Country A.
Premier X is killed by an unknown assailant.
Identified evidence
Evidence 1 June 30
Evidence 2 July 10
Evidence 3 July 15
Gap-filler or hypothetical to be determined
Gap-filler 1 July 1-12
Gap-filler 2 July 3-14
The leadership in Country B regards Premier X of Country A as a fool whose repeated threats against Country B are not taken seriously. The break of diplomatic relations between B and A is due to unfavorable economic pressures by A against B.
The Premier X of Country A was involved in repeated harassment of the wives of X's own governmental officials
Scenario 3: Premier X not only says reckless things that Country B does not worry about, but X is also unpopular among members of his own government who are given reasons for getting rid of X.

Figure 87. A third scenario hypothesized from the events in Figure 30.

The Case of General Alpha

32. Determine which items from Table 4 are relevant to which of the following sub-hypotheses:
    H_1: The Blues enjoy popular support among the citizens of country Orange.
    H_2: The Blues will have the military capability necessary for the insurgency to succeed.
    H_4: The Blue group leadership is adequate to make the insurgency successful.

Answer

First consider the items of evidence are relevant to H_1: The Blues enjoy popular support among the citizens of country Orange. Items E14 and E31 suggest that the Blues are widely known throughout the Orange country and that there are no rival possible insurgency groups. Items E2, E7, and E8 suggest that the deteriorating living conditions due to General Alpha's taxation increases and land takeover are widely known throughout the Orange country. Items E16, E20, E21, E23, and E27, show public resentment over Gen Alpha's arrest and detention of popular religious leader L. And Item E29 shows religious leaders urging resistance movements involving the Blues. Next, consider the five items of evidence we judged relevant on sub-hypothesis H_2: The Blues will have the military capability necessary for the insurgency to succeed. Item E1 shows that the Blues have active communications capabilities, items E18, E19, E22, and E23 show Blue training in progress involving heavy weapons. And items E3 and E4 show that the weapons Blue has are now being assembled in neighboring country Green. Finally, consider the five items of evidence under H_4: The Blue group leadership is adequate to make the insurgency successful. Items E9, E12, and E13 show that Person X's former soldiers will join him in insurgency efforts. Item E28 allows the inference that Person Y is respected in rural areas in Orange. But, item E24 suggests that Person Y may be vulnerable to threats made to injure or kill his family.

33. Develop an argument structure showing how the evidence supports the hypothesis H_1: The Blues enjoy popular support among the citizens of country Orange.

Answer

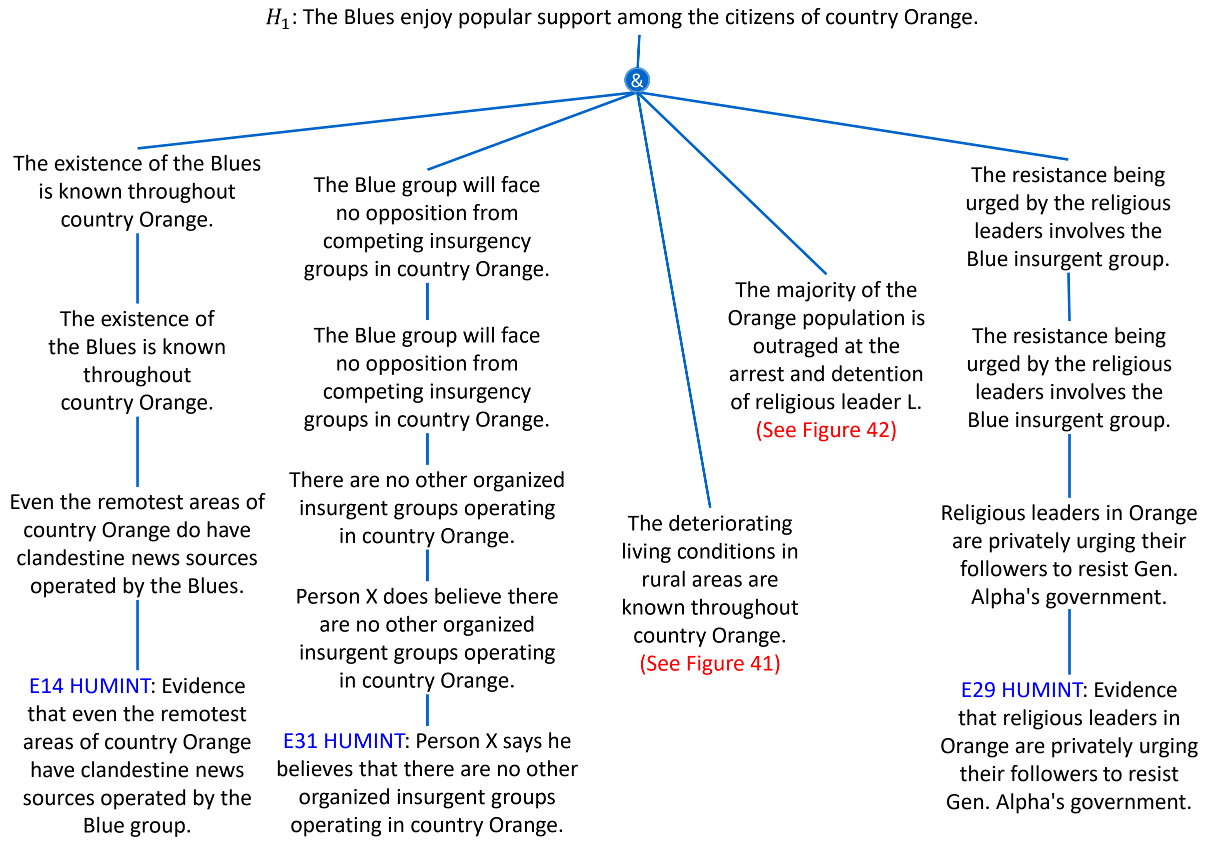
H_1: The Blues enjoy popular support among the citizens of country Orange.
&
The existence of the Blues is known throughout country Orange.
The existence of the Blues is known throughout country Orange.
Even the remotest areas of country Orange do have clandestine news sources operated by the Blues.
E14 HUMINT: Evidence that even the remotest areas of country Orange have clandestine news sources operated by the Blue group.
The Blue group will face no opposition from competing insurgency groups in country Orange.
The Blue group will face no opposition from competing insurgency groups in country Orange.
There are no other organized insurgent groups operating in country Orange.
Person X does believe there are no other organized insurgent groups operating in country Orange.
E31 HUMINT: Person X says he believes that there are no other organized insurgent groups operating in country Orange.
The deteriorating living conditions in rural areas are known throughout country Orange. (See Figure 41)
The majority of the Orange population is outraged at the arrest and detention of religious leader L. (See Figure 42)
The resistance being urged by the religious leaders involves the Blue insurgent group.
The resistance being urged by the religious leaders involves the Blue insurgent group.
Religious leaders in Orange are privately urging their followers to resist Gen. Alpha's government.
E29 HUMINT: Evidence that religious leaders in Orange are privately urging their followers to resist Gen. Alpha's government.

Figure 88. The top part of the argumentations for H_1.

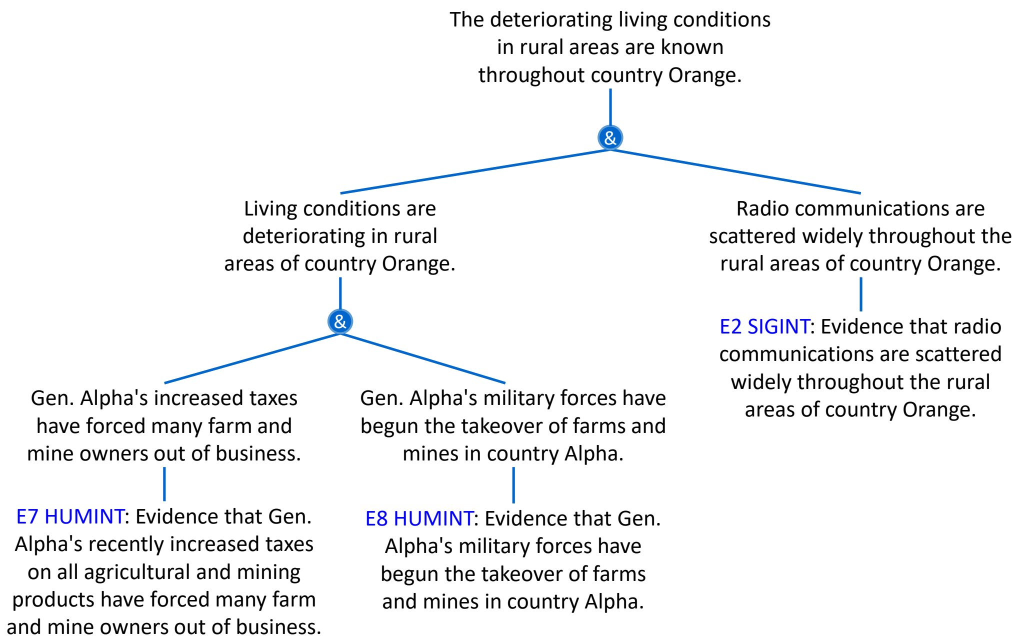
The deteriorating living conditions in rural areas are known throughout country Orange.
&
Living conditions are deteriorating in rural areas of country Orange.
Radio communications are scattered widely throughout the rural areas of country Orange.
&
E2 SIGINT: Evidence that radio communications are scattered widely throughout the rural areas of country Orange.
Gen. Alpha's increased taxes have forced many farm and mine owners out of business.
Gen. Alpha's military forces have begun the takeover of farms and mines in country Alpha.
E7 HUMINT: Evidence that Gen. Alpha's recently increased taxes on all agricultural and mining products have forced many farm and mine owners out of business.
E8 HUMINT: Evidence that Gen. Alpha's military forces have begun the takeover of farms and mines in country Alpha.

Figure 90. Continuation of the argumentation from Figure 88.

34. Develop an argument structure showing how the evidence supports the hypothesis H_2: The Blues will have the military capability necessary for the insurgency to succeed.

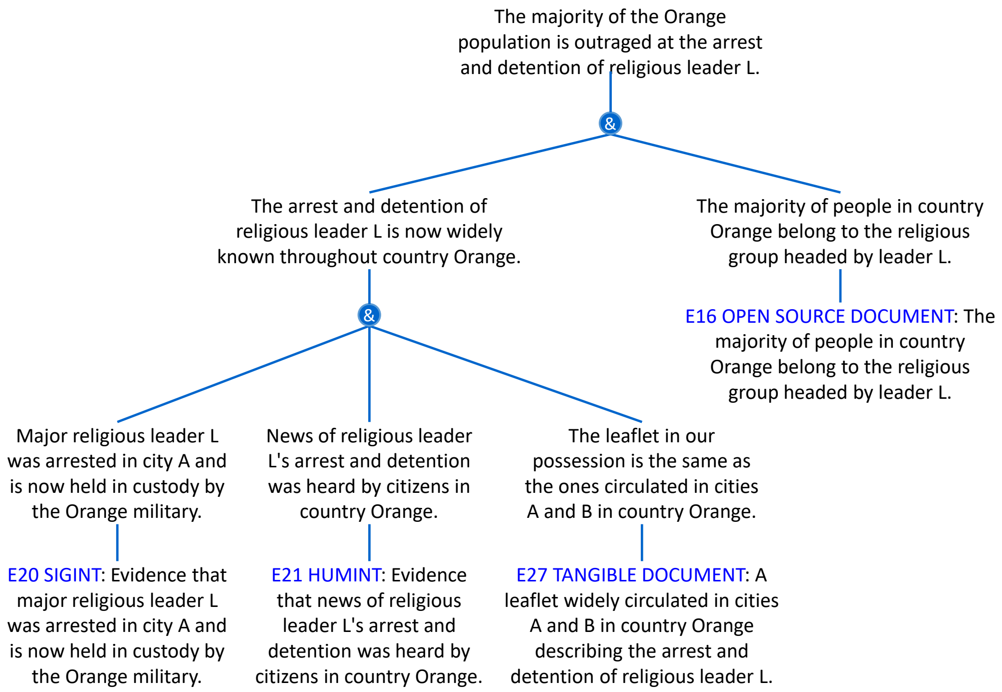
The majority of the Orange population is outraged at the arrest and detention of religious leader L.
&
The arrest and detention of religious leader L is now widely known throughout country Orange.
The majority of people in country Orange belong to the religious group headed by leader L.
&
E16 OPEN SOURCE DOCUMENT: The majority of people in country Orange belong to the religious group headed by leader L.
Major religious leader L was arrested in city A and is now held in custody by the Orange military.
News of religious leader L's arrest and detention was heard by citizens in country Orange.
The leaflet in our possession is the same as the ones circulated in cities A and B in country Orange.
E20 SIGINT: Evidence that major religious leader L was arrested in city A and is now held in custody by the Orange military.
E21 HUMINT: Evidence that news of religious leader L's arrest and detention was heard by citizens in country Orange.
E27 TANGIBLE DOCUMENT: A leaflet widely circulated in cities A and B in country Orange describing the arrest and detention of religious leader L.

Figure 89. Continuation of the argumentation from Figure 91.

35. Develop an argument structure showing how the evidence supports the hypothesis H_4: The

Blue group leadership is adequate to make the insurgency successful.

Answer

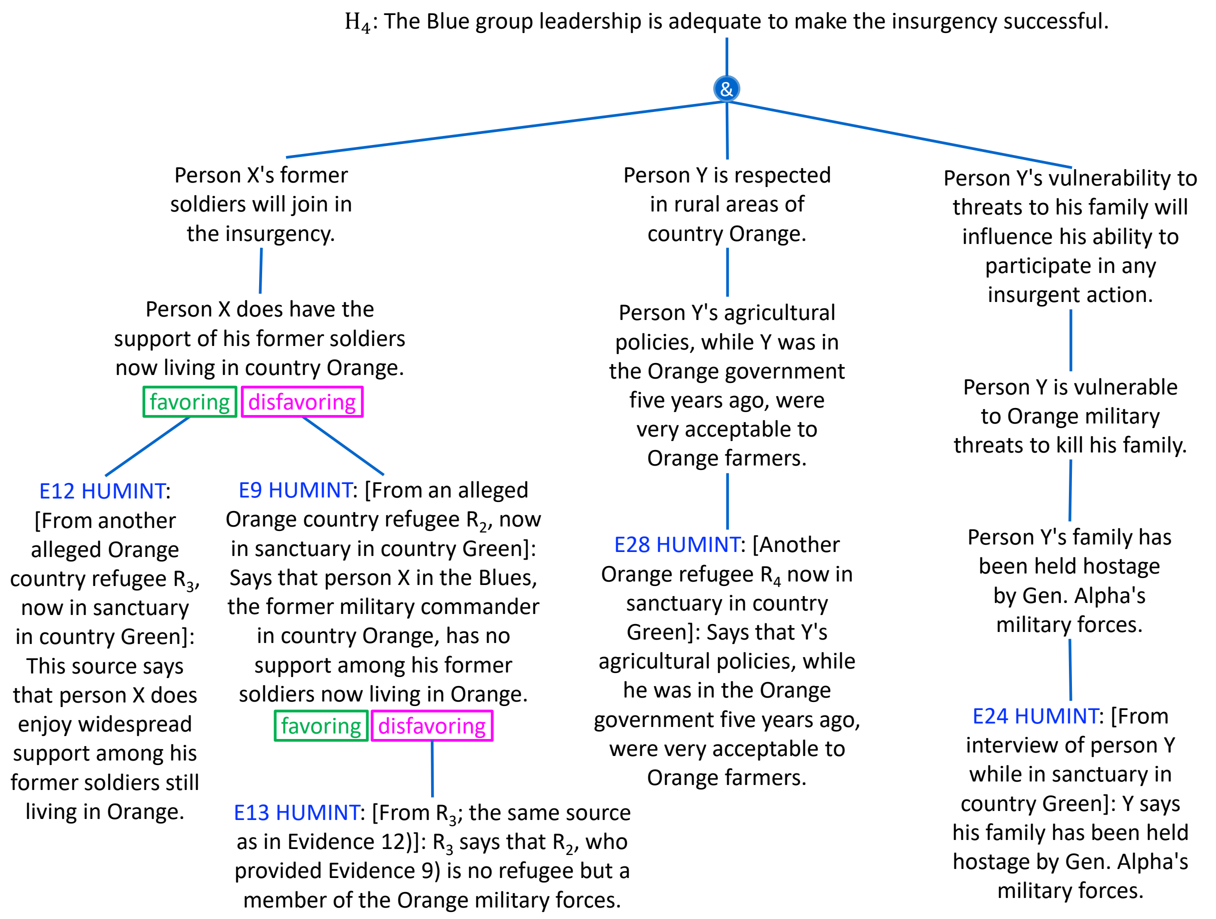
H_4: The Blue group leadership is adequate to make the insurgency successful.
&
Person X's former soldiers will join in the insurgency.
Person Y is respected in rural areas of country Orange.
Person Y's vulnerability to threats to his family will influence his ability to participate in any insurgent action.
Person X does have the support of his former soldiers now living in country Orange.
favoring
disfavoring
Person Y's agricultural policies, while Y was in the Orange government five years ago, were very acceptable to Orange farmers.
Person Y is vulnerable to Orange military threats to kill his family.
E12 HUMINT: [From another alleged Orange country refugee R_3, now in sanctuary in country Green]: This source says that person X does enjoy widespread support among his former soldiers still living in Orange.
E9 HUMINT: [From an alleged Orange country refugee R_2, now in sanctuary in country Green]: Says that person X in the Blues, the former military commander in country Orange, has no support among his former soldiers now living in Orange.
favoring
disfavoring
E28 HUMINT: [Another Orange refugee R_4 now in sanctuary in country Green]: Says that Y's agricultural policies, while he was in the Orange government five years ago, were very acceptable to Orange farmers.
Person Y's family has been held hostage by Gen. Alpha's military forces.
E24 HUMINT: [From interview of person Y while in sanctuary in country Green]: Y says his family has been held hostage by Gen. Alpha's military forces.
E13 HUMINT: [From R_3; the same source as in Evidence 12)]: R_3 says that R_2, who provided Evidence 9) is no refugee but a member of the Orange military forces.

Figure 91. Argumentation for H_4.

Evidence

What is Evidence?

36. Evidence, especially testimonial evidence (e.g., HUMINT), often relates the occurrence of several events. For example, here is an item of evidence coming from a human source where the source tells us several things. The source says: "I observed person P in company with a known Al-Qaeda operative in Vienna, Austria on August 21, 2002. During their conversation, I observed the Al-Qaeda operative taking a stack of $100 bills and a document from Person P. The document looked like a flight manual." This report, carefully parsed, contains several events. Can you identify them?

Answer

The report contains the following events:

- Person P in company with the Al-Qaeda operative;
- The location of their meeting;
- The time of their meeting;
- The transfer of money; and
- The transfer of a document looking like a flight manual.

It would be quite incorrect to say that we have just one item of evidence here; in fact we have five items that might be considered separately.

What frequently happens is that a person evaluating evidence such as this may fail to parse the evidence. The time and location of the meeting between P and the Al Qaeda operative is one matter. What transpired between them is quite another. That they met on a certain date and place may bear on one line of argument. What transpired between them might bear on another line of argument. So, it is usually advisable to decompose evidence items into the finest grain details that seem separate to you. Failure to do this can cause lots of further trouble. The reason is that, in our arguments, we want lines of reasoning to be as specific as we can make them. This same problem can occur with other items of evidence apart from HUMINT. For example, we might encounter the interpretation of a photograph that reveals the occurrence of many events. We should consider evidence about each event separately so that we will not be confused, or confuse others. There are no definite ground rules for parsing complex items of evidence. However, each event being revealed jointly may have special meaning on its own.

37. We have emphasized the fact that evidence about some event is not the same as knowing that this event actually occurred. Suppose we have some evidence E* that event E occurred. We gave an example involving HUMINT evidence E* from a source named Mouse that event E occurred, where E is the event that Amad M. attended an al Qaeda weapons training class near Madyan in Northwest Pakistan in October, 2013. Here we had the task of inferring E based on evidence E*. Can you think of another example in which we infer an event E, based on evidence E*?

Answer

Suppose we have IMINT evidence E* in the form of a photograph concerning event E, where E is the event that person P was running from a car just before a bomb in the car exploded at 10AM yesterday in Kabul, Afghanistan. Our inference that E actually occurred depends on the authenticity, accuracy, and reliability of this sensor evidence. Was the photo taken in Kabul yesterday at 10AM? And was the photo sharp enough for us to identify person P?

The Credentials of All Evidence

Types of Evidence

Recurrent Substance-Blind Combinations of Evidence

38. Indicate and justify what type of evidence is each of the following items:

    (a) A spent shell casing.

    (b) Human source X reports to us that military coup is to be expected in Country A within the next two weeks.

    (c) A captured document.

    (d) You take your car for an oil change, expecting the bill to be about $25. Instead, the bill is $350. You ask the mechanic why an oil change costs so much. The mechanic tells you that you needed a new fuel pump and a new water pump, which he changed in the interests of your safety. You ask the mechanic to let you see these two pumps which you believed were working perfectly. The mechanic tells you how sorry he is that these two items have gone missing.

    (e) Human source Y reports to us that the morale among combat troops in Country B is at an all-time low.

    (f) A sensor image (radar, IR, photo) of some ground installations in a certain territory.

    (g) A table showing the reliability of a certain system after various numbers of hours of operation.

Answer

(a) Real tangible evidence: You can examine this casing to determine what kind of weapon it was fired from.

(b) This is testimonial evidence based on opinion since the source's report concerns an event that hasn't happened yet, if it will at all.

(c) The document itself is real tangible evidence because you can examine it for yourself to see what it says (if you can read the language in which it was written). Whether or not you believe what this document says depends upon its authenticity and what you know about the believability of any sources cited in this document.

(d) This is an example of missing evidence. What inference might you draw from your inability to obtain these parts the mechanic said were defective? One possibility is that the parts were not defective and that you were cheated. The mechanic has these two parts and intends to install them as new parts in someone else's car.

(e) This should be considered testimonial evidence based on opinion since Y does not tell us what he himself directly observed. He must have drawn an inference from various kinds of information.

(f) Demonstrative tangible evidence: the authenticity of the image you see and the accuracy of the sensor that produced it are credibility matters of concern to you.

(g) This is an example of authoritative record. Though you would not be asked to prove that any information you extracted from this table is correct, you might have to consider asking someone to do such calculation if it turned out that this information was in violent conflict with other information you have. Authoritative records can sometimes be incorrect.

Major Sources of Uncertainty in Masses of Evidence

39. In discussing the relevance of evidence we noted that this credential of evidence answers the question: So what? How is this evidence linked to hypotheses we are trying to prove or disprove? Consider evidence E* and event E (that M did rent the car from Quick car rental company on 24 November) in the answer to Question 12. From E we infer F that M drove the rented car the next day, 25 November.) Then consider the hypothesis “H: Person P was acting on behalf of Al Qaeda in this car bombing incident.” How would you defend the relevance of this evidence E* on hypothesis H?

Answer

A possible argument you could construct that P was acting on behalf of Al Qaeda in this car bombing incident is shown in Figure 92. This argument or chain of reasoning from E →A→B→C→H shows what links we might consider in a linkage between E* and H.

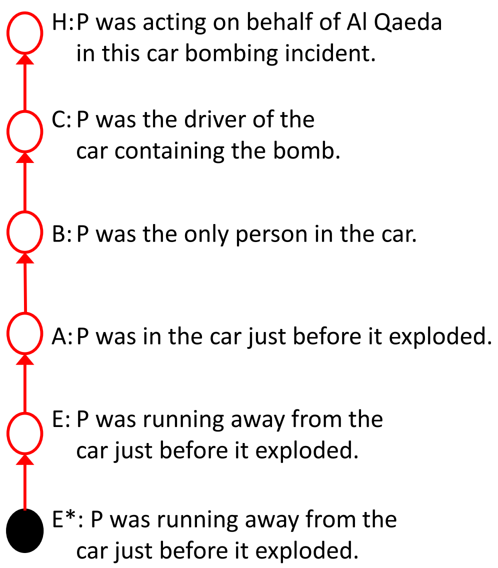
H: P was acting on behalf of Al Qaeda in this car bombing incident.
C: P was the driver of the car containing the bomb.
B: P was the only person in the car.
A: P was in the car just before it exploded.
E: P was running away from the car just before it exploded.
E*: P was running away from the car just before it exploded.

Figure 92. A possible argument.

40. Consider our answer to Question 39 in which we proposed a chain of reasoning between evidence E* and hypothesis H. The links we considered in this relevance argument consisted of the events E, A, B, C, and H. All these links are sources of doubt or uncertainty about these links. In other words, any of these events might not be true. Provide some reasons why these events might not be true.

Answer

The easiest way to answer this question is to consider binary events at each stage consisting of an event and its complement; i.e. the non-occurrence of the event. Here are all the events and their complementary events.

a) For event E and its complement E^C, that event E did not occur. As the photo shows, P was running away from a car just before it exploded. But P might have been running away from something else besides the exploding car. Perhaps P was running away from another car at this scene and was trying to avoid being run down. In addition, of course, we might have been wrong to identify P as the person of interest. For example, suppose the photo was not authentic and was taken two weeks ago and not two days ago. In such a case P cannot be the person in the photo in Kabul, since we know for sure that P was in Pakistan and not in Afghanistan two weeks ago.

b) For event A and its complement A^C, that event A did not occur. Perhaps P was not even in the car before it exploded; someone else was in the car before it exploded.

c) For event B and its complement B^C, that event B did not occur. Perhaps P was not the only person in the car before it exploded.

d) For event C and its complement C^C, that event C did not occur. Person P was not the driver of the car containing the bomb. P may have been in the car but was not the driver.

e) For hypothesis H and its complement H^C, that H is not true. Person P was not acting on behalf of Al Qaeda, but acting on behalf of himself to satisfy some personal grievance. He may also have been acting on behalf of a militant jihadist group other than Al Qaeda.

41. There is no such thing as a perfect argument or one that is absolutely correct and complete. For example, someone can always find one or more missing links in an argument that should be considered. Can you find any missing links in our argument shown in the answer to Question 40?

Answer

As an interested critic, you might easily argue that there are some things missing from the C→H link. You ask how does P being the driver of the car containing the explosives justify the suggestion that P was acting on behalf of Al Qaeda? You argue this requires another chain of reasoning to go along with our existing argument. This additional argument is shown in Figure 93.

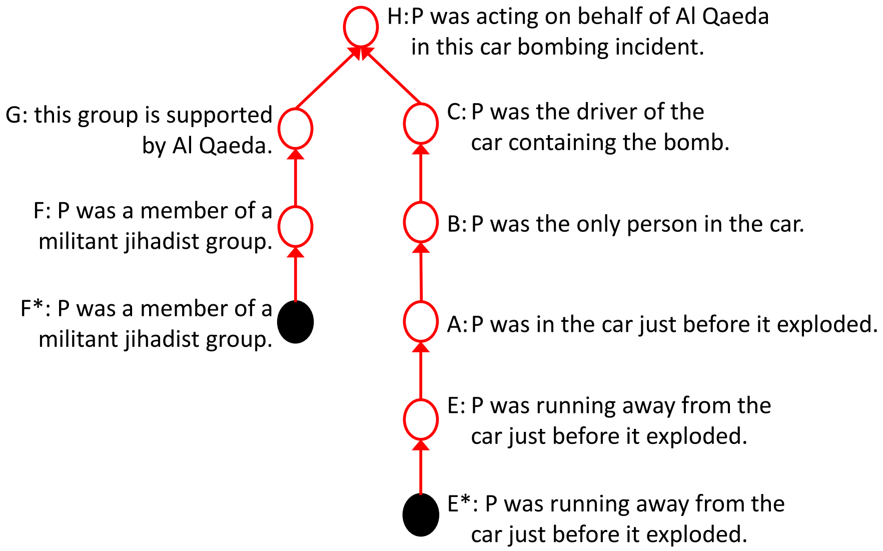
H: P was acting on behalf of Al Qaeda in this car bombing incident.
G: this group is supported by Al Qaeda.
C: P was the driver of the car containing the bomb.
F: P was a member of a militant jihadist group.
B: P was the only person in the car.
F*: P was a member of a militant jihadist group.
A: P was in the car just before it exploded.
E: P was running away from the car just before it exploded.
E*: P was running away from the car just before it exploded.

Figure 93. A revision of the argument in Figure 97.

In your revised argument you suggest that we should consider any evidence, such as F* about event F, that person P was a member of a militant jihadist group, and that F suggests the event G, that this group is supported by Al Qaeda. Then your G taken together with C does a better job of inferring that P was acting on behalf of Al Qaeda in the car bombing incident.

42. Give some examples from your own experience when you have heard people providing information about which they hedge or equivocate.

Answer

You should be able to think of many examples. How many times have you heard your spouse say, “I don't know,” or “I don't remember” in response to a question you have asked? How many times have you said, “I don't know” or “I don't remember” in response to a question your spouse has asked you?

43. Here’s an interesting and important question for you to consider. If you have a weak link in an argument from evidence to some major hypothesis, is it worse to have this weak link at the top or at the bottom of your argument?

Answer

There are different possible answers to this question. On some views, a weak link in an argument may be equally serious no matter where it occurs. However, a weak link at the bottom of a chain of reasoning from evidence to some hypothesis is especially serious since it can invalidate the reasoning at all subsequent links in the chain. Consider the argument from evidence E* to hypothesis H inFigure 92. If it was not Person P who was running away from the car that exploded, this destroys the inferences of events A, B, C, and H, all of which refer to Person P. Similarly, if person P was not a member of any militant jihadist group, as evidence F* asserts in your revised argument in Figure 93, this acts to destroy the inferences of events G and H, since they also concern Person P. What this question illustrates is that the credibility of evidence forms the very basis or foundation for our arguments from evidence to major hypotheses. If we cannot believe what the evidence is telling us, no subsequent argument from this evidence to major hypotheses can be defended.

44. Why are credibility questions different for different forms of evidence and its sources?

Answer

The answer to this question first depends on the nature of the interface between the source and the events reported by a source. First consider an item of HUMINT evidence provided by a human source's report. This report is an example of testimonial evidence. The interface here is this person's relevant sensory systems as it is processed by this person's mental capacities. We must first answer questions concerning this person's competence. Did this person actually make the observation he/she claims to have made? and Did this person understand what he/she was observing? The second question involves this person's credibility attributes: Was this person truthful, objective, and accurate in what this person reported?

Now, many other forms of evidence are produced by a variety of sensing devices and can be viewed as tangible evidence, since their products can be viewed by an analyst to determine what the evidence reveals. The interface here is whatever sensory system(s) the sensor employs. We have different credibility-related questions to answer here than we do for human sources. In the first place we would not be concerned about the veracity or objectivity of a sensing device. What we are concerned about is the authenticity of this tangible item: Is it actually what it is represented to be? Then we are also interested in the accuracy of the sensory device itself.

There is a final matter here that concerns the believability of both testimonial and tangible evidence. This matter concerns the persons and devices that may have done various things to the evidence between the time it was collected and the time when an intelligence analyst first receives it. This collection of persons and their devices is commonly called a chain of custody. Both competence and credibility questions arise at links in chains of custody.

45. The third credential of evidence, its force or weight, depends upon our beliefs about the other two credentials: relevance and believability. Give some examples of this relationship. You can do this in words and without any equations.

Answer

Here are some examples illustrating the dependence of evidential force or weight on relevance and believability. In all of these examples our task involves inferring whether or not hypothesis H is true. First, suppose that if event E occurred, this would strongly favor H being true. In other words E is a very inferentially important event and strongly favors hypothesis H. Now, suppose it also happens that we have a very competent and credible source, Mary, who tells us that event E did occur. The inferential force of Mary's evidence is very strong because she is a very credible source about an important event.

Now suppose we have an event F that only very weakly favors hypothesis H. In addition, the credibility of our source of evidence about event F, Frank, is also very weak. From experience, we have observed that Frank makes lots of observational errors. The weight of Frank's evidence will also be very weak because Frank is not a very credible source of a rather unimportant event.

Now, here come some examples of the matters discussed in our answer to Question 28 about weak links in chains of reasoning. Suppose again that we have the inferentially important event E that strongly favors hypothesis H. But this time our source of evidence about event E is not Mary, but Clyde. Unfortunately, we believe that Clyde could rarely observe the difference between event E and event E^C, not-E. In this case, Clyde's evidence E* would have very little value even though he reports a very important event.

Consider Mary again and suppose she is very competent and credible in reporting event F, that only very weakly favors hypothesis H. The force of Mary's evidence F* cannot be very strong in spite of her strong credibility because she is reporting an event with weak importance. Comparing these last two examples it might seem that a weak link is equally damaging whether it occurs at the top or at the bottom of a reasoning chain. But remember our argument in Answer 28 above when we considered several stages of a relevance argument. In general, it might be more preferable to have a strong believability foundation for a weaker relevance argument than a weak believability foundation for a stronger argument.

46. Indicate and justify what type of evidence is each of the following items:

(a) A spent shell casing.

(b) Human source X reports to us that military coup is to be expected in Country A within the next two weeks.

(c) A captured document.

(d) You take your car for an oil change, expecting the bill to be about $25. Instead, the bill is $350. You ask the mechanic why an oil change costs so much. The mechanic tells you that you needed a new fuel pump and a new water pump, which he changed in the interests of your safety. You ask the mechanic to let you see these two pumps which you believed were working perfectly. The mechanic tells you how sorry he is that these two items have gone missing.

(e) Human source Y reports to us that the morale among combat troops in Country B is at an all-time low.

(f) A sensor image (radar, IR, photo) of some ground installations in a certain territory.

(g) A table showing the reliability of a certain system after various numbers of hours of operation.

47. Give some examples from your own experience when you have heard people providing information about which they hedge or equivocate.

48. Can you provide other examples of mixtures of evidence from your own experience?

49. What inferences might we draw from Omar al-Massari's refusal to provide us with his laptop computer?

50. Can you make up some examples of evidence that corroborates other evidence in our dirty bomb scenario? Ask yourself what items of evidence we now have that you would like to see corroborated.

51. You have two sources reporting the current location of a certain tank column. One source says it is five miles away to the north; the other says it is three miles away to the east. How would you characterize these two items of evidence?

Answer
They are contradictory since this tank column cannot be in two locations at the same time. However, you would certainly inquire about whether or not there are two different tank columns.

52. Two human sources each report observing person X in company yesterday with a known distributor of narcotics. Is this evidence corroborative or convergent?

Answer
These evidence items are corroborative since they both report the same event.

53. You have an aerial photograph you believe shows three surface to surface missiles of a certain sort at map coordinates (x, y); this photograph was taken one week ago. You also have HUMINT from a source who reports observing three missiles of this sort one week ago at nearly the same coordinates. Is this redundant evidence?

Answer
If you believe what the photo tells you, then the HUMINT tells you nothing new and is redundant evidence. The HUMINT is valuable to the extent that you believe the photo to be inauthentic or misleading.

54. One source tells us that T, a known terrorist, was observed at location X at 10:00PM last Friday. Then another source tells us that T was at location Y (100 miles away from X) at this same time. What kind of evidence is this?

Answer

This is contradictory evidence because it concerns evidence about two events that cannot have happened together. Terrorist T cannot have been at both of these locations at the same time.

55. From one source we receive information that Country A is moving military forces in the direction of Country B; we believe this favors hostile action between countries A and B. But another source tells us of recent secret negotiations between representatives of countries A and B that were successful in resolving major differences between countries A and B. This evidence you believe favors the possibility that there will be no hostile action involving A and B. What kind of evidence is this?

Answer

This is divergent evidence because it points to different hypotheses.

56. In assessing these three forms of evidence combinations (i.e., harmonious, dissonant, and redundant), show why it is so necessary to carefully consider the believability of sources of the evidence being considered.

Answer

Careful considerations of the believability of sources of evidence are absolutely vital in assessing the joint impact or force of evidence combinations. In some cases we will observe reversals in the inferential direction of these combinations. Here is a simple example. Suppose we have corroborating evidence from three sources who all report the same event E. However, suppose we have strong evidence that one of these sources does not believe what he just reported; he actually believes that event E did not occur. In other words, this source is not being truthful. If this happens, our evidence pattern is not harmonious but dissonant.

57. We naturally encounter instances of harmonious, dissonant, and redundant combinations of evidence in our daily lives. Provide some examples.

Answer

Here is an interesting pattern of evidence you have observed. Suppose you have a very young daughter who is the light of your life. You look forward to seeing her smiling face as she very excitedly meets you when you come home from work in the evenings. But last evening when you got home she manages only a weak smile and seems subdued. You first notice that her little face seems flushed. So, you place your hand on her forehead and it feels warm. You then take her temperature and it measures 102 degrees F. Here is a case in which you have three different sensors telling you the same thing in different ways. The first is visual and involves your observation; the second involves the tactile sensitivity of your hand. You note that these two items of evidence are not very precise in indicating what you suspect, namely that your child has a fever. So you take a measurement of her bodies temperature using a device designed for this purpose. This example illustrates the many instances in which we employ different sensors leading to combinations of evidence having interesting properties.

In the example just given, you had an example of harmonious evidence leading you to conclude that your child has a fever. You might say that this is a story of corroboration since the three sensors have told you in various ways that your child has a fever. What now worries you is what is causing this fever, a cold or something more serious? But now suppose that your child's measured temperature was 98.5 degrees F. So now this apparent dissonance leads you to explore reasons for the difference in her behavior you have observed. Her weak smile and her subdued behavior might indicate reasons why her day was not all that happy for her and you would wish to explore various possible reasons hoping that you could correct them in various ways.

58. Here is a problem that involves resolving dissonant patterns of contradictory evidence. Suppose you have six persons who say that event E occurred, and only two persons who say that E did not occur. Can you resolve this contradiction by simply counting heads and siding with the majority?

Answer

The trouble with simply counting heads on each side of a contradiction is that it assumes that all witnesses have the same believability, a consequence rarely observed. As Wigmore also observed, our courts do not accept a majority rule interpretation. What matters is the aggregate believability of the witnesses on either side. So, we would be entirely justified in siding with the two witnesses who said that E did not occur if their aggregate believability is greater than that of the six witnesses who said E did occur.

59. Give an example of corroborative redundance.

Answer

If we obtain the same reports from two or more sources, these reports will be corroboratively redundant to some degree.

60. We suspect that person P is a double agent and is presently passing classified information to a certain potential adversary. We thought he worked only for us but now have some grounds for a belief that he is also employed by this adversary. What would constitute absolutely complete evidence that P is in fact a double agent?

Answer

Suppose we begin to list all of the factors we should consider before we can say that we are able to give a <u>complete</u> account of P's presently being a double agent. This list grows very long since we would need a complete account of all of P's activities in the past and a complete account of all associations he might have had with any potential adversary. Even if we were willing to say that we had listed all relevant factors we should consider, we would certainly not be able to obtain evidence bearing on all of these factors. For example, it seems very unlikely that any potential adversary would willingly provide us with information bearing on any association they have had with P.

61. One of our military transport aircraft made a stop at a civilian airfield in Country C, with whom we have had friendly relations over the years. Two days ago this aircraft was destroyed on the ground by an explosive device. There are identified groups in Country C that do not favor C's continued friendly associations with us; one of these groups is called the "Purples." Person Q, believed to be associated with the "Purples" was observed, by a usually believable source, in an aircraft parking area of this airfield just one hour before our aircraft was destroyed. Why is this just inconclusive evidence that the "Purples" were involved in this incident?

Answer

Evidence is inconclusive to the extent that it could be plausibly explained in more than one way. Here are some plausible reasons why this evidence about Q is only inconclusive evidence that the "Purples" were involved in this incident:

i. Q might not in fact have any association with the "Purples;" we only believe he has;
ii. Q might have been in the aircraft parking area in order to board an airplane himself and not with any mischief in mind;
iii. It might not have been Q that our source said he observed; all we have said is that this source is "usually" believable.

All of these factors, plus others that could be identified, make this evidence about Q only inconclusive in nature. Although this evidence does provide <u>some grounds</u> for a belief that the "Purples" were involved in this incident, it is certainly explainable on other grounds as well.

62. We believe the person shown in the photograph is P. However, the figure is blurred. Why is this evidence ambiguous?

Answer

Because the figure is blurred, we cannot be sure that it is P who is the photograph. Therefore this evidence is ambiguous.

63. A radar image shows the possibility of one or more aircraft at a certain location. Why is this evidence ambiguous?

Answer

This evidence is ambiguous because we cannot tell how many aircrafts are at this location. There might be just one aircraft, but there might be several.

64. An observer tells us that he saw a tall man with very dark hair driving away in an old car from the sight of a terrorist incident shortly after it occurred. Why is this evidence ambiguous?

Answer

This evidence is imprecise or ambiguous since we do not know what this observer means by the words: "tall," "very dark," and "old." These words may mean something different to her than they do to you.

65. From your own experience, can you recall other items of ambiguous evidence you have received?

Answer

How about the following: Our wives tell us that they are fixing something we will really like for dinner tonight. That's all they will tell us.

66. Identify the "dots," details, or "trifles" [as Sherlock Holmes called them] in the following intelligence report:

    FBI Report 1: [1 April, this Year. Abdul R is the owner of a Gourmet Foods shop in City A, in Virginia. [Phone number 703-abc-defg]. BB Union National Bank lists Gourmet Foods as holding account number 10701xxxxxx Six checks totaling $35,000 have been deposited in this account in the past four months and are recorded as having been drawn on accounts at the Pyramid Bank of Cairo, Egypt and the Central Bank of Dubai, United Arab Emirates. Both of these banks have just been listed as possible conduits in money laundering schemes.

Answer

Here are the specific dots or trifles in this report:

- A man named Abdul R. owns the Gourmet Foods shop in in City A. in Virginia.
- The phone number at Ramazi's shop is 703-abc-defg.
- Ramazi's shop has account number 1070xxxxxxx at the BB National Bank, in City A.
- Six checks totaling $35,000, drawn on accounts at the Pyramid Bank of Cairo, Egypt and the Central Bank of Dubai, UAI, have been deposited in the above account during the past six months.
- The two foreign banks in Egypt and the UAI are possible conduits in money laundering schemes.

An important fact is that any one of these dots or trifles may not seem important at all until they are connected with dots or trifles in other intelligence reports you receive. What we must learn to do is to carefully parse each intelligence report to observe the particular trifles each report reveals. Considered separately or independently, individual dots may have little meaning. Taken together, however, some combination of dots or trifles may allow you to generate some very productive hypotheses as well as new lines of inquiry or questioning.

67. Identify the "dots," details, or "trifles" in the following intelligence report:

> **FBI Report 2:** Report Date: 12 May, this year [From MI-5, UK]. Riyad Yasser, a UK citizen, was arrested on 1 May, this year following an accident on the M4 Motorway near the Heston Service Area outside of London. Yasser has been an air traffic controller at Heathrow Airport for the past six years. Two kilos of Semtex were found in the trunk of his car. A videocassette of a sermon given by Omar Mahmoud Othman, formerly a Salafi jihad preacher at the Baker St. Mosque in London, was found in Yasser's apartment at # 44, Northumberland Circle, East Bedfont, London. Also found in Yasser's apartment was a note containing several addresses in Canada, the USA, and in Nassau in the Bahamas. The addresses are: 7xx St. Clare St., Montreal; 4xx 11th Street, Miami Beach, FL; 17xx Ferry Ave., Camden, NJ, and **xx** Apple St. in Nassau, The Bahamas.

Answer

Here are the specific dots or trifles in this report:

- Riyad Yasser is a British citizen.
- Riyad Yasser lives in an apartment at # 44 Northumberland Circle, East Bedfont, London.
- Riyad Yasser has been an air traffic controller at Heathrow Airport for the past six years.
- Riyad Yasser was arrested on 1 May, 2003 following an accident on the M4 Motorway near the Heston Service Area in London.
- Two kilos of Semtex were found in the trunk of Yasser's car during his arrest.
- A videocassette of a sermon given by Omar Mahmoud Othman was found in Riyad Yasser's apartment.
- Omar Mahmoud Othman was formerly a Salafi jihad preacher at the Baker St. Mosque in London.

68. What are some causes of evidence to be ambiguous, and how does ambiguity differ from inconclusiveness?

Answer

As illustrated in Example 16 (p. 116), human observers often provide <u>ambiguous</u> accounts of observed events because they cannot make precise judgments under many conditions and are trying to do their best in honestly reporting what they observed. As illustrated in Example 18 (p.

117), analysts are also often criticized for providing ambiguous conclusions in their analyses.

The terms ambiguous, inconclusive, and dissonance are often mixed together when they mean different things. This so often occurs when we hear someone say, “We need to disambiguate the evidence.” For example, we may have very precise and unambiguous evidence **E*** about event **E**, precisely defined, but where event **E** is consistent with any one of several hypotheses and therefore is inconclusive. Similarly, we may have precise and unambiguous evidence about events **E** and **F** where these two events are divergent and therefore are dissonant. The term disambiguate is only used correctly when applied to evidence that is imprecise in some way and that we cannot say what it is telling us.

Relevance of Evidence and Arguments

69. Given a hypothesis that a new missile called X was flight tested on 1 July, which of the items of evidence below would have the highest relevance to this hypothesis:

    E1: The missile that was flight tested demonstrated the same range as missile X.

    E2: A source reported in August that missile X was flight tested on 1 July.

    E3: A source reported in April that a flight test of missile X was scheduled for 1 July.

Answer

E2 has the highest relevance because a report in August that missile X was flight tested on 1 July, if true, would guarantee that missile X was tested on 1 July.

On the other hand, if the missile that was flight tested indeed demonstrated the same range as missile X (evidence E1) then it is possible but not certain that missile X was flight tested on 1 July, provided that there are other missiles with the same range.

Also, a report in April that a flight test of missile X was scheduled for 1 July (evidence E3), even if true, would not guarantee that X was flight tested on 1 July. Plans may have changed between April and July.

70. Given a hypothesis that Mark robbed the bank, which of the items of evidence below would have the lowest relevance to this hypothesis:

    E1: The first five digits of the license plate number of Mark’s car matches the first five digits of a six-digit license plate number of the car that the bank robber used in his escape

    E2: Mark was not at home at the time of the bank robbery

    E3: Mark told a friend a week before the bank robbery that he was planning to rob a bank

Answer

E1 has the highest relevance. It is not certain because someone else may have taken or borrowed Mark's car, and there is a small chance that the sixth digit in the license plate number is different that the license plate number on Mark's car.

E2 proves almost nothing: Mark could have been away from home for any number of possible reasons.

E3 has high relevance but is not conclusive because Mark may have changed his mind; nonetheless, it is much more incriminating than evidence of Mark not being home when the bank was robbed.

Thus, E2 has the lowest relevance.

71. Consider the following arguments:

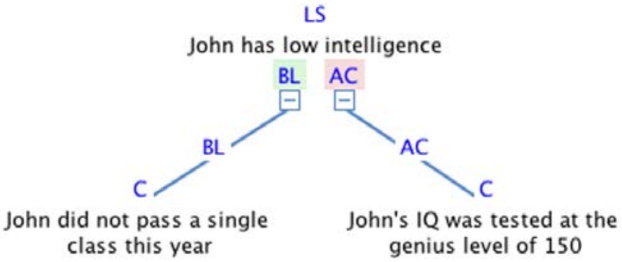
LS
John has low intelligence
BL
AC
BL
AC
C
C
John did not pass a single class this year
John's IQ was tested at the genius level of 150

Figure 94. Cogent argumentation.

The relevance of the sub-hypothesis *"John did not pass a single class"* was assessed as BL because:

a) If John did not pass a single class, it shows conclusively that John has low intelligence.

b) There may be a host of other reasons why John did not pass a single class.

Answer

a) *Incorrect:* The relevance of this sub-hypothesis is much less than conclusive because there are other possible explanations that could sufficiently explain John's failure to pass a single class.

b) *Correct:* John's failure to pass a single class could be explained by many things, including inattention, indifference, and illness. BL (barely likely, 50-55%) allows for these other possibilities.

Credibility of Testimonial Evidence

72. However credible you believe a person might be, you also must give consideration to this person's competence. Provide an example of a credible source which is not believable.

Answer

You might believe that your best friend is certainly a credible source of information. But you would not consult your best friend regarding a medical problem you face unless your friend were also a physician trained in an appropriate specialty.

73. How does intelligence analysis differ from evidential analyses in law trials as far as the completeness of evidence is concerned?

Answer

The answer concerns the process of discovery in intelligence analysis and in law. In law trials, whether criminal or civil, we need conclusions to be reached as quickly as possible, especially when a person's freedom or life is at stake. Of course, having a speedy trial does not always happen but it is one objective in our legal system. At some point, discovery comes to an end and a trial is held in order to reach a verdict from the evidence presented by the parties in contention. You may be familiar with the once-popular *Perry Mason* series on television in which Mason so often discovered new and startling evidence during a criminal trial. In real trials, however, this would be a very rare event. The reason is that attorneys at trial, whichever side they represent, do not want any surprises during a trial. It is said that a good attorney will know exactly what evidence the other side will present. Of course, it is true that an attorney may be surprised by a witness's testimony. For example, an attorney may encounter a hostile witness who says something at trial that is much different from what this witness previously told the attorney he/she would testify. This different testimony may favor the attorney's opposition.

In intelligence analysis, however, it seems accurate to say that discovery is ongoing and never ceases; the world continually changes as we are trying to understand parts of it of interest to our nation's security. As a result, we have evidence in search of hypotheses; hypotheses in search of evidence, and the testing or hypotheses all going on at the same time. In other words, we ask new questions all the time and have no hope of obtaining complete answers to any of our questions. We might add here, that continual discovery and evidential incompleteness is a feature of so many other areas including medicine, science, business, history, and political affairs. For example, you would certainly not retain the services of your physician if she said, "I am going to quit giving you any more physical examinations since I have enough evidence about you now."

74. Can you think of instances in which you might say you have conclusive evidence when this is actually not correct?

Answer

We might describe evidence as being conclusive in some situation by saying that this evidence makes some hypothesis or proposition certain. In the field of law, however, the definition of conclusive evidence does not go quite so far. In law, evidence is said to be conclusive if it is so strong and convincing that it overcomes all proof to the contrary and establishes the proposition in question beyond any reasonable doubt. Notice that this definition does not say "beyond all shadow of doubt," or "with certainty."

Here is something you ought to keep in mind concerning evidence you might say is conclusive. As we have noted in several places, there is a necessary distinction to be made between the actual occurrence of an event and evidence for this event. For example suppose we have evidence **E*** saying that event **E** occurred. The trouble is that **E*** and **E** are not the same. Just because we have evidence **E*** this does not entail that event **E** actually occurred. We might have **E*** when event **E** did not occur. What happens on so many instances is that event **E**, if it really happened, would be conclusive on some proposition or hypothesis **H**, but evidence **E*** is not conclusive on **H**. The reason involves the believability of the source of evidence **E***. In such cases it would be correct to say that evidence **E*** is conclusive regarding hypothesis **H** provided that the source of evidence **E*** is perfectly believable in reporting event **E**. What this says is that what is conclusive would be the actual occurrence of the event reported and we could not be sure that event **E** did happen unless the source of evidence **E*** were perfectly believable.

75. The leadership in Country T has embarked upon an aggressive track regarding its relationship with neighboring countries. We are presently assessing the capability of Country T to wage war on a country with whom we have very friendly relations. We suspect that policy makers in T are considering the development of a certain tactical weapon system we will call W. If they are successful in developing system W, this would give T a decided advantage in any armed conflict they might have with this friendly country. We presently have a source S, a national of Country T, who is an engineer and an expert on the design of weapon systems such as W. Further, she meets regularly with policy makers in Country T regarding the development of tactical weapon systems. Source S has agreed to inform us about deliberations made by policy makers in T regarding the development of system W.

    (a) Does what we know so far about S bear on her <u>competence</u> or <u>credibility</u>?
    Source S now reports to us the following information. She says she was just told by a ranking policy maker in Country T that all plans to develop system W have been suspended because it was thought that such development would be far too expensive.

    (b) What S has told us seems to be good news, but can we believe what she says? What general kinds of evidence should we consider about the credibility of what she has just told us?

76. Show how evidential dissonance and selectivity are related in ways that can be inferentially hazardous.

Answer

Intelligence analysts always have choices about which evidence they will select to include in a reported analysis. They also have choices about what evidence they will seek to discover in their work. Such selectivity also occurs in other areas including law, medicine, science, and elsewhere. In law, the parties in contention often seek to withhold evidence, if they can get by with it, that is unfavorable to their case. In medicine, physicians may withhold certain evidence they believe may be disturbing to their patients. In reports of scientific studies, researchers will sometimes not mention the findings of other researchers that go against their favored conclusions.

Unfortunately, the policy-making consumers of intelligence analyses can also be quite selective in their use of intelligence analyses. Here is an intelligence analysis that is well balanced in the sense that the analysts have incorporated dissonant evidence on many sides of a complex issue involving several major hypotheses. But the consumers of this analysis only focus on and publicize the evidence on the one hypothesis they favor; this has been termed “cherry picking.” Such behavior invites inferential catastrophe since the world does not behave in ways that necessarily correspond to what our policy-makers wish or expect to happen. In some cases, persons and organizations will deliberately act in ways counter to what our leaders desire or expect.

So, any failure to give a balanced account of dissonant evidence available on rival or competing hypotheses invites trouble when we selectively emphasize the evidence on just one favored hypothesis; we may be favoring the wrong hypothesis.

77. It can be argued that of all the inferential issues involved in intelligence analysis, the most important and interesting ones involve the believability of evidence and its sources. Give some reasons why this is so.

Answer

Here are some very good reasons why careful attention to the believability of evidence and its sources is absolutely vital in intelligence analysis as it is in other contexts. First, as we have noted, inferences concerning the believability of our evidence, whatever kind of evidence we have, form the very foundation of all subsequent arguments we make concerning the relevance and force or weight of the evidence. If we cannot believe what the evidence is telling us, there is no point in troubling about its relevance and force or weight. But we may not have a complete lack of belief about what the evidence is telling us. In such cases combining our inferences about the believability of our evidence with our ideas about the relevance and force of this evidence involves many complex and subtle judgments. These difficulties are compounded when we have masses of evidence and complex inference networks to consider, as we do in most intelligence analyses.

When we consider what forms our evidence items take, matters become very interesting in addition to being very difficult. The reason is that we must ask and try to answer different questions for different forms of evidence. If the evidence is tangible, we are concerned about its authenticity, accuracy, and reliability. If the evidence is testimonial from human sources, we are first concerned about the competence of these sources and whether they had access and understanding. But we are also interested in the credibility of these sources and must answer questions concerning their veracity, objectivity, and observational sensitivity. Trying to answer these questions about human observers always involves some interesting and sensitive matters.

78. The leadership in Country T has embarked upon an aggressive track regarding its relationship with neighboring countries. We are presently assessing the capability of Country T to wage war on a country with whom we have very friendly relations. We suspect that policy makers in T are considering the development of a certain tactical weapon system we will call W. If they are successful in developing system W, this would give T a decided advantage in any armed conflict they might have with this friendly country. We presently have a source S, a national of Country T, who is an engineer and an expert on the design of weapon systems such as W. Further, she meets regularly with policy makers in Country T regarding the development of tactical weapon systems. Source S has agreed to inform us about deliberations made by policy makers in T regarding the development of system W.

(a) Does what we know so far about S bear on her <u>competence</u> or <u>credibility</u>?

Source T now reports to us the following information. She says she was just told by a ranking policy maker in Country T that all plans to develop system W have been suspended because it was thought that such development would be far too expensive.

(b) What S has told us seems to be good news, but can we believe what she says? What general kinds of evidence should we consider about the credibility of what she has just told us?

Answer

(a) What we know so far about S bear only upon her <u>competence</u>. All we know so far is that she is an expert who seems to have appropriate contacts with policy makers in T.

(b) Evidence bearing upon S's credibility including her veracity, objectivity, and observational sensitivity. Though she may have given us reasonably accurate information in the past, we wonder whether or not she is still to be believed. You may recognize that what she has reported to us is information she says she received at second hand from another source. So, we have the credibility of her source of information to consider. We should, of course, compare this evidence from S with other evidence we have about the possibility of system W developments in Country T.

79. Why is the competence of sources of testimonial evidence so important and why is competence not the same as credibility?

Answer

Human sources of testimony report the occurrence of events they claim to have observed in some way using their senses. In order to do so, they must have been in a relevant place to make these observations or to have had access to these events. Here is a source S who tells us that he observed some important event E. If there is evidence that S was not in a place where E could have occurred or not, we would naturally question the believability of S's report concerning event E. Also, we would naturally consider the extent to which S could have understood what he was observing. For example, we would also question the believability of S's report if his account of event E is muddled or incoherent in some way. Now, suppose we can perfectly believe that source S had entirely appropriate access and understanding regarding what he was observing, but we also believe that S's report of event E is incorrect in some way regarding S's credibility. Event E may not have occurred counter to what S reports. S may not have been truthful, objective, or accurate as far as his observation and reporting are concerned. But these are credibility attributes quite different from competence attributes. So, it is a violent non sequitur to say that we can believe a source's report just on competence grounds since competence and credibility involve entirely different attributes, different questions and different evidence about a human source.

80. The credibility attribute, veracity or truthfulness, is widely discussed and often widely misunderstood or mistakenly attributed. It seems obvious that the veracity attribute is a property of human sources of evidence. It is very hard to imagine a mechanical or electronic sensor attempting to mislead us, willfully or not, in providing a report. Such reports can of course be incorrect but for reasons not involving truthfulness. Provide some examples of the uses and misuses of the attribute veracity.

Answer

Mistakes are often made in discussing a human source's veracity or truthfulness. We may be tempted to say that a source is truthful only if the event he/she reports did in fact occur. But, this overlooks the fact that, if a reported event did not occur, the person reporting this event might simply have been mistaken or not objective in his/her observation. We would not say that a person is being truthful in making a report unless this person believed what he/she reported to us. In many situations people report against their beliefs; they tell us things they do not themselves believe have happened or will happen. There are situations in intelligence analysis in which human sources are instructed what to tell us. In such cases, a human source may have no particular belief, one way or the other, about the event(s) they are instructed to report. So, veracity or truthfulness concerns whether a human source believes what he/she reports to us. Establishing what a human source actually believes cannot, of course, be done directly. But there are various time-tested strategies for making inferences about whether a human source believes what he/she reports to us.

Here are some examples of situations in which a human source's veracity could reasonably be questioned. First, consider a source S who is inconsistent in his/her reports. Source S first tells us that event E occurred but later tells us that E did not occur. A source S may tell different things to different persons. Source S tells the FBI that event E occurred, but has also told CIA that E did not occur. In another situation, we have evidence that source S is under the influence of a hostile group and can be exploited by this group. Examining the past records of the behavior of source S, we discover instances of possible untruthfulness in S's behavior. As yet another example, we find S's report of event E inconsistent with what a different source T has said: this source T reports that event E did not occur. Of course we would be just as concerned about the veracity of source T as we are about the veracity of source S.

We now consider situations in which are either doubtful concerning veracity or are misleading as far as veracity is concerned. First, it is quite common to use the demeanor and bearing of a human source as an indicator of his/her veracity. Is the source agitated or composed while telling us about event E? The idea is that any hint of agitation or discomfort on the source's part is an indicator of lack of veracity. The truth is that demeanor and bearing is actually a very poor indicator of veracity. Very agitated and uncomfortable sources may be entirely truthful and very composed sources being untruthful. Polygraph analyses are also doubtful as indicators of veracity. The results of polygraph analyses are taken seriously by intelligence agencies but are inadmissible in our courts. One argument is that the results of a polygraph examination may say more about the examiner than it does about the person being examined. Finally, the word untruth is often misapplied. Here is a source S who faces a critic who says the following: "This source S is not being truthful in reporting that event E occurred, since he has such poor eyesight and made only a brief observation under a low level of ambient illumination." This evidence concerns S's observational capacities and not his veracity. What this evidence shows is that S may only have been mistaken since S truly believed that event E occurred.

81. The objectivity attribute of the credibility of a human source is widely overlooked in spite of its importance. What is odd is that lack of objectivity is much discussed in common discourse. We so often hear that we all from time to time believe things because we want to believe them in spite of having little or no evidence for them. One matter of interest concerns the possibility that objectivity is not only a property exclusive to human sensors but also relevant to sensing devices. Provide some examples of the objectivity attribute in various situations.

Answer

Here comes a human source S whose veracity was wrongly questioned. What happened was that source S made a competent observation two weeks ago and told us then that event E occurred. At a hearing today S now tells us that event E did not occur and so we question S's veracity saying that his testimony has been inconsistent. What we say is that S has not gotten his story straight. But S now explains this inconsistency. Here is what S tells us today. S says: "I remember telling you that event E occurred two weeks ago. However, on careful reflection, I have come to the view that I did believe that E occurred simply because I expected to observe E and not because my senses told me that E occurred. My belief now is that E did not occur." Our dictionary defines an objective observer as a person who bases a belief on sensory evidence and not on what this observer wished or expected to observe. Since we have no other evidence of S's lack of veracity we conclude that his inconsistency is based on his admitted lack of objectivity in his observation.

As we noted, objectivity concerns the formation of our beliefs. But our beliefs are elastic and so frequently change over time, often in response to new evidence and later experiences. What is also true is that we often cannot remember what we believed at some time in the past. What is also true is that objectivity involves various biases that can easily plague all of us. In certain situations we may be biased on favor of believing certain events because we want them to occur or expect them to occur.

But there is another species of observer bias that may concern veracity, we may call this bias "testimonial bias" instead of "objectivity bias." Suppose a source S who we believe would tell us that event E occurred regardless of what S truly believed. We might discover, for example, that S thinks the occurrence of event E would be good new to us and that S, believing this, would enjoy being the bearer of good news. It happens that if we believed that S would report the occurrence of event E whether S believed E occurred or did not believe that E occurred, then S's report of event E has no value at all.

It is possible that we can link the attribute objectivity to sensing devices as well as to human observers, as least as far as bias is concerned. It is possible, for example to adjust the gain or amplification of a sensor in some situations to display images of certain events when captured energy is beyond certain thresholds. These thresholds can be adjusted or biased to capture certain events at the expense of capturing other events.

82. It is no mystery that the sensory sensitivity of human observers is an important credibility attribute. But this sensitivity analysis must always be accompanied by evidence concerning other elements of the physical condition of the observer as well as the particular environmental situation in which a sensory observation was made. Provide some examples of the importance of these situations.

Answer

Here is a source S whose visual acuity is 20/20 and whose physical condition was excellent at the time of her observation. Further, S was not inebriated or under the influence of any narcotic. S has reported that event E occurred. We ask her to provide an account of how sure she was that event E did occur. Here's what she says: "Actually I was distracted as I was observing what I believe was E happening. Other events were happening and I did not devote my entire attention to event E." How certain should we be in believing that E happened, based on S's report? Here is a different human source T, who reports that event F occurred. T reports overhearing a discussion between two high level officials in an adversary country discussing the occurrence of event F in two days. Source T is asked how sure he is that these two officials were discussing event F. We know that source T has very acute hearing. But T tells us that last evening he overheard this discussion in a very crowded bar in the capital city of this country. Finally, we know that Source T did make a relevant observation to ground his report that event F occurred yesterday. We know that T is in perfect physical condition. But we have evidence that T was quite intoxicated in the bar at the time he says he overheard the two officials discussing event F. If true, this evidence puts obvious constraints our beliefs that event F occurred, based on what we know about Source T.

83. Why is the competence of sources of testimonial evidence so important and why is competence not the same as credibility?

Answer

Human sources of testimony report the occurrence of events they claim to have observed in some way using their senses. In order to do so, they must have been in a relevant place to make these observations or to have had access to these events. Here is a source S who tells us that he observed some important event E. If there is evidence that S was not in a place where E could have occurred or not, we would naturally question the believability of S's report concerning event E. Also, we would naturally consider the extent to which S could have understood what he was observing. For example, we would also question the believability of S's report if his account of event E is muddled or incoherent in some way. Now, suppose we can perfectly believe that source S had entirely appropriate access and understanding regarding what he was observing, but we also believe that S's report of event E is incorrect in some way regarding S's credibility. Event E may not have occurred counter to what S reports. S may not have been truthful, objective, or accurate as far as his observation and reporting are concerned. But these are credibility attributes quite different from competence attributes. So, it is a violent non sequitur to say that we can believe a source's report just on competence grounds since competence and credibility involve entirely different attributes, different questions and different evidence about a human source.

84. The credibility attribute, veracity or truthfulness, is widely discussed and often widely misunderstood or mistakenly attributed. It seems obvious that the veracity attribute is a property of human sources of evidence. It is very hard to imagine a mechanical or electronic sensor attempting to mislead us, willfully or not, in providing a report. Such reports can of course be incorrect but for reasons not involving truthfulness. Provide some examples of the uses and misuses of the attribute veracity.

Answer

Mistakes are often made in discussing a human source's veracity or truthfulness. We may be tempted to say that a source is truthful only if the event he/she reports did in fact occur. But, this overlooks the fact that, if a reported event did not occur, the person reporting this event might simply have been mistaken or not objective in his/her observation. We would not say that a person is being truthful in making a report unless this person believed what he/she reported to us. In many situations people report against their beliefs; they tell us things they do not themselves believe have happened or will happen. There are situations in intelligence analysis in which human sources are instructed what to tell us. In such cases, a human source may have no particular belief, one way or the other, about the event(s) they are instructed to report. So, veracity or truthfulness concerns whether a human source believes what he/she reports to us. Establishing what a human source actually believes cannot, of course, be done directly. But there are various time-tested strategies for making inferences about whether a human source believes what he/she reports to us.

Here are some examples of situations in which a human source's veracity could reasonably be questioned. First, consider a source S who is inconsistent in his/her reports. Source S first tells us that event E occurred but later tells us that E did not occur. A source S may tell different things to different persons. Source S tells the FBI that event E occurred, but has also told CIA that E did not occur. In another situation, we have evidence that source S is under the influence of a hostile group and can be exploited by this group. Examining the past records of the behavior of source S, we discover instances of possible untruthfulness in S's behavior. As yet another example, we find S's report of event E inconsistent with what a different source T has said: this source T reports that event E did not occur. Of course, we would be just as concerned about the veracity of source T as we are about the veracity of source S.

We now consider situations in which are either doubtful concerning veracity or are misleading as far as veracity is concerned. First, it is quite common to use the demeanor and bearing of a human source as an indicator of his/her veracity. Is the source agitated or composed while telling us about event E? The idea is that any hint of agitation or discomfort on the source's part is an indicator of lack of veracity. The truth is that demeanor and bearing is actually a very poor indicator of veracity. Very agitated and uncomfortable sources may be entirely truthful and very composed sources being untruthful. Polygraph analyses are also doubtful as indicators of veracity. The results of polygraph analyses are taken seriously by intelligence agencies but are inadmissible in our courts. One argument is that the results of a polygraph examination may say more about the examiner than it does about the person being examined. Finally, the word untruth is often misapplied. Here is a source S who faces a critic who says the following: "This source S is not being truthful in reporting that event E occurred, since he has such poor eyesight and made only a brief observation under a low level of ambient illumination." This evidence concerns S's observational capacities and not his veracity. What this evidence shows is that S may only have been mistaken since S truly believed that event E occurred.

85. The objectivity attribute of the credibility of a human source is widely overlooked in spite of its importance. What is odd is that lack of objectivity is much discussed in common discourse. We so often hear that we all from time to time believe things because we want to believe them in spite of having little or no evidence for them. One matter of interest concerns the possibility that objectivity is not only a property exclusive to human sensors but also relevant to sensing devices. Provide some examples of the objectivity attribute in various situations.

Answer

Here comes a human source S whose veracity was wrongly questioned. What happened was that source S made a competent observation two weeks ago and told us then that event E occurred. At a hearing today S now tells us that event E did not occur and so we question S's veracity saying that his testimony has been inconsistent. What we say is that S has not gotten his story straight. But S now explains this inconsistency. Here is what S tells us today. S says: "I remember telling you that event E occurred two weeks ago. However, on careful reflection, I have come to the view that I did believe that E occurred simply because I expected to observe E and not because my senses told me that E occurred. My belief now is that E did not occur." Our dictionary defines an objective observer as a person who bases a belief on sensory evidence and not on what this observer wished or expected to observe. Since we have no other evidence of S's lack of veracity we conclude that his inconsistency is based on his admitted lack of objectivity in his observation.

As we noted, objectivity concerns the formation of our beliefs. But our beliefs are elastic and so frequently change over time, often in response to new evidence and later experiences. What is also true is that we often cannot remember what we believed at some time in the past. What is also true is that objectivity involves various biases that can easily plague all of us. In certain situations we may be biased on favor of believing certain events because we want them to occur or expect them to occur.

But there is another species of observer bias that may concern veracity, we may call this bias "testimonial bias" instead of "objectivity bias." Suppose a source S who we believe would tell us that event E occurred regardless of what S truly believed. We might discover, for example, that S thinks the occurrence of event E would be good new to us and that S, believing this, would enjoy being the bearer of good news. It happens that if we believed that S would report the occurrence of event E whether S believed E occurred or did not believe that E occurred, then S's report of event E has no value at all.

It is possible that we can link the attribute objectivity to sensing devices as well as to human observers, as least as far as bias is concerned. It is possible, for example to adjust the gain or amplification of a sensor in some situations to display images of certain events when captured energy is beyond certain thresholds. These thresholds can be adjusted or biased to capture certain events at the expense of capturing other events.

Credibility of Tangible Evidence

86. The most important attribute of the credibility of tangible evidence is its authenticity: Is this evidence what it is claimed to be? Provide some examples of real and demonstrative tangible evidence items that are not authentic.

Answer

Here first are some examples of inauthentic real tangible evidence. In a murder trial a prosecutor shows the jury a bullet he says was taken from the body of the murder victim during a post mortem examination. In reality, this bullet was one test-fired during the trial by the police through the defendant's automatic pistol. In another murder trial the prosecutor shows a blood sample she says was found in defendant's bedroom shortly after a murder. In reality, this blood sample was one taken from the defendant two weeks after the murder during a physical examination of the defendant. In the investigation of a terrorist incident involving a car bombing, a men's jacket was found at the scene and said by the investigators to belong to a major suspect in this incident. In reality, this jacket did not belong to this suspect but to a by-stander who dropped it while fleeing from the bombing scene. In another terrorist incident, a document was circulated, allegedly by a suspect who claimed responsibility for this incident. In reality, this document was circulated by another person hoping to link this suspect to the incident.

Now here are some examples of inauthentic demonstrative tangible evidence. A photo showing two men at the scene of a terrorist incident led investigators to label these two men as suspects since investigators believed this photo was taken on the day of the incident. In reality, this photo was wrongly labeled as to the time it was taken; it was actually taken two weeks before the terrorist incident. A map showing the locations of Taliban units in eastern Afghanistan was taken from a captured Taliban fighter. In reality, this map was contrived to mislead our forces about the locations of these units. A statistical analysis is presented to investigators regarding the frequency of terrorist incidents around the globe. This analysis shows a general decrease in such incidents. In reality, this analysis is faulty since a variety of terrorist incidents have gone unreported at several locations. Finally, a photo of the body of a wanted Al Qaeda leader in Yemen is made available to the press showing the leader's demise. In reality, this photo does depict this leader but was staged to show him being the victim of a drone attack. We are led to believe this leader is dead when he is still quite alive.

87. The leadership in Country T has embarked upon an aggressive track regarding its relationship with neighboring countries. We are presently assessing the capability of Country T to wage war on a country with whom we have very friendly relations. We suspect that policy makers in T are considering the development of a certain tactical weapon system we will call W. If they are successful in developing system W, this would give T a decided advantage in any armed conflict they might have with this friendly country. We presently have a source S, a national of Country T, who is an engineer and an expert on the design of weapon systems such as W. Further, she meets regularly with policy makers in Country T regarding the development of tactical weapon systems. Source S has agreed to inform us about deliberations made by policy makers in T regarding the development of system W. Does what we know so far about S bear on her competence or credibility?

Answer

What we know so far about S bear only upon her competence. All we know so far is that she is an expert who seems to have appropriate contacts with policy makers in T.

88. Source T now reports to us the following information. She says she was just told by a ranking policy maker in Country T that all plans to develop system W have been suspended because it was thought that such development would be far too expensive. What S has told us seems to be good news, but can we believe what she says? What general kinds of evidence should we consider about the credibility of what she has just told us?

Answer

Evidence bearing upon S's credibility including her veracity, objectivity, and observational sensitivity. Though she may have given us reasonably accurate information in the past, we wonder whether or not she is still to be believed. You may recognize that what she has reported to us is information she says she received at second hand from another source. So, we have the credibility of her source of information to consider. We should, of course, compare this evidence from S with other evidence we have about the possibility of system W developments in Country T.

89. The most important attribute of the credibility of tangible evidence is its authenticity: Is this evidence what it is claimed to be? Provide some examples of real and demonstrative tangible

evidence items that are not authentic.

Answer

Here first are some examples of inauthentic real tangible evidence. In a murder trial a prosecutor shows the jury a bullet he says was taken from the body of the murder victim during a post mortem examination. In reality, this bullet was one test-fired during the trial by the police through the defendant's automatic pistol. In another murder trial the prosecutor shows a blood sample she says was found in defendant's bedroom shortly after a murder. In reality, this blood sample was one taken from the defendant two weeks after the murder during a physical examination of the defendant. In the investigation of a terrorist incident involving a car bombing, a men's jacket was found at the scene and said by the investigators to belong to a major suspect in this incident. In reality, this jacket did not belong to this suspect but to a by-stander who dropped it while fleeing from the bombing scene. In another terrorist incident, a document was circulated, allegedly by a suspect who claimed responsibility for this incident. In reality, this document was circulated by another person hoping to link this suspect to the incident.

Now here are some examples of inauthentic demonstrative tangible evidence. A photo showing two men at the scene of a terrorist incident led investigators to label these two men as suspects since investigators believed this photo was taken on the day of the incident. In reality, this photo was wrongly labeled as to the time it was taken; it was actually taken two weeks before the terrorist incident. A map showing the locations of Taliban units in eastern Afghanistan was taken from a captured Taliban fighter. In reality, this map was contrived to mislead our forces about the locations of these units. A statistical analysis is presented to investigators regarding the frequency of terrorist incidents around the globe. This analysis shows a general decrease in such incidents. In reality, this analysis is faulty since a variety of terrorist incidents have gone unreported at several locations. Finally, a photo of the body of a wanted Al Qaeda leader in Yemen is made available to the press showing the leader's demise. In reality, this photo does depict this leader but was staged to show him being the victim of a drone attack. We are led to believe this leader is dead when he is still quite alive.

90. Give some examples from your own experience when you have heard people providing information about which they hedge or equivocate.

Answer

You should be able to think of many examples. How many times have you heard your spouse say, "I don't know," or "I don't remember" in response to a question you have asked? How many times have you said, "I don't know" or "I don't remember" in response to a question your spouse has asked you?

91. What inferences might we draw from Omar al-Massari's refusal to provide us with his laptop

computer?

Answer
We are entitled here to infer that Omar al-Massari has information stored on his laptop about his and others' activities that he does not wish us to see.

92. Can you provide some examples of mixtures of evidence from your own experience?

Answer
The easiest ones perhaps are tangible documents you have received that make reference to other tangible documents, or conversations you have had with persons who tell you about tangible objects they have seen.

Credibility of Chains of Custody

93. Human source Y reports to us that the morale among combat troops in Country B is at an all-time low. We ask Y to give us some specifics. He then reports seeing a classified document at a military installation in B that describes the increasing rate of defections and AWOL (Absent Without Official Leave) over the past year. What kind of evidence is this and how should it be analyzed?

Answer
We recognize this as secondhand (or hearsay) testimonial evidence since it comes to us through a chain of sources. We have Y's believability, the authenticity of the document, and the believability of the persons writing this document (if they can be identified) to consider before we believe this evidence.

94. It has been noted for years, and by many persons, that intelligence analysts are hampered by not being the persons who assess the believability or credibility of much of their evidence; this is particularly true of HUMINT evidence. In many cases, sources of HUMINT are under deep cover and their identities are not revealed to analysts. In addition, evidence bearing on the competence and credibility of these HUMINT sources is not always made available to intelligence analysts who will use this HUMINT evidence. Show how this competence and credibility burden on intelligence analysts is made so much heavier when we consider the chains of custody discussed in this chapter.

Answer

To illustrate the problems introduced in this question, consider the situation facing analyst Clyde concerning the chain of custody shown in Figure 41. In this case, Clyde is given an item of HUMINT provided by a source Clyde only knows as Wallflower. Clyde has had no contact with Wallflower, who is someplace in Iran. All Clyde knows about Wallflower comes from Case Officer Bob and, perhaps, Reports Officer Marsha. Suppose all Clyde is told about Wallflower’s competence and credibility is that Wallflower is a “usually reliable source.” It is highly likely that neither Bob nor Marsha would reveal their sources of evidence regarding Wallflower’s competence and credibility. It is very likely that these sources are also under deep cover.

Now, Wallflower is an Iran national who only speaks Farsi. Case Officer Bob is limited in speaking and understanding Farsi. Clyde may have had no contact with Bob or Marsha and may not even know where they are located, presumably in Iraq but not in Iran. How Wallflower communicates with Bob is not revealed. Now, Wallflower’s report concerning the Iraqi Emir Z. is translated by a person who may or may not be known to Clyde. We have given this translator a name, Husam A., that Clyde may know or not. Now, Husam A. does not send his translation of Wallflower’s report to Clyde but to Reports Officer Marsha. As far as Husam A.’s competence as a translator, and credibility as a person, Clyde can only rely upon Marsha. How much Marsha knows about Husam A. is not obvious to Clyde. She may have had to draw upon several other persons to assess the competence and credibility of Husam A. It is even possible that Marsha does not speak Farsi and has to rely upon other persons to assess how good has been Husam A.’s translation of Wallflower’s report. An interesting and troublesome point emerges here. It may simply not be possible to list all the persons who have in fact been links in the chain of custody of Wallflower’s original report to Bob.

Now, as a Reports Officer, Marsha is authorized to edit incoming HUMINT reports and comment on their believability. Marsha passes her edited version of Wallflower’s report to Clyde. How much analyst Clyde knows about Marsha’s competence and credibility is an important matter. Unless Clyde has Husam’s translation of Wallflower’s report, he cannot judge the merits of Marsha’s edited version of it, and whether she altered or omitted anything in Husam A.’s translation.

In this answer we have mentioned only the human links in chain of custody of Wallflower’s report to Bob. But there are actually two other links in this chain: the Sony recording device Bob used to record Wallflower’s report and the SN247 transmission system Marsha used to send to Clyde her edited version of Husam A.’s translation of Wallflower’s report. The accuracy and reliability of these systems are also potential sources of doubt since neither system is probably error-free.

So, what it comes to is that analyst Clyde is hardly in a position to assess Wallfower’s credibility attributes: veracity, objectivity, and sensory sensitivity. And Clyde is also limited in his ability to assess Wallflower’s competence attributes: access and understanding. The basic trouble is that, even if Clyde knows to ask relevant questions about Wallfower’s credibility and competence, he may not receive any answers from persons he may know about in the chain of custody of Wallflower’s report. As we noted, there may be other persons in this chain that Clyde will never know about. In summary, Clyde will experience doubts initiated by Wallflower, but he may experience even more doubts initiated by others involved in the chain of custody of Wallflower's report.

95. Intelligence analysts may choose to ignore the heavy burden mentioned in Question 117. Analysts might prefer to accept versions of HUMINT reports they receive without questioning anything about chains of custody of these reports. Show some of the consequences of failure to assess possible sources of doubt that may lurk in a chain of custody.

Answer

To begin to answer this question we must consider common recurrent constraints facing analysts such as Clyde. The first involves time. What we do know is that Clyde will not have an unlimited amount of time to dwell upon the authenticity of the report he has received from Marsha concerning Wallflower's report. First, Clyde may be working on other important inferential tasks besides the one involving Emir Z. Second, Clyde may be pressured to report his conclusions regarding Emir Z. in a short time. But even in the absence of such temporal constraints, Clyde faces a virtually impossible task of considering all the sources of doubt he knows about in the chain of custody regarding Wallflower's report; but there are sources of doubt in this chain Clyde does even not know about.

So, what is Clyde to do with the translated and edited report he has received from Reports Officer Marsha? One possibility is that Clyde chooses to ignore entirely what he may know about the chain of custody and focus entirely only on what Marsha says about Wallflower's believability and other matters. Here is the exact message Clyde has received from Marsha on 22 September:

"On 18 September we received the following report from a usually reliable Iranian source, code-named *Wallflower*. Wallflower says that on 13 September he observed an Iraqi man he knows as Emir Z. leaving a building at 231 Palm Street in Ahwaz, Iran at 2:30 PM. The first three floors of this building are occupied by the IRGC (Iranian Revolutionary Guards Corps). Wallflower says that this man hurried away from this building and disappeared in a crowd. We know that Wallflower would recognize Emir Z. since Wallflower stayed in Emir Z.'s home in Iraq in April, 2002."

We now recall the inferential objectives of analyst Clyde regarding Emir Z. Clyde suspects that Emir Z. is not the respected Iraqi government official he pretends to be but is actually collaborating with Iran and its universally despised IRGC. Clyde certainly finds this report from Wallflower relevant and interesting and so he accepts it entirely. He is especially interested about where Wallflower says he observed Emir Z. Here are some consequences that can occur as a result of Clyde's ignoring the chain of custody of Wallflower's report about Emir Z.

First, what Clyde does not know is that the translator Husam A. made an error in translating the date on which Wallflower said he observed Emir Z. in Ahwaz, Iran. Actually, Wallflower said he made this observation of Emir Z on 8 September and not on 13 September. This error was not caught by Reports Officer Marsha. This would certainly cause Clyde troubles in assessing the value of Wallflower's report since it is known for sure that Emir Z. was in Baghdad all day on 8 September. Second, it was Marsha who inserted the information about the IRGC occupying the first three floors of the building Wallflower said he saw Emir Z. leaving. What Marsha apparently does not know is that the IRGC vacated all its offices from this building four months ago in May. These revelations, if noted, would make Wallfower's evidence virtually valueless to analyst Clyde. In fact, it seems that Wallflower fabricated this report about Emir Z.

96. One thing analysts are trained to do is to assess the consistency of one item of evidence with other items of evidence they may also have. Show how even this consistency assessment is affected by ignoring chains of custody.

Answer

Now, suppose we temporarily forget about these possible chain of custody difficulties regarding Wallflower's testimonial evidence. Suppose Clyde is quite confident that Emir Z. was in Ahwaz, Iran coming out of a building that houses IRGC offices. But we suppose Clyde is a careful analyst who would investigate the consistency of Wallflower's report with other evidence Clyde presently has. There could be three forms of consistent evidence, the first two concern corroborative evidence. We might have evidence from a source who provides evidence favorable to Wallflower's veracity, objectivity, and observational sensitivity. Or, we might have evidence from another source who provides the same report as Wallflower did about seeing Emir Z. But we have no such corroborative evidence.

What we do have in our example is convergent evidence; this is evidence about other events that favor the same hypothesis as the evidence whose consistency is at issue. An important consideration about convergent evidence is that it may also be synergistic in enhancing the value of the evidence whose consistency is at issue. In our example we have provided Clyde with some tangible evidence that apparently converges with Wallflower's testimonial evidence concerning Emir Z.'s Iranian connections. This evidence comes in the form of a photo taken by another source, code-named Stovepipe. This photo allegedly shows Emir Z. at an IRGC base outside Dezful, Iran. Here is a description of this demonstrative tangible evidence Clyde received on 1 September. The evidence was transmitted to Clyde from a photo interpreter named Mike.

"1 September. Here is my interpretation of the photo we received from Stovepipe on 28 August. Stovepipe says he took this photo (attached) eight days ago (on 20 August). I can first verify that this photo was taken about three miles outside Dezful, Iran at what we believe is an IRGC base. The photo shows three persons emerging from a building on this base. The person on the left I could not identify, the person in the middle, I believe to be Major Omar P, an IRGC officer. The person on the right I believe to be an Iraqi named Emir Z. One problem is that, as you can see in this photo, this person was not directly facing the camera; we only have a side view of his face. Our facial recognition software provides a modest degree of confidence in this identification. I myself have interpreted other photos taken of Iraqi government officials. I am reasonably confident that this person is Emir Z."

Figure 42 shows the chain of custody concerning Stovepipe's photo evidence. We have shown only two human links in this chain: Case Officer Bob and Photo Interpreter Mike, but there might be others. There are two other device links: a camera and a transmission system. You can observe the credibility and competence attributes involved in this chain of custody that analyst Clyde may know very little about. Suppose Clyde completely accepts Mike's interpretation and believes it is reasonably probable that Emir Z. was in Dezful, Iraq coming out of a building at the IRDC base. Clyde may think to himself: "Wow, this photo evidence says the same thing as Wallflower's report that Emir Z. was coming out of an IRDC building in Ahwaz, Iran. Together, these two items of evidence converge to make me extremely confident that Emir Z. is collaborating with the Iranian IRDC. I am going to suggest to Iraqi officials that they should fire Emir Z. and place him under custody."

But what analyst Clyde does not know is that Mike's photo interpretations have been matters of concern for quite some time. Records show that he has made several interpretations that have been incorrect and that have led us to take actions that had unfortunate consequences. Now, let us reconsider items that were also unknown to Clyde when he assessed Wallflower's report that he learned from Reports Officer Marsha. Recall the evidence not taken account of by Clyde, that Emir Z. was not even in Ahwaz, Iran when Wallflower said he was, and that the IRDC had vacated the building four months ago that Emir Z. was allegedly exiting.

It may be possible to say that analyst Clyde brought these troubles on himself by not questioning competence and credibility attributes of links in the chain of custody of both of the items of evidence he has considered. However, there is no guarantee that Clyde may have obtained answers to his questions when in fact he knew so little about the links in these chains of custody. And, even if Clyde did know which links to question, there is no guarantee either that he would have obtained useful and helpful answers to his questions. What these examples illustrate is the extreme importance of considering chains of custody of intelligence information.

Methods of Assesing Uncertainty

Subjective Bayesian View

97. As we noted, the subjective Bayesian view of probability lets us assess probabilities for

> singular, unique, or one-of-a-kind events, provided that our assessed probabilities obey the three Kolmogorov axioms we discussed above regarding enumerative probabilities. First, is there any way of showing that these axioms for enumerative probabilities also form the basis for ideal or optimal probability assessments in the non-enumerative case? Second, can this really be the rational basis for all probability assessments based on evidence?

Answer

Here is one argument that subjective probabilities are ideal or optimal only when they obey the Kolmogorov axioms. This argument involves consideration of what are termed "Dutch books." A Dutch book is a combination of wagers you might accept that guarantees you will lose no matter what happens. Here is a non-enumerative event: the winner of the next Superbowl. Just before the next Superbowl, you and another person $\mathcal{P}$ are discussing possible outcomes involving teams X and Y. As we know, there is always a winner in the Superbowl; there are no ties. Suppose your team is X and $\mathcal{P}$'s team is Y. Suppose $\mathcal{P}$ offers to pay you \$3 if X wins; i.e. $\mathcal{P}$ is betting on Y to win. The main issue you now face is: how much should you pay $\mathcal{P}$ to buy into this wager? Clearly, this involves your view of the probability of X or Y winning.

In discussing your probabilities for the outcome, suppose $\mathcal{P}$ hears you say, "I think it is just as likely that X wins or Y wins and so I will say that P(X wins) = 2/3 and P(X loses) = 2/3." This is a clear violation of Kolmogorov's axioms since P(X wins) + P(X loses) must equal 1.0. Now, a probability of 2/3 corresponds to betting odds of 2/3/1/3 = 2:1, according to Kolmogorov's axioms. What $\mathcal{P}$ is hearing you say is that you would pay \$2 to play each of the wagers implied by your assessed probabilities. In other words, you are betting for and against team X at the same odds 2:1. If you accept both of these wagers you will lose no matter what happens since you will win \$3 on one of these wagers but you have paid \$4 to play them both. So, your net loss is one dollar regardless of which team wins.

Now, it has been shown that such a Dutch book cannot be made on anyone whose probabilities conform to the Kolmogorov axioms. Many persons have argued that this shows that the only basis for rational or coherent probability assessment rests on these axioms. Suppose you recognize the trouble with your initial probability assessment and so you now say P(X wins) = 2/3 and P(X loses) = 1/3 and so you know feel entirely rational. So, you offer person $\mathcal{P}$ \$2 to play the wager he offers and P accepts your offer. You accept $\mathcal{P}$'s wager since you are paying only \$2 for a wager that will get you the \$3 you expect. The Superbowl is played and $\mathcal{P}$'s team Y has won; you have lost your \$2. You say, "How can this have happened since I was so rational in my probability assessment?" Jonathan Cohen whose work on Baconian probability we have discussed, answers your question.

Cohen would tell you that just having coherent or consistent probability assessment is not nearly enough to claim rationality in reasoning based on evidence. Among other things, he says that you will not escape ruin if you continue to bet against persons who are better informed than you are; i.e., persons whose coverage of relevant evidence is much more complete than yours. Here is what happened: person 𝒫's inference was based on several matters that you overlooked and evidence on these matters favored team Y's winning. There is nothing in conventional probability and Bayes' rule that takes account of how complete is the coverage of evidence and considers how many relevant questions are left unanswered. Cohen's Baconian probabilities do account for evidential incompleteness.

98. Show how Bayes' rule supplies no method for incorporating "pure evidence" as does the Belief Function system.

Answer

Dating back to at least the time of the Jacob Bernoulli, a distinction was made between mixed evidence and pure evidence. Mixed evidence is that which has some probability of occurrence under every hypothesis being considered. Pure evidence, on the other hand, says nothing at all about one or more hypotheses being considered. It is evident that Bayes' rule assumes that all evidence is mixed; it cannot capture inferences based on pure evidence; but the Belief Function system can capture inference based on pure evidence. This situation is best illustrated using the scales shown in Figure 46

The middle scale shows Shafer's support (s) or belief (b). What is crucial is that s = 0 and b = 0 simply mean "lack of support" or "lack of belief." These are judgments freely allowed in the Belief Function system. For example, here's an inference involving three hypotheses: H_1, H_2, and H_3. Suppose an analyst 𝒜 assesses the following for some evidence E: $s(H_1) = b(H_1) = 0.2$; $s(H_2) = b(H_2) = 0.3$; and $s(H_3) = b(H_3) = 0$. What this assignment says is that analyst 𝒜 judges evidence E to offer some support to H_1 and H_2, but E offers no support at all for H_3, which says that the analyst's $b(H_3) = 0$ means a lack of belief in H_3 and not a disbelief in H_3. Now, this lack of belief in H_3 does not make it completely dead in the Belief Function system since support may come from further evidence and some belief in H_3 is restored.

Now, let's try to capture this situation using Bayes' rule. As we have shown in the top scale in Figure 46, the conventional probability P = 0 means impossibility or disbelief. Here is another analyst ℬ who attempts to capture this pure evidence situation using Bayes' rule. As we all know, the force or weight of evidence is captured in Bayes' rule by likelihoods. Analyst 𝓑 says, "I have three of them to consider for evidence E: $P(E|H_1)$, $P(E|H_2)$, and $P(E|H_3)$. To make my friend Analyst 𝒜 happy I will use the same values 𝒜 did in assessing the force of evidence E. My friend 𝒜 agrees with Shafer in saying that evidential support and evidential weight are the same. So, here are my likelihoods for evidence E: $P(E|H_1) = 0.2$; $P(E|H_2) = 0.3$; and $P(E|H_3) = 0$. Well, I see what my problem is: When I set $P(E|H_3) = 0$, Bayes' rule says that evidence E makes H_3 impossible and that H_3 will forever have zero probability regardless of whatever further evidence I obtain. I cannot resuscitate H_3 as my friend $\mathcal{A}$ can do using Belief Functions. I now see the difference between lack of belief and disbelief. This distinction is crucial in capturing pure evidence."

Belief Functions

99. Consider the following support assignments:

	$\{H_1\}$	$\{H_1^C\}$	$\{H_1, H_1^C\}$	Ø
Normalized $\mathbf{s}_1 \oplus \mathbf{s}_2$	0.23	0.63	0.14	0
S_3	0.6	0.1	0.3	0

Apply the Dempster's rule and determine what our analyst's support for the new orthogonal sum $\mathbf{s}_1 \oplus \mathbf{s}_2 \oplus \mathbf{s}_3$ should be.

Answer

If you follow the steps for using Dempster's rule and do the arithmetic correctly, you should end up with the following orthogonal sum:

	$\{H_1\}$	$\{H_1^C\}$	$\{H_1, H_1^C\}$	Ø
Normalized $\mathbf{s}_1 \oplus \mathbf{s}_2 \oplus \mathbf{s}_2$	0.23	0.63	0.14	0

Fuzzy Probability

100. Provide an example showing how an analyst's numerical assessment of a probability applied to a conclusion can invite criticism.

Answer

Here is analyst $\mathcal{A}$ whose inferential task has been to report her conclusion about which one of four competing or rival hypotheses is true: H_1, H_2, H_3, and H_4. This analysis has involved a mass of evidence of different types: some HUMINT, and some tangible evidence such as IMINT, SIGINT, and MASINT. Analyst $\mathcal{A}$ might have been assisted by other analysts who are experts in various areas relevant in this analysis. But it is analyst $\mathcal{A}$ who bears responsibility for drawing a final conclusion in this important inferential assignment. Analyst $\mathcal{A}$ is required to provide her conclusion at a hearing involving a group of policy-makers having vital interests in the conclusion $\mathcal{A}$ will report.

Analyst $\mathcal{A}$ begins her report by reviewing the four hypotheses she has entertained in the matter of interest and the lines of evidence and arguments she has considered that bear upon these hypotheses. She concludes her report by the following comment. She says, "My judgment is that there is an 83% chance that H_2 is true; as you have seen, several lines of evidence converge in favoring H_2."

Hearing $\mathcal{A}$'s conclusion, the lead policy-maker $\mathcal{PM}_1$ says: "I agree that the evidence points most strongly to H_2 but what troubles me is the precise probability you have used to hedge your conclusion. We know that you have not made 100 observations and discovered that the event described in H_2 occurred on exactly 83 of these occasions. All the events in your hypotheses are unique and have never happened before. So, where did your precision come from? You have mentioned the array of subjective judgments you and your team of analysts had to make."

Analyst $\mathcal{A}$ replies: "Sir, you are correct that my conclusion was based on a variety of subjective judgments. The number 83% that I gave you is just my best estimate of this probability of H_2; it could be anywhere between 75% and 87%."

Another policy-maker $\mathcal{PM}_2$ then replies: "But you now give us additional precise numbers showing the upper and lower limits of your assessed probabilities. Again, where did this precision come from? This would still be based on a variety of imprecise judgments."

Analyst $\mathcal{A}$ then says: "I agree that these upper and lower probability estimates are imprecise; the lower limit could be anywhere between 70% and 76% and the upper limit could be anywhere between 85% and 89%."

Hearing this response by $\mathcal{A}$, the group of policy-makers are silent. Hearing no comments, $\mathcal{A}$ says: "The truth is that I expected you to require me to hedge my conclusion using numbers and would criticize me for hedging my conclusion in words. The second truth is that what I wished to report is that my analysis shows that H_2 is quite probable and greater than the probability of any of the other hypotheses."

In response to $\mathcal{A}$'s final comment, the lead policy-maker $\mathcal{PM}_1$ says: "We actually are quite satisfied with your verbal assessment of the probability of H_2. We caused you lots of trouble trying to defend the numerical assessments you made trying to justify these numbers. We agree that you have done a very fine analysis and accept your conclusion that H_2 is quite probable and the most likely of any of the hypotheses you have entertained. I am going to recommend that we take a course of action consistent with H_2 being true. We thank you for your very helpful analysis."

The main thing this example shows is that precise numerical probability hedges on an analytic conclusion cannot be defended because of the rampant imprecision of the judgments necessary in intelligence analysis. One virtue of probability assessments made in words is that they do not require commitments to any particular numbers when such commitments cannot be justified.

Complementarity of the Prbability View

101. Think back to the very first time you were ever tutored about probability, what it means, and how it is determined. What were you told about these matters? Then, describe your present views about these probability matters.

Answer

If you are like most of us, the very first things you were told about probabilities is that they are numbers between zero and one (inclusive) that indicate how certain we are that some event has occurred. At the end points of this interval, zero probability means that the event in question is impossible or cannot occur. At the other end of this interval, a probability of one means that the event in question is certain to occur. Numbers in between zero and one indicate gradations in our certainty that an event has occurred. Further, you were told that if two events are mutually exclusive, then we simply add together their separate probabilities to determine whether one or the other of these events has occurred. In short, what you and the rest of us were initially given was a verbal account of the Kolmogorov axioms. More than likely, you were first given examples of these axioms in aleatory situations involving games of chance. Further examples would then follow involving relative frequencies and conventional statistics; not Bayesian statistics mind you, since the use of Bayes' rule was usually regarded as inapplicable.

After these initial introductions to probability, you may have had any number of subsequent courses on probability and statistics in which you certainly appreciated the richness of what you had learned in a very wide array of applications in science and engineering, most of which assume repeatable or replicable situations. It is entirely possible that you, like so many others, regarded what you have learned and applied as **THE** theory of probability, meaning that there is no other possible theory of probability. Until you read this chapter in this book, you may have retained a belief that there is only one system of probability and that this system applies in all situations in which we need to express and combine our uncertainties.

But, as we have explained, the conventional theory of probability applies to enumerative situations when we can count the occurrence of events. However, we have to ask whether this conventional view of probability applies when we have no events to count. This will happen in any situation in which the events of concern to us are singular, unique, or one-of-a-kind. Such situations abound in intelligence analysis, law, history, and many other situations. Here is one example from the field of law.

You may be familiar with the trial of O. J. Simpson, the American football star who was accused of first-degree murder in the slaying of his wife Nicole Brown Simpson. O.J. Simpson was acquitted of this charge but opinions still differ about whether he was actually guilty of this charge. You may have some opinion about the probability that O.J. did kill Nicole. However, one thing you could never do is to play the world over again 100 or 1000 times to see how many times O.J. did it. He either did or did not do it on exactly one occasion. To reason in these non-enumerative situations in which we have nothing to count, we told you about four different probability systems in this chapter, each of which says some valuable things about reasoning based on evidence, but no one of these systems says all there is to be said. We hope we have given you adequate grounds for believing that there is more than one system of probability to be considered in intelligence analysis. In short, probabilistic reasoning based on evidence is far too rich an intellectual activity to have all of its richness captured by any single probability theory we might devise.

Hypothessis Analysis with Wigmorean Argumentations

102. Which of the following is an **assumption** in the argument that John has a higher IQ than Mark:

a) John scored higher than Mark on the SAT.
b) Individuals that score higher on the SAT have a higher IQ.
c) John scored higher than Mark on an IQ test.

Answer

a) is an item of evidence or a hypothesis based on the SAT results of John and Mark.
b) is an assumption based on “commonsense” reasoning about cause and effect. An assumption is a statement taken to be likely true, without having any supporting evidence in the current problem.
c) is an item of evidence, or a hypothesis based on the IQ test results of John and Mark.

103. Possible answers to a question about a situation are considered:

a) assumptions
b) hypotheses
c) items of evidence

Answer

a) *Incorrect:* An assumption is a statement taken to be likely true, without having any supporting evidence.
b) *Correct:* Possible answers to a question about a situation are considered hypotheses.
c) *Incorrect:* Evidence is any item of information that favors or disfavors the truthfulness of a hypothesis.

104. Consider the hypothesis
Mark’s grades have improved.

and the sub-hypotheses

Last year Mark maintained a C average.

Mark is maintaining a B average this year.

How should you represent them in Cogent?

a. As two alternative (separate) arguments
b. As a single AND argument

Answer
The hypothesis requires both sub-hypotheses to be considered as an AND argument. Mark's grades could not have improved unless there was a lower grade level from which his grades could improve. You could not conclude anything about the direction of Mark's grades just from the sub-hypothesis *"Mark is maintaining a B average this year."* Similarly, you could not conclude anything about the direction of Mark's grades just from the sub-hypothesis *"Last year Mark maintained a C average."*

105. True or false:
If "Last year Mark maintained a C average" and "Mark is maintaining a B average this year",
then it is <u>likely</u> that "Mark's grades have improved."

Answer
False: If both sub-hypotheses are true (John had a C average and progressed to a B average), the hypothesis that his grades improved is <u>certain</u>.

106. Consider the hypothesis
Material on the web site of terrorist group X persuaded John to become a terrorist.

and the sub-hypotheses

John did not harbor any pro-terrorist views prior to March.

John visited the terrorist web site 22 times in March.

John offered his services to terrorist group X in April.

How should you represent them in Cogent?

a. As three alternative (separate) arguments.
b. As a single AND argument.

Answer
The hypothesis requires all three sub-hypotheses to be considered as an AND argument. The sub-hypotheses *"John did not harbor any pro-terrorist views prior to March"* and *"John offered his services to terrorist group X in April"* support the notion that John's views on terrorism evolved. The sub-hypothesis *"John visited the terrorist web site 22 times in March"* supports the notion

that the website was the driver of this new perspective on terrorism.

107. Assuming that "John did not harbor any pro-terrorist views prior to March" and "John visited the terrorist web site 22 times in March" and "John offered his services to terrorist group X in April", how certain are you that "Material on the web site of terrorist group X persuaded John to become a terrorist"?

   a. certain
   b. almost certain

Answer

Almost certain is a better answer than certain because we do not know exactly what John thought about the material on the website. It appears the material on the website was highly influential but we cannot be 100% certain. It is possible that he harbored some doubts about the material on the website but was influenced to act by other factors.

108. Consider the hypothesis

   President Doe's intelligence service assassinated dissident James Fairley at
   the behest of Doe

and the sub-hypotheses

   The intelligence service of Doe assassinated Fairley

   The intelligence service would only assassinate Fairley if Doe gave the order

How should you represent them in Cogent?

   a. As two alternative (separate) arguments
   b. As a single AND argument

Answer

As a single AND argument. Both sub-hypotheses need to be true in order to conclude that the hypothesis is true. If either the intelligence service did not assassinate Fairley or the service assassinated Fairley on its own initiative, then we cannot conclude that President Doe's intelligence service assassinated dissident James Fairley at the behest of Doe.

109. True or false: If "The intelligence service of Doe assassinated Fairley" and "The intelligence service would only assassinate Fairley if Doe gave the order" then it is certain that "The intelligence service assassinated dissident James Fairley at the behest of President Doe."

Answer

True: If the intelligence service did in fact assassinate Fairley and would only do so at Doe's direction, then the hypothesis that the intelligence service assassinated Fairley and acted at Doe's behest is 100% certain.

110. True or false: If “The intelligence service of Doe assassinated Fairley” and “Doe maintains very strict control over his intelligence service” then it is certain that “President Doe gave the order to his intelligence service to assassinate dissident James Fairley.”

Answer
False: The sub-hypothesis that *“Doe maintains very strict control over his intelligence service”* is not the same as total or complete control and allows for the possibility that the intelligence service acted on its own initiative. Thus, the relevance of the AND argument would be less than certain.

111. Consider the hypothesis

John conducted the terrorist attack against the government on 1 May

and the sub-hypotheses

John was planning to conduct a terrorist attack against the government on 1 May

John was involved in several attacks against the government last year

How should you represent them in Cogent?

a. As two alternative (separate) arguments
b. As a single AND argument

Answer
As two alternative (separate) arguments. Either of the sub-hypotheses can support the hypothesis. Even if we lacked information that *“John was planning to conduct a terrorist attack against the government on 1 May”,* the information that *“John was involved in several attacks against the government last year”* suggests John could be involved in the 1 May attack.

112. True or false: Given the hypothesis “John conducted the terrorist attack against the government on 1 May”, the sub-hypothesis “John was planning to conduct a terrorist attack against the government on 1 May” should be assessed as having lower relevance than the sub-hypothesis “John was involved in several attacks against the government last year.”

Answer
False: The relevance of *“John was planning to conduct a terrorist attack against the government on 1 May”* should be higher because that sub-hypothesis very specifically connects John to the 1 May attack. In contrast, the information that *“John was involved in several attacks against the government last year”* does not state anything conclusively about what John is doing or planning regarding the 1 May attack.

113. Consider the hypothesis

John, who is in his early 40’s, is capable of embezzling money from his firm and the sub-hypotheses

John was convicted of shoplifting as a teenager

John has a reputation of having little or no integrity among most of his current co-workers

How should you represent them in Cogent?

c. As two alternative (separate) arguments
d. As a single AND argument

Answer
As two alternative (separate) arguments. Either sub-hypothesis can independently support the hypothesis. If one of the sub-hypotheses is not true, the other still provides some basis for concluding that John is capable of embezzling the money.

114. True or false: Given the hypothesis "John is capable of embezzling money", the sub-hypothesis "John was convicted of shoplifting as a teenager" should be assessed as having lower relevance than the sub-hypothesis "John has a reputation of having little or no integrity among most of his current co-workers."

Answer
True. The sub-hypothesis *"John was convicted of shoplifting as a teenager"* should be assessed as having lower relevance. The relevance of this sub-hypothesis speaks to John's moral compass as a youth, not his current moral compass. It is possible that John has turned his life around. In contrast, the sub-hypothesis that *"John has a reputation of having little or no integrity among most of his current co-workers"* speaks to his current moral compass and suggests few of his co-workers trust John to do the right thing.

Analytic Bias

115. An intelligence analysis has miscarried on an important matter concerning national security and a post mortem hearing is now in progress to determine what went wrong. Attention is focused on the work of analyst 𝒜, who provided key judgments during the analysis process. At the hearing a critic notes, "Our main trouble was that we paid too much attention to analyst 𝒜 who gave us a biased assessment of the force of evidence E*. A said this evidence very strongly favored hypothesis H_2, which we now know did not occur. H_4 really happened and we have all been embarrassed since we reported that H_2 was true." What could have happened that led this critic to say that 𝒜 was biased? Who or what determines analytic bias? And, can analytic bias be prevented?

Answer
An episode of intelligence analysis can go wrong for many reasons. On many accounts we have

read, assorted alleged analytic biases are the major reasons why an analysis has gone wrong. In some cases it seems that it is argued that analytic bias is the only reason why an intelligence analysis can go wrong. However, an analysis may go wrong for other reasons not involving bias but rather for an assortment of analytic errors that might be made. What is the distinction between bias and error in intelligence analysis and why is this distinction so important to recognize and discuss?

116. Are there sources of bias that cannot be linked to individual analysts or teams of analysts?

Answer

In discussing bias in intelligence analysis, we cannot overlook what we may refer to as "institutional bias." This is a form of bias generated or originated by entire intelligence offices or agencies. Persons with experience in intelligence analysis will surely have heard the term: "institutional memory." This term seems to have arisen in connection with repeated experiences with an actual or a potential adversary. In such experiences, there have been attempts to capture the thought processes and the decisional or inferential behavior of leaders or decision makers in adversary countries or groups. Such attempts have led to expectations and indeed preferences in believing certain things about what these adversaries will do or not do in the future. In older times, we might have heard assertions such as, "The Soviets would never do X, Y or Z; they would prefer to do W." In more recent times the assertion might read, "The Pakistani Taliban would never take actions S, T or U; they would prefer to do R." This certainly has the positive effect of speeding up the analysis process and will generate plausible conclusions on the future behavior of adversaries. But it also creates the possibility of bias. For example, here is an analyst or team of analysts trying to draw a conclusion about possible Taliban actions in Pakistan based upon recent evidence. Here is what the lead analyst says to his team, "The evidence we have points most strongly to the Taliban doing T. But we better build a stronger case for their doing R, since this is what our bosses will expect the Taliban to do." In Question 139 we said that analyst $\mathcal{A}$ had a preference for believing H_2 but we did not discuss why $\mathcal{A}$ may have had this preference. $\mathcal{A}$'s preference may have been the result of an institutional bias inflicted upon analyst $\mathcal{A}$ and was not a bias $\mathcal{A}$ generated himself.

117. In discussions of bias, so much attention has been based on numerical assessments of the probability of hypotheses considered in intelligence analysis. What other properties of intelligence analysis represent a much more important emphasis in assessing the quality of an analysis?

Answer

The answer to this question is quite easy. The defensibility of an analyst's arguments is so much more important than the numerical probabilities this analyst might use to hedge a conclusion. If the analyst's arguments based on evidence are defective, it should not matter what numbers the analyst uses to hedge a conclusion based on this evidence; these numbers cannot be taken seriously. This is why we have emphasized the view that in probabilistic reasoning, arguments are more important than numbers. As an example, here is analyst $\mathcal{B}$ who says that H_3 is the most likely hypothesis having a probability of 58% of being true. Critic C_1 says, "Well, $\mathcal{B}$ is usually biased, being conservative in his probability assessments, and so we should believe the probability of H_3 is much higher." Critic C_2 says, "I have found many of $\mathcal{B}$'s arguments in this analysis to be triumphs of the non sequitur. I could not take seriously any numbers $\mathcal{B}$ used in his analysis." Critic C_2 then points out a few of these non sequiturs to C_1 and they both agree that $\mathcal{B}$'s analysis cannot be taken seriously.

118. An episode of intelligence analysis can go wrong for many reasons. On many accounts we have read, assorted alleged analytic biases are the major reasons why an analysis has gone wrong. In some cases it seems that it is argued that analytic bias is the only reason why an intelligence analysis can go wrong. However, an analysis may go wrong for other reasons not involving bias but rather for an assortment of analytic errors that might be made. What is the distinction between bias and error in intelligence analysis and why is this distinction so important to recognize and discuss?

Answer

We have attempted to provide an accurate and useful account of analytic biases that holds up in different situations. As we noted, the concept of bias involves subjective judgments made by people with reference to their views, beliefs, and opinions. Among the indicators of bias are such factors as partiality, prejudice, favoritism, one-sidedness, intolerance, and narrowness. But intelligence analyses, as well as analyses in other contexts, can miscarry for many reasons not involving these subjective judgmental matters. Incorrect or unproductive analyses can be rooted in an assortment of definite errors that are simply impossible to enumerate. In so many instances, analytic errors can arise because of analysts' unawareness of an important evidential or inferential complexity that has a profound effect on the correctness of a conclusion based on available evidence. We will mention reasons why analysts are unaware of these complexities. But before we begin to discuss errors due to analysts' innocence of evidential and inferential complexities, it seems advisable to discuss the correctness of an intelligence analysis and how this bears on analytic bias and analytic error.

Intelligence analyses can involve future or past events. It seems fair to say that most intelligence analyses involve predictions about possible future events. But they can also involve possible explanations for past events. Here is an adversarial group G whose activities demand our attention. We may be vitally interested in what actions G might take in some situation in the future. Among the actions we consider that G might take are actions: a_1, a_2, a_3, and a_4. As a result of our evidential analysis we predict that group G will take action a_3. But we may also be interested in trying to explain why group G did take some particular action in the past and so we consider an array of possible explanations: e_1, e_2, e_3, e_4, and e_5. As a result of our evidential analysis we conclude that e_4 is the most likely explanation for group G having taken this action in the past.

We now consider the questions: "How correct or accurate was our prediction that group G will take future action a_3?" and "How correct or accurate was our explanation e_4 of group G's taking a particular past action?" Here comes a depressing truth: in so many instances in intelligence analysis we may never be able to tell whether a prediction or an explanation was correct or accurate. In both cases we may simply lack sources of evidence that could verify whether predicted action a_3 occurred or whether explanation e_4 is correct. And, as we have discussed, the world changes continually and we learn new things all the time. One result of this nonstationarity of the world is that the possibilities we must consider in our predictions or explanations have to be continually revised.

There is an important message here that bears on analytic biases and errors. Waiting to be concerned about bias and error until we find out whether a prediction or explanation is correct or not is not a good strategy. This would be an after-the-fact strategy that we have already criticized; we may never know what the "facts" are in any situation. The time to be concerned about bias and error is when an intelligence analysis is in progress and before conclusions are reached and reported.

Here are some comments about possible evidential and inferential errors. We notice right away that the errors we mention are not made just by intelligence analysts but by anyone trying to draw defensible and persuasive conclusions from masses of evidence of different sorts and coming from a variety of different sources. As we have noted, evidence has three important credentials: relevance, credibility, and inferential force or weight. So many persons are innocent of the requisites for defending the relevance of items and collections of evidence by arguments. So many errors occur because of innocence of what is involved in assessing the credibility of evidence of different types and combinations. And, there is considerable innocence regarding the alternative view of probability necessary to capture the incompleteness, inconclusiveness, ambiguity, dissonance, and imperfect credibility of evidence. A major objective of this book has been to provide an effective tutorial on these matters for intelligence analysts.

119. Can analysts ever be criticized for having drawn incorrect conclusions? Or, are some alleged "intelligence failures" actually failures after all?

Answer

The answer to this question involves matters associated with a nonstationary world and some recent developments in information technology. Here comes an example that illustrates problems that analysts have experienced on some occasions. One thing this example illustrates is that the real or potential adversaries we face are not mindless and can react in various ways to our analyses and actions.

Suppose a team of analysts have been following the actions of a known terrorist group in a certain part of the world. The actions of this group indicate that this group is planning a major terrorist action in the near future. The analysts form three hypotheses about when and where this action will occur; the hypotheses are: H_1, H_2, and H_3. As a result of a careful analysis of current evidence, the analysts predict that H_3 will occur in a short time. Police and military forces in the location specified by H_3 are alerted. The time approaches and the analysts are distressed to learn that there has been a very violent terrorist action involving the loss of many lives and the destruction of several large buildings housing many occupants. The trouble is that this action occurred not in the location specified by H_3, but in the location specified by H_1. This violent and very costly terrorist action is widely reported in the press. Press accounts record the fact that this terrorist group was being monitored by an intelligence agency and the result was labeled as an "intelligence failure", a label commonly applied when thing go wrong. The analyst team that predicted the terrorist action at the location and time specified by H_3 are called on the carpet to see where they went wrong in this intelligence failure. But was this actually an intelligence failure?

Suppose the following events happened. The terrorist group had every intention of launching the action at the time and location correctly predicted by the analysts' hypothesis H_3. But at the very last minute, the terrorists changed their minds and decided to attack at the location specified by H_1. What could have happened in this alleged "intelligence failure"? First, the terrorist group may have noticed the alerting of the police and military forces in the location in H_3. This action on our part may have caused the terrorists to vary their intentions. Second, the intended or unintended leakage of information is a common occurrence, even regarding intelligence matters. One possibility here is that the terrorists used classified intelligence information criminally stolen from intelligence agencies to tap into ongoing intelligence activities. A third possibility is that terrorists, like others, simply change their minds for a variety of reasons. Perhaps the terrorist group suddenly thought that the location under H_1 would be a more promising target than the location under H_3.

What this example shows is that intelligence analyses are always exceedingly difficult in an ever-changing world, and that useful and correct assessments of the quality of analyses are no less difficult. The analysis of "intelligence failures" is itself an exceedingly difficult task when we learn new things all the time.

Anticipatory Intelligence with Cogent

Types of Evidence

120. What type of evidence item is E9 MDDOT Record in Table 21 on page 226?

Answer

Tangile: The Maryland DOT records, in the form of a tangible document, could be given to the analyst to verify that the vehicle carrying MD license plate number MDC-578 is registered in the name of the TRUXINC Company in Silver Spring, MD.

121. What type of evidence item is E8 GuardReport in Table 21?

Answer

Tangile: Here we have a document in the form of a log showing when the truck bearing license plate number MDC-578 exited the STEMEQ parking lot at 8:30 PM on the day in question. This tangible item could also be made available to analysts investigating this matter.

122. In new evidence regarding the dirty bomb example, suppose we have a source code-named “Yasmin.” She tells us that she knew a man in Saudi Arabia named Omar al-Massari. Yasmin says she is “quite sure” that Omar spent two years “somewhere” in Afghanistan “sometime” in the years 1998-2000. What type of evidenc is this?

Answer

Equivocal testimonial evidence by Yasmin who says she is “quite sure” that Omar spent two years “somewhere” in Afghanistan “sometime” in the years 1998-2000.

123. We return to our asset “Yasmin” who has given us further evidence about Omar al-Massari in our cesium-137 example. Suppose we have a tangible document recording Yasmin’s account of her past experience with Omar al-Massari. In this document Yasmin tells us about having seen a document detailing plans for constructing weapons of various sorts that was in Omar al-Massari’s possession. What kind of evidence is this and how should it be analyzed?

Answer

This is a mixture of tangible evidence (the document recording Yasmin’s account of her past experience with Omar al-Massari) and testimonial evidence (Yasmin’s actual account). As far as analyisis is concerned, we first have the authenticity of the transcription of her testimony to consider. If Yasmin speaks only in Arabic we should also wonder how adequate the translation of her testimony has been. Also, we have concerns about Yasmin’s credibility to consider in her recorded testimony. Finally, we have further interest in the authenticity of the document she allegedly saw in Omar al-Massari's possession.

124. Consider our discussion on the cesium-137 canister. Upon further investigation we identify the person who rented the truck as Omar al-Massari, alias Omer Riley. We tell him that we wish to see his laptop computer. We are, of course, interested in what it might reveal about the terrorists he may be associating with. He refuses to tell us where it is. What kind of evidence is this?

Answer
We refer to this as the non-production of evidence, a type of missing evidence.

125. What type of evidence item is E6 Clyde in Table 21?

Answer
E6 Clyde is unequivocal testimonial evidence. It represents positive evidence.

126. What type of evidence item is E14 Grace in Table 21?

Answer
This is Grace's unequivocal testimony that no one at the STEMEQ Company had checked out the canister for work on any project. Moreover, she has given negative evidence saying the cesium-137 was not being used by the STEMEQ Company. This negative evidence is very important because it strengthens our inference that the cesium-137 canister was stolen.

Recurrent Substance-Blind Combinations of Evidence

127. Convergent evidence involves evidence about different events, all of which point to the same conclusion. Look at evidence items from Table 21 and identify convergent evidence.

Answer
E9 MDDOT Record, E10 TRUXINC Record1, E11 TRUXINC Record2 and E12-Silver Spring Record, all point toward the conclusion that Omar al-Massari was the person who rented the truck from the TRUXINC Company in Silver Spring, MD.

128. Can you make up some examples of evidence that corroborates other evidence in our dirty bomb scenario? Ask yourself what items of evidence we now have that you would like to see corroborated.

Answer
Our asset, Santa, has told us some very valuable evidence regarding Omar al-Massari. He is obviously an informant or plant inside (what we believe to be) terrorist groups in the area. To be on the safe side, we would like corroborating evidence from other sources concerning what Santa

has told us. Additionally, at the risk of being thought paranoid, we might seek corroboration on what John Walsh has told us. Perhaps he is seeking to protect the interests of a valuable employee. Then of course we have Willard, whose report started this whole thing off. Can we be sure that Willard himself was not involved in some way with the missing cesium-137 canister?

Major Sources of Uncertainty in Masses of Evidence

129. One of the unrealistic features about our cesium-137 example is that all the evidence we have so far is harmonious in pointing toward the hypothesis that a dirty bomb containing cesium-137 will be set off somewhere in Washington, DC. In short, we have no contradictory or divergent evidence so far. Could you imagine what some items of dissonant evidence might be?

130. Answer

Consider E14 Grace from Table 21 (p. 226), Grace's evidence that no one at the STEMQ Company was using cesium-137 on a current project. We might have a source who contradicts Grace. We might have evidence that the STEMQ Company has had radioactive materials stolen in the past by persons working for competitors; this would be divergent evidence on E14 Grace because it would point to a different hypothesis.

131. Consider the evidence provided by John Walsh, the President of the Ultratech company in Silver Spring, MD, at which Omar al-Massari (alias Omer Riley) works (i.e., E020-Walsh in Table 21). Walsh tells us that Omer Riley's job does not require him to handle any radioactive materials such as cesium-137. But suppose we interview another executive at the Ultratech Company, Dan Moore, who is more directly familiar with the work Omer Riley actually performs. This person tells us that Omer Riley has worked recently on devices for measuring soil dampness gradients and this work does involve the use of radioactive materials. What kind of evidence combination do we have here?

Answer

Contradictory evidence because Moore contradicts what Walsh tells us and might be used to reduce the force of E020-Walsh concerning the traces of cesium-137 on Omar al-Massari's body when he was interviewed.

132. Recall the evidence item E17 Santa Adr in Table 23, where Santa is telling us that Omar al-Massari lives with two other males at 403 Winston Road in Silver Spring, MD. We interview a person named Martha, who says she saw a person she knows as Omer Riley (the alias Omar al-Massari uses) park a panel truck in the driveway of the house Omer shares with two other males. Then we have another neighbor, Paul, who tells us that he saw someone, who looked like Omer Riley, park a panel truck in the driveway of the house at 403 Winston Road in Silver Spring, MD. What kind of evidence combination do we have in this case?

Answer

Martha’s report would be additional evidence bearing on the event that the terrorist organization (to which Omar al-Massari belongs) will use, or has used, the stolen cesium-137 to build a dirty bomb. But Martha and Paul report different events: Martha says it was Omer Riley, and Paul says it was someone who looked like Omer Riley. First, suppose we consider Martha to be perfectly believable when she says that it was Omer Riley who parked the panel truck in the driveway at 403 Winston Road. In this case, what Paul tells us would be inferentially valueless. If it was Omer Riley (alias of Omar al-Massari) who parked the truck in the driveway, then it follows necessarily that someone who looked like Omer Riley parked the truck. However, Paul’s report becomes important to the extent to which we believe Martha not to be believable.

133. Look again at I17 Yasmin in Table 24 (p. 234) in which we have “Yasmin” telling us that she knew a man in Saudi Arabia named Omar al-Massari. Yasmin says she is “quite sure” that Omar spent two years “somewhere” in Afghanistan “sometime” in the years 1998-2000. Why is this information ambiguous?

Answer

Because, in addition to Yasmin hedging her assessment by saying she is “quite sure,” she is additionally vague or imprecise regarding where and when Omar al-Massari was in Afghanistan. If we knew where and when Omar al-Massari was in Afghanistan we could make a better assessment of what he was doing there and whether he was actually involved in jihadist activities.

134. For what other evidence items in Table 23(p. 232) would you wish to have other sources available which could help in deciding whether or not to believe what the items tell us?

Answer

We would certainly like to have more evidence concerning the circles Santa travels in, if it is available. We would obviously pump Santa for all the information he can give about this dirty bomb situation.

135. Show why all the evidence in Table 21 (p. 226) and Table 23 (p. 232) is inconclusive.

Answer

To see that all the evidence items in Table 21 and Table 23 are inconclusive, just look at the sources of doubt separating these items from what we are trying to prove from them.

136. Give an example of cumulative redundance.

Answer

Cumulative redundance involves evidence about different events. We already have evidence of minute cesium-137 traces on the hair and skin of Omar al-Massari, but if we now get evidence of traces of cesium-137 on his clothes we would be inclined to say, “So what?” Finding such traces on his body makes finding traces on his clothes pretty likely and would tell us little we don’t already believe, namely, that Omar al-Massari was exposed to cesium-137.

137. Table 21 presents additional items of evidence related to the missing cesium-137 scenario. What examples of unequivocal testimonial evidence do you see in this table?

Answer

E020-Walsh is one example, because John Walsh does not equivocate in saying that Omar al-Massari’s work does not involve handling radioactive substances.

E015-SantaAlias, E014-SantaWork, E014-SantaAdr, and E014-SantaTerOrg are also items of unequivocal testimonial evidence because Santa does not equivocate in what he is telling us about Omar al-Massari.

138. Do you see any example of mixtures of evidence in Table 21?

Answer

E008-GuardReport is tangible evidence about testimonial evidence, because it is a record of Sam’s testimony.

139. Can you provide other examples of mixtures of evidence from your own experience?

Answer

The easiest ones perhaps are tangible documents you have received that make reference to other tangible documents, or conversations you have had with persons who tell you about tangible objects they have seen.

140. What other items of evidence are missing so far in our discussion of the cesium-137 case?

Answer

What has happened to the cesium-137 canister? Where is it and who has it? How did Omar al-Massari obtain the canister from the XYZ warehouse? Was he assisted by someone at the XYZ company? Did anyone at the XYZ company see Omar al-Massari the day the cesium-137 canister was stolen?

141. What items of tangible evidence do you see in Table 21?

Answer

E006-Clyde and E008-GuardReport are both tangible evidence about testimonial evidence. E007-Camera is demonstrative tangible evidence. E009-MDDOTRecord, E010-TRUXINCRecord1, E011-TRUXINCRecord2, and E012-SilverSpringRecord are all items of real tangible evidence. E013-

InvestigativeRecord is demonstrative tangible evidence.

Additionally, Grace could supply documents showing the absence of any use of the cesium-137 in any on-going XYZ project, and those documents would be tangible evidence as well.

142. Table 21presents additional items of evidence related to the missing cesium-137 scenario. What examples of unequivocal testimonial evidence do you see in this table?

Answer

E020-Walsh is one example, because John Walsh does not equivocate in saying that Omar al-Massari's work does not involve handling radioactive substances.

E015-SantaAlias, E014-SantaWork, E014-SantaAdr, and E014-SantaTerOrg are also items of unequivocal testimonial evidence because Santa does not equivocate in what he is telling us about Omar al-Massari.

143. What other items of evidence are missing so far in our discussion of the cesium-137 case?

Answer

What has happened to the cesium-137 canister? Where is it and who has it? How did Omar al-Massari obtain the canister from the XYZ warehouse? Was he assisted by someone at the XYZ company? Did anyone at the XYZ company see Omar al-Massari the day the cesium-137 canister was stolen?

144. Do you see any example of mixtures of evidence in Table 9?

Answer

E008-GuardReport is tangible evidence about testimonial evidence, because it is a record of Sam's testimony.

145. Consider E8 Guard Report and other items of tangible evidence in Table 21(p. 226). What kind of questions would you ask about these other tangible items?

Answer

What is the authenticity, reliability, and accuracy of the records at STEMQ Company logs recording the license numbers of trucks that enter and leave the parking area?

Improving the Analysis of Competing Hypotheses

146. Some intelligence analysts may look upon Heuer's ACH methods, as well as other methods, as being ways of simplifying intelligence analyses. There are some problems associated with such views; can you think of some of these problems?

Answer

One thing that none of us, including Heuer and ourselves, can do is to simplify the world in which intelligence analysts must function every day. Just one element of the world's myriad of complexities is that it changes every day, every hour, and every minute. One way of saying this is to say that the world is not stationary; new events and situations emerge all the time. One result of this lack of stationarity is that the capabilities and intentions of real or potential adversaries may change, often in unanticipated and unrecognized ways. This is why we have emphasized the fact that discovery in intelligence analysis is a continuous process that never ceases. One obvious result of this characteristic of discovery is that it makes the accurate prediction of future events and explanations of past events extremely difficult or even impossible. As we have noted, intelligence analysts have evidence in search of hypotheses, hypotheses in search of evidence, and the testing of hypotheses, all going on at the same time during the life-cycle of an analytic problem. This raises the following interesting question.

The question is, "When does an analytic problem really end?" Stated another way, "Can we ever say when the life-cycle of an analytic problem has terminated?" There are some answers to these questions but none of them seems satisfactory. First, we might say that an analysis problem ends when we see what it has produced; have its conclusions been correct or incorrect? There are several troubles with this answer. The most obvious trouble is that it may not be possible to determine whether a conclusion has been correct or incorrect. In some cases we may have to wait a long time to answer these questions and in other cases we may never be able to discover whether a past conclusion has been correct or incorrect.

Suppose we have determined that an analytic conclusion has been correct. In addition to considering what possible further analyses are suggested by the implications of our being correct, we may also be interested in determining why our conclusion was correct. This may be useful as far as work on other related analytic problems are concerned. But, if a conclusion has been incorrect, a further post mortem analysis might reveal reasons for our failure. This will also be useful in efforts to prevent failures in other analyses.

A second answer to the question concerning the termination of an analytic problem life-cycle is that we have simply run out of time and resources to continue to perform the analysis. But the termination here may only be temporary and not final. One thing that happens in our nonstationary world is that analytic time and resources also change over time; we may later have time and resources to continue the discovery process and inferential work on a problem we have halted. A third reason analysts may provide for the termination of a problem is lack of continued interest. The analysts may say, "We stopped further work on this problem because the questions we were trying to answer are no longer relevant. An event has recently happened that now makes these questions uninteresting." The trouble here is that relevancy determination is a complex process. Questions thought to be irrelevant at one time, can be thought relevant at a later time. Perhaps this new event that caused the analysts to lose interest in the questions they were asking can be explained away. This may result in a restoration of interest in these questions and the necessity of further work on the problem.

So, we have examined three reasons for the difficulty in determining when an analytic problem really ends. We now return to the original question of simplifying intelligence analyses. It is quite true that intelligence analysts have been assisted in various ways by advancements in information technology. For example, using modern computer facilities, analysts can marshal and retrieve information in a variety of useful ways. Using such facilities, analysts can rapidly communicate information and analyses to other interested agencies. In addition, there are software systems that can assist in the structuring of analytic inferential problems involving inference networks. But the problem is that being assisted in performing a complex analytic task is not the same as simplifying these tasks. There are some very good reasons for saying that, regardless of how many new tools analysts have, they still do not simplify the basic evidential and inferential tasks they must perform; here are some reasons for this argument.

Intelligence analyses commonly rest on emerging masses of evidence of different forms and coming from a variety of sources. Trying to establish the meaning of these masses of evidence is no simple task. Analysts first have to decide what problems should be addressed and what hypotheses should be entertained. Then, some often exceedingly difficult problems arise in the careful construction of arguments justifying the relevance, credibility, and force or weight of patterns of evidence bearing on hypotheses being considered. There is no computer-based system that can inform analysts about all the parts necessary to construct defensible and persuasive arguments for the probabilistic conclusions that are sought by the analysts. These matters require an array of very difficult subjective judgments on the part of the analysts.

147. What approaches can be taken if there are no ways of simplifying the requirements for the analysis intelligence professionals face in an ever-changing world?

Answer

As we have just discussed in the answer to question 96, there is no way we can simplify the requirements for reasoning based on emerging masses of intelligence evidence in an ever-changing world. As the title of this book announces, we are concerned about the process of "connecting the dots" in intelligence analysis; the first chapter dwells on how difficult this process is whether an analysis concerns the prediction of some future events or the explanation of some past events. It would certainly help if one or more analysts were truly clairvoyant and possess what the Scots call "second sight." Such a person would be able, without any form of assistance, to accurately predict important future events and provide correct explanations for some pattern of past events. Every now and then a person will appear who claims to be clairvoyant or to have "second sight." But such persons are rarely taken seriously.

So, barring dependable and recurrent clairvoyance or "second sight", which we believe no one has, a natural question is: what do analysts bring with them in their performance of complex tasks involving the connection of massive numbers of dots? A superficial answer is that analysts bring with them their native intellectual capacity, their past educational experiences, and the results of their specific analysis training, including on-the-job training. From our experience over many years now, we have observed that intelligence analysts are so often very poorly tutored in their very analytic stock-in-trade, namely <u>evidence as the foundations of all intelligence analyses</u>.

As you have seen in the chapters of this book, we have directed the reader's attention to a study of the properties, uses, discovery, and marshaling of evidence in the probabilistic reasoning encountered in intelligence analysis. Further, we have used our system Cogent to illustrate how intelligence analyses problems involving these four matters might be addressed. This system is indeed "knowledge-based" since it already knows a substantial amount of knowledge about these four matters. But at no point in our work have we argued that we have made intelligence analysis simpler. To do so would be very misleading. What we do offer is an array of knowledge analysts might bring with them as they attempt to make sense out of masses of evidence in a world that keeps changing all the while they are trying to understand parts of it associated with our nation's security.

GLOSSARY OF TERMS

Abductive reasoning (imaginative, creative, or insightful reasoning)**.** A form of reasoning that makes some conclusion possibly true, used to generate hypotheses from data.

Access. An attribute of the competence of a human source characterizing the extent to which that source actually made the observation he or she claims to have made or had access to the information that he or she reports.

Accuracy. Also termed *sensitivity.* An attribute of the credibility of certain kinds of tangible evidence such as those provided by sensing devices and tabled information.

Ambiguous evidence. Evidence which is imprecisely stated and we cannot determine exactly what it is telling.

Analysis. A reasoning operation by which we break down a hypothesis into components to better assess it. This operation is complementary to synthesis.

Anomaly. Used with reference to evidence that seems unexplainable or out of place.

Argument. A chain of reasoning connecting observed evidence with a hypothesis of interest; The links in such chains represent ordered sequences of sources of doubt the analyst believes to be interposed between evidence and hypothesis.

Argument magnet. Magnet that attracts trifles that will form relevant evidence on major arguments for some hypothesis being entertained.

Assumption. Refers to inferences made in the absence or deficits of evidence.

Authenticity. An attribute of the credibility of tangible evidence referring to whether a tangible item is what it is claimed to be.

Baconian probability system. A probability system based on Sir Francis Bacon's eliminative and variative views on evidential reasoning. This is the only system that captures how complete our evidence covers matters that should be covered in an inference problem.

Balance of probability. A probability standard used in law to refer to cases in which the evidence just favors one hypothesis over another even by the smallest amount.

Bayesian probability system. A probability system based on the conventional axioms of probability that concern games of chance and statistics, both of which involve repeatable phenomena. This system has many deficits when applied to cases involving non-repeatable events and situations.

Belief functions. A probability system very useful for non-repeatable events and situations. In this system beliefs are associated with any combination of the considered hypotheses.

Credibility. Credential of evidence indicating the degree to which we can believe what the evidence is telling us.

Credibility magnet. Magnet that attracts trifles we have concerning the competence and credibility of our sources of intelligence information.

Beyond reasonable doubt. A probability standard used in law to refer to the highest grade of support evidence can provide some hypothesis of interest.

Bias. Used with reference to conclusions reached by a person just on the basis of personal preference rather than on a careful consideration of the evidence.

Big data. Data sets that are too large or too complex for traditional data processing applications.

Chain of custody. The sequence of the persons or devices that had access to the original source evidence, the time at which they had such access, and what they did to the original evidence when they had access to it.

Chronology magnet. Magnet that attracts inferred times at which reported events have occurred and allows inferences about the temporal ordering of these events.

Circumstantial evidence. Evidence that makes the existence of a hypothesis in an argument more or less probable "indirectly," in that at least one further inferential step is involved.

Clear and convincing evidence. A standard of proof required in congressional hearings and other tribunals.

Competence. Credibility credential of human sources of evidence. A source of evidence is competent if she/he had access to what was reported and was able to understand it. Competence also refers to any skills a person might have to do some required job.

Composition. See Synthesis.

Conclusive evidence. If believable, such evidence would make some conclusion certain.

Conjunction. The combination of judgments about the probability of individual hypotheses into a single hypothesis as a whole, where all the individual hypotheses need to be true to make the single hypothesis true. See also Disjunction.

Connecting the dots. The task of marshaling thoughts and evidence in the generation or discovery of productive hypotheses and new evidence, and in the construction of defensible and persuasive arguments on hypotheses we believe to be most favored by the evidence we have gathered and evaluated.

Contradictory evidence. Type of dissonant evidence involving events that are mutually exclusive, i.e. they cannot occur jointly. For example, one evidence item says that event E occurred and another evidence item says that event E did not occur.

Convergent evidence. Two or more evidence items that concern different events which point toward or favor the same hypothesis

Corroborative evidence. Evidence that reports the same event.

Corroborative redundant evidence. Repeated evidence about the same event.

Credential of evidence. A term used to describe a property of evidence that needs to be established or justified. Three major credentials of evidence are: relevance, credibility or credibility, and inferential force or weight.

Credibility. Concerns the extent to which an item of evidence or a source of evidence may be believed. On occasion, this term is wrongly equated with the term reliability (q.v., which has a more restricted definition). As a credential of evidence, credibility has several different attributes that depend upon the form of evidence, whether it is tangible or testimonial.

Credibility attributes. For tangible evidence, these attributes are: authenticity, accuracy, and reliability. For testimonial evidence, these attributes are: veracity, objectivity, and observational sensitivity.

Critical reasoning. Reasoning represented by an argument that is logically coherent, free of disconnects and non sequiturs.

Cumulative redundant evidence. Redundant evidence about different events.

Current intelligence. Refers to cases in which an analyst’s customer requires a conclusion in a very short time.

Data. Refer to un-interpreted signals, raw observations, or measurements, such as such as the number 6 or the color red. See also Evidence and data or item of information.

Deductive reasoning. A form of reasoning that makes some conclusion necessarily true or certain.

Defensible and persuasive argument. An argument is defensible if it is free from logical disconnects. It is persuasive if it is compelling. The trouble is that not all persuasive arguments are defensible and not all defensible arguments persuasive.

Demonstrative tangible evidence. Evidence not of a thing itself but of a representation or image of this thing.

Direction. Refers to the hypothesis we believe our evidence favors most.

Directly relevant evidence. Evidence is said to be directly relevant if a defensible chain of reasoning can be constructed that links this evidence with a hypothesis whose proof is at issue.

Discovery. Refers to the process of generating new hypotheses or new lines of inquiry and new evidence.

Disjunction. The combination of judgments about the probability of individual hypotheses into a single hypothesis as a whole, where only one of the individual hypotheses need to be true to make the single hypothesis true.

Disfavoring evidence. Evidence that argues against the truth of some hypothesis.

Dissonant evidence. Directionally inconsistent items of evidence pointing toward different hypotheses.

Divergent evidence. Type of dissonant evidence where the evidence points to different hypotheses as opposed to contradictory evidence which involves events that are mutually exclusive.

Divide and conquer. The act of decomposing a complex reasoning task into its simpler ingredients and of combining their conclusions. See Analysis and Synthesis.

Dots. Details in the observable information or data about an intelligence situation, as well as potential links in chains of reasoning or arguments we may construct to link dots to hypotheses we are trying to prove or disprove.

Eliminative induction. A method of proof in which a variety of evidential tests are employed in an effort to eliminate alternative hypotheses being considered. The hypothesis that best resists our eliminative attempts is the one that can be taken most seriously.

Eliminative magnet. Magnet that attracts trifles representing evidence relevant in showing why some hypothesis can be safely eliminated.

Epistemology. A branch of philosophy concerning the acquisition of and validity of knowledge.

Evidence. Evidence is any observable sign, indicator, or datum we believe is relevant in deciding upon the extent to which we infer any hypotheses we have entertained as being correct or incorrect.

Evidence and data or item of information. Evidence differs from data or items of information. Data or Items of information only become evidence when their relevance is established regarding some matter to be proved or disproved.

Evidence and events. There is an important distinction to be made between evidence of some event and the event itself. Having evidence that an event occurred does not entail that this event did occur. What is at issue is the credibility of the evidence and its source(s).

Evidence-based hypothesis assessment. The process of determining the probability of a hypothesis based on the available evidence.

Evidence custodian. Person designated of making careful records of every person who had access to an evidence item from the time it was received, what they did with this item, how long they held the item, and who next received the item before it was finally introduced at trial.

Evidence in search of hypotheses. The “bottom-up” generation of a new hypothesis from evidence.

Evidential dots. One of two forms of dots that must be connected. The other form of dots concern ideas about the meaning of an evidential dot.

Evidentiary testing of hypothesis. See Evidence-based hypothesis assessment.

Fact. Any event or act or condition of things, assumed (for the moment) as having happened or having existed.

Favoring evidence. Evidence that is directionally consistent in favoring the same hypothesis.

Force. See Inferential force or weight.

Fuzzy probability system. A probability system where the uncertainty about a conclusion reached is expressed in words (such as “likely” or “almost certain”), each “fuzzy” word being related to a range of numerical probabilities by a possibility function.

Generalization. A general proposition claimed to be true which is used implicitly or explicitly to argue that a conclusion has been established.

Harmonious evidence. Two or more items of evidence that are directionally consistent in the sense that they all point toward, or favor, the same hypothesis or possible conclusion.

Heuristic. A rule of thumb that aids you in any discovery, inference, learning, or decision problem.

Holistic approach to analysis. Work on an analysis problem where you do all in your own head without decomposing it in any way or seeking the assistance of others.

HUMINT. Testimonial evidence given by a human source about some matter of intelligence interest.

Hypotheses magnet. Magnet that uses generated hypotheses to attract information items that could become relevant evidence in their favor or against.

Hypothesis. A general proposition put forward as a possible explanation for known facts from which additional investigations can be planned to generate evidential data that will tend to strengthen or weaken the basis for accepting the proposition as the best or strongest explanation of the available data. Hypotheses commonly refer to possible alternative conclusions we could entertain about matters of interest in an analysis.

Hypothesis in search of evidence. The “top-down” generation of evidence believed to be consistent with the hypothesis, and therefore useful in testing this hypothesis.

Idea dots. One of two types of dots that must be connected, having the form of links in chains of reasoning or arguments we construct to link evidential dots to hypotheses. The other type of dots are the evidential dots.

Imaginative reasoning. See Abductive reasoning.

IMINT. Tangible evidence gathered from satellite, aerial photography, or mapping/terrain data.

Inconclusive evidence. Evidence which is consistent with the truth of more than one hypothesis or possible explanation.

Indirectly relevant evidence. See Ancillary evidence.

Inductive reasoning. A form of reasoning that makes some conclusion probably true, used to test hypotheses based on evidence.

Inference. The process of deriving logical conclusions from premises.

Inference network. Network consisting of multiple lines of argument that connect many different kinds of evidence to the hypothesis under consideration.

Inferential force or weight. Credential of evidence indicating how strong the evidence is in favoring or disfavoring hypotheses we are considering.

Information. Data equipped with meaning provided by a certain context, such as “6 a.m.” or “$6”. See also Evidence and data or item of information.

Knowledge. Justified true belief. We say that Person $\mathcal{A}$ knows that event B occurred if the event B did occur (true), the person $\mathcal{A}$ got non-defective evidence that B occurred (justified), and $\mathcal{A}$ believed this evidence (belief).

Likelihood. Probability of evidence E* given some hypothesis H, written P(E*|H).

Likeliness. Probability of a hypothesis H given some evidence E*, written P(H|E*).

Marshaling. Refers to the bringing together of thoughts and evidence during hypotheses generation and argument construction. Having useful strategies for marshaling thus helps advance the processes of hypotheses generation and analysis.

Marshaling magnet. A metaphoric description of an evidence marshaling operation that serves to attract particular combinations of evidence from some collection of data or trifles and that can assist in generating new hypotheses or that can open up new lines of inquiry and evidence.

MASINT. Measures and signatures intelligence. Evidence of the traces left behind by objects and processes.

Meta-evidence. See Ancillary evidence.

Nugget. A term used by intelligence agencies with reference to believable/credible evidence that would make some conclusion certain.

Objectivity. An attribute of the credibility of a human source characterizing the extent to which that source based her/his belief that the reported belief occurred on her/his sensory evidence rather than on what this source expected or desired to observe.

Observational sensitivity. An attribute of the credibility of a human source characterizing how good was the sensory evidence this source received under the conditions in which her/his observation was made.

Posterior belief. Belief assessed after we receive and incorporate the evidence we have.

Posterior probability. Probability of hypothesis after we receive and incorporate the evidence we have. See Likeliness.

Prior probability. Probability used to indicate the initial conditions of our uncertainty before we consider evidence that begins to emerge.

Probability. Characterizes the uncertainty about a given event, statement, hypothesis or conclusion. Differing conceptions of probability are a matter of considerable controversy and debate within statistics and the logic of proof.

Proposition. A statement that is true or false, that can be affirmed or denied.

Question magnet. Magnet that attracts trifles representing possible answers to any question that comes to mind as an intelligence analysis proceeds.

Real tangible evidence. Evidence of the thing itself that can be directly examined.

Reduction. See Analysis.

Redundant evidence. Two or more evidence items that either say the same thing over again or do not add anything to what we already have.

Relevance. Credential of evidence indicating how a datum or information item is linked to something we are trying to prove or disprove.

Relevant evidence. Evidence having any tendency to make the existence of any fact that is of consequence to the determination of the action more probable or less probable than it would have been without the evidence. Federal Rule of Evidence 401.

Reliability. Reliability is especially relevant to various forms of sensors that provide us with many forms of demonstrative tangible evidence. A system, sensor, or test of any kind is reliable to the extent that the results it provides are repeatable or consistent. This term is often used incorrectly as a synonym for the term *credibility* or *credibility,* which involves other attributes.

Scenario magnet. Magnet that attracts a temporally-ordered sequence of trifles forming relevant evidence about events that will form the basis for a story or scenario about what has happened in some situation of interest.

SIGINT. Signals intelligence used with reference evidence obtained with sensors and recording devices.

Standard of proof. The degree of persuasion required to establish a particular fact. The standard of proof in civil cases is typically "the preponderance of the evidence" or "the balance of probabilities." In criminal cases the prosecution has to satisfy the standard of "beyond reasonable doubt" in order to succeed. In some non-criminal cases, the standard of proof is said to be "clear and convincing."

Substance-blind. A term used to describe a particular way of categorizing forms and combinations of evidence without regard to its substance or content. Such a categorization is based on the inferential properties of evidence and not on its content.

Synergistic evidence. Two or more evidence items which have greater inferential force or weight than they would have if considered separately or independently.

Synthesis. A reasoning operation by which we combine the assessments of the sub-hypotheses of a hypothesis in order to obtain an assessment of the hypothesis. This operation is complementary to analysis.

Tangible evidence. Evidence that can be directly examined by persons drawing conclusions to see what event(s) this evidence reveals. Examples include objects, documents, images, measurements, and charts.

Task decomposition. See Analysis.

Testimonial evidence. Evidence provided by a human source. Testimonial evidence about some event can be based on direct observations, secondhand reports from another source, or on the basis of opinion or inferences based on information about the occurrence of other events.

Trifles. A term used by Sherlock Holmes to refer to the many details (dots) he observed that formed the basis for his investigations.

Understandability. An attribute of the competence of a human source characterizing the extent to which that source understood what was being observed well enough to provide us with an intelligible account.

Veracity. An attribute of the credibility of a human source characterizing the extent to which that source believes that the reported event occurred

Weight of evidence. See Inferential force or weight.

REFERENCES

Anderson, T., Schum, D., and Twining, W. (2005). *Analysis of Evidence,* Cambridge University Press.

Baring-Gould, W. S. (1967). *The Annotated Sherlock Holmes*, Vol. I and Vol. II, Clarkson N. Potter, New York, NY.

Betham, J. (1810). An Introductory View of the Rationale of the Law of Evidence for Use by Non-lawyers as well as Lawyers (vi works 1-187 (Bowring edition, 1837-43) originally edited by James Mill circa 1810).

Boicu, M., Tecuci, G., Marcu, D., Bowman, M., Shyr, P., Ciucu, F., and Levcovici, C. (2000) "Disciple-COA: From Agent Programming to Agent Teaching", in *Proceedings of the Seventeenth International Conference on Machine Learning* (ICML), Stanford, California 2000, Morgan Kaufman. http://lac.gmu.edu/publications/data/2000/2000_il-final.pdf

Boicu, M., Tecuci, G., Marcu, D. (2012). Rapid Argumentation Capture from Analysis Reports: The Case Study of Aum Shinrikyo, in *Proceedings of the Seventh International Conference on Semantic Technologies for Intelligence, Defense, and Security – STIDS 2012*, Fairfax, VA, 23-26 October. http://ceur-ws.org/Vol-966/STIDS2012_T05_BoicuEtAl_RapidArgumentation.pdf

Bruce, J. B. (2008). Making Analysis More Reliable: Why Epistemology Matters to Intelligence, in George R.Z., Bruce J.B., eds., *Analyzing Intelligence: Origins, Obstacles, and Innovations*, Georgetown University Press, Washington, DC, pp. 184-185.

Chantrell, G., ed. (2004). *Oxford Dictionary of Word Histories*. Oxford University Press, Oxford UK.

Clemen, R. T. (1995). *Making Hard Decisions*. Duxbury Press, Belmont. CA.

Cohen, L. J. (1977). *The probable and the provable.* Oxford, UK: Clarendon Press.

Cohen, L. J. (1989). An introduction to the philosophy of induction and probability. Oxford, UK: Clarendon Press.

Dale, A. (2003). Most Honourable Remembrance: The Life and Work of Thomas Bayes, Springer-Verlag, New York, NY.

Danzig, R., Sageman, M., Leighton, T., Hough, L., Yuki, H., Kotani, R., and Hosford Z.,M. (2011). Aum Shinrikyo: Insights Into How Terrorists Develop Biological and Chemical Weapons, Center for a New American Security, Washington, DC, July.

David, F.N. (1962). *Gods, Games and Gambling,* Griffin, London.

DOD J. P. 1-02 (2010). Department of Defense Dictionary of Military and Associated Terms, *Joint Publication 1-02*, 8 November 2010.

Drogin, B. (2007). CURVEBALL: Spies, Lies, and the Con Man Who Caused a War. Random House, New York, NY.

Emerson, S. (2006). *Jihad Incorporated: A Guide to Militant Islam in the US*. Prometheus Books, Amherst, NY, 468 – 469

Forbus, K. (2015). Analogical abduction and prediction, *2015 AAAI Fall Symposium - Technical Report*.

George, R. and Bruce, J. B., eds. (2008). *Analyzing Intelligence: Origins, Obstacles, and Innovations.* Georgetown University Press, Washington. DC.

Heuer, R. J. (1999). *Psychology of Intelligence Analysis*, Center for the Study of Intelligence, Central Intelligence Agency, Washington, DC.

Heuer, R. J. (2008). Computer-Aided Analysis of Competing Hypotheses, in George R. Z., Bruce J. B., eds., *Analyzing Intelligence: Origins, Obstacles, and Innovations*, Georgetown University Press, Washington, DC.

Heuer, R. J. Jr., and Pherson, R. H. (2011). *Structured Analytic Techniques for Intelligence Analysis*, CQ Press.

Hintikka, J. (1983). Sherlock Holmes Formalized. In: Eco, U., Sebeok, T., *The Sign of Three: Dupin, Holmes, Peirce.* Indiana University Press.

Howe, M. (1999). *Genius Explained*, Cambridge University Press.

Johnston, R. (2005). *Analytic Culture in the U. S. Intelligence Community*. Central Intelligence Agency, Washington, DC.

Josephson, J. R., and Josephson, S. G. (1994). *Abductive Inference: Computation, Philosophy, Technology*. Cambridge University Press.

Kahneman, D., Tversky, A. (1974). Judgment under Uncertainty: Heuristics and Biases*. Science*, 27 September, Vol. 185, pp. 1124-1131.

Kahneman, D., Slovic, P., Tversky, A. (1982). *Judgment under Uncertainty: Heuristics and Biases.* Cambridge University Press.

Kent, S. (1994). Words of Estimated Probability, in Steury, D. P. (ed.), *Sherman Kent and the Board of National Estimates: Collected Essays*, Center for the Study of Intelligence, Central Intelligence

Agency, Washington, DC.

Kolmogorov, A. N. (1956). *Foundations of a Theory of Probability* (1933), 2nd English edition, Chelsea Publishing, New York, NY., pp. 3-4.

Kolmogorov, A. N. (1969). The Theory of Probability. In: Aleksandrov, A. D., Kolmogorov, A. N., Lavrentiev, M. A. (eds.), *Mathematics: Its Content, Methods, and Meaning*. MIT Press, Cambridge, MA, Volume Two, Chapter XI, pp. 231-264.

Langley, P. (2019). Heuristic Construction of Explanations Through Associative Abduction, *Proceedings of the Seventh Annual Conference on Advances in Cognitive Systems*, Technical Report Number COLAB2-TR-4, 21-36, Massachusetts Institute of Technology, Cambridge, Massachusetts, August 2-5.

Lempert, R. O., Gross, S. R., and Liebman, J. S. (2000). *A Modern Approach to Evidence*, 3rd edition, West Publishing, St. Paul, MN.

Lindberg, C. A., ed. (2004). *The Oxford American Writer's Thesaurus*. Oxford University Press, Oxford, UK.

Marrin, S. (2011). Improving Intelligence Analysis: Bridging the gap between scholarship and practice. Routlege, London and New York.

Marrin, S., and Clemente J. D. (2005). Improving Intelligence Analysis by Looking to the Medical Profession, *International Journal of Intelligence and CounterIntelligence,* Vol. 18, No. 4, pp. 707–729.

Martin, D. C. (1980). *Wilderness of Mirrors*, Ballantine, New York, NY.

Meckl, S., Tecuci, G., Boicu, M., and Marcu, D. (2013). Towards an Operational Semantic Theory of Cyber Defense Against Advanced Persistent Threats, in Laskey, K. B., Emmons, I., Costa, P. C. G., and Oltramari, A., *Proceedings of the Tenth International Conference on Semantic Technologies for Intelligence, Defense, and Security – STIDS 2015*, pp. 58–65, Fairfax, VA, November 18-20. lac.gmu.edu/publications/2015/APT-LAC.pdf

Moore, D. T. (2011). Sensemaking: A Structure for Intelligence Revolution. NDIC Press.

Mueller, C. B., and Kirkpatrick L. C. (2009). Federal Rules of Evidence, 2009 edition, West Publishing Co., St. Paul, MN.

Murphy, P. (2003). *Evidence, Proof, and Facts: A Book of Sources*, Oxford University Press, Oxford, UK.

Negoita, C. V., and Ralescu, D. A. (1975). *Applications of Fuzzy Sets to Systems Analysis*. Wiley: New York.

Nilsson, N.J. (1971). *Problem Solving Methods in Artificial Intelligence*, McGraw-Hill, New York, NY.

Nolte, W. (2005). *Keeping Pace with the Revolution in Military Affairs,* Center for the Study of Intelligence, Central Intelligence Agency, Washington D.C.

Paul, George, P. (2007). The "Authenticity Crisis" in Real Evidence. *Law Practice Today*. 2007.

Peirce, C. S. (1898). *Reasoning and the Logic of Things*, 1898, in Ketner, K. (ed.), Harvard University Press, Cambridge, MA, 1992.

Peirce, C. S. (1901). Abduction and Induction, in Buchler, J. (ed.) (1995). *Philosophical Writings of Peirce*, 150-156, Dover: New York, NewYork.

Pherson, R., Boardman, M. (2017). *Cognitive Biases and Intuitive Traps Most Often Encountered by Analysts: Which Structured Analytic Techniques Best Mitigate Their Impact?* 2017 International Studies Association Annual Convention.

Pherson, K.H., Pherson, R.H. (2021). *Critical Thinking for Strategic Intelligence,* CQ Press.

Pomerol, J.C., Adam, F., On the Legacy of Herbert Simon and his Contribution to Decisionmaking Support Systems and Artificial Intelligence, in Jatinder N.D. Gupta, Giusseppi A. Forgionne and Manuel Mora T. (eds.), *Intelligent Decision-making Support Systems: Foundations, Applications and Challenges*, Springer-Verlag, London, 2006, pp. 25-43.

Schum D. A. (1987). *Evidence and Inference for the Intelligence Analyst* (2 volumes), University Press of America, Lanham, MD.

Schum, D. A. (1989). Knowledge, Probability, and Credibility, *Journal of Behavioral Decision Making*, Vol. 2, pp. 39-62.

Schum, D. A. (1991). Jonathan Cohen and Thomas Bayes on the Analysis of Chains of Reasoning, in Eells, E., and Maruszewski, T. (eds.), *Probability and Rationality: Studies on L. Jonathan Cohen's Philosophy of Science*, Editions Rodopi, pp. 99 – 145.

Schum, D. A. (1999). Marshaling Thoughts and Evidence During Fact Investigation. *South Texas Law Review* (pp. 419-423), Vol. 40, No. 2, Summer 1999, pp 401-454.

Schum, D. A. (1994/2001a). Northwestern University Press.

Schum, D. A. (2001a). *The Evidential Foundations of Probabilistic Reasoning*. Evanston, IL:

Northwestern University Press.

Schum, D. A. (2001b). Species of Abductive Reasoning in Fact Investigation in Law. *Cardozo Law Review*, 22 (5-6), pp. 1645–1681.

Schum, D. A., and Morris, J. (2007). Assessing the Competence and Credibility of Human Sources of Evidence: Contributions from Law and Probability, *Law, Probability and Risk,* Vol. 6, pp. 247-274.

Schum, D. A. (2009). Classifying Forms and Combinations of Evidence: Necessary in a Science of Evidence, in *Evidence, Inference and Inquiry,* The British Academy, Oxford University Press.

Schum, D. A. (2009). Science of Evidence: Contributions from Law and Probability. Law Probab Risk 8 197–231.

Schum, D. A., Tecuci, G, and Boicu, M. (2009a). Analyzing Evidence and its Chain of Custody: A Mixed-Initiative Computational Approach, *International Journal of Intelligence and Counterintelligence,* Vol. 22, pp. 298-319. http://lac.gmu.edu/publications/2009/Schum et al - Chain of Custody.pdf

Shafer, G. (1976). *A Mathematical Theory of Evidence,* Princeton University Press, Princeton, NJ.

Shafer, G. (1988). Combining AI and OR. University of Kansas School of Business Working Paper No. 195, April.

Simonite, T. (2013). “Bill Gates: Software Assistants Could Help Solve Global Problems”, MIT Technology Review, 16 July, http://www.technologyreview.com/news/517171/bill-gates-software-assistants-can-save-the-world/

Tecuci, G. (1988). Disciple: A Theory, Methodology and System for Learning Expert Knowledge, *Thèse de Docteur en Science,* University of Paris-South.
http://lac.gmu.edu/publications/1988/TecuciG_PhD_Thesis.pdf

Tecuci, G., and Michalski, R. S. (1991). A Method for Multistrategy Task-adaptive Learning Based on Plausible Justifications, in Birnbaum L., and Collins G. (eds.), *Machine Learning: Proc. of the Eighth International Conference,* pp. 549-553, Chicago, June, Morgan Kaufmann.
http://lac.gmu.edu/publications/1991/TecuciG_Multistrategy_Learning_Method.pdf

Tecuci, G. (1993). Plausible Justification Trees: a Framework for the Deep and Dynamic Integration of Learning Strategies, *Machine Learning Journal*, vol.11, pp. 237-261.
http://lac.gmu.edu/publications/1993/TecuciG_Plausible_Justification_Trees.pdf

Tecuci, G. (1998). *Building Intelligent Agents: An Apprenticeship Multistrategy Learning Theory,* *Methodology, Tool and Case Studies*. London, England: Academic Press. http://lac.gmu.edu/publications/1998/TecuciG_Building_Intelligent_Agents/default.htm

Tecuci G. and Kodratoff Y. eds. (1995). *Machine Learning and Knowledge Acquisition: Integrated Approaches,* Academic Press. http://lac.gmu.edu/publications/1995/TecuciG_MLKA_Integrated_Approaches.pdf

Tecuci, G., Boicu, M., Wright, K., Lee, S. W., Marcu, D., and Bowman, M. (1999). "An Integrated Shell and Methodology for Rapid Development of Knowledge-Based Agents," in *Proceedings of the Sixteenth National Conference on Artificial Intelligence* (AAAI-99), July 18-22, Orlando, Florida, AAAI Press, Menlo Park, CA. http://lac.gmu.edu/publications/data/1999/ismrdkba.pdf

Tecuci, G. and Keeling, H. (1999). Developing an Intelligent Educational Agent with Disciple, *International Journal of Artificial Intelligence in Education*, vol. 10, no. 3-4. http://lac.gmu.edu/publications/1999/TecuciG_Intelliget_Educational_Agent.pdf

Tecuci, G., Boicu, M., Wright, K., Lee, S. W., Marcu, D. and Bowman, M. (2000). A Tutoring Based Approach to the Development of Intelligent Agents, in Teodorescu, H.N., Mlynek, D., Kandel, A. and Zimmermann, H.J. (editors). *Intelligent Systems and Interfaces,* Kluwer Academic Press. http://lac.gmu.edu/publications/data/2000/2000_Disciple-Planning.pdf

Tecuci, G., Boicu, M., Bowman, M., Marcu, D., with commentary by Burke, M. (2001). An Innovative Application from the DARPA Knowledge Bases Programs: Rapid Development of a Course of Action Critiquer, *AI Magazine,* **22(2)**, 43-61. http://lac.gmu.edu/publications/2001/TecuciG_Disciple_COA_IAAI.pdf

Tecuci, G., Boicu, M., Marcu, D., Stanescu, B., Boicu, C., Comello, J., Lopez, A., Donlon, J., Cleckner, W. (2002a) Development and Deployment of a Disciple Agent for Center of Gravity Analysis, in *Proceedings of the Eighteenth National Conference of Artificial Intelligence and the Fourteenth Conference on Innovative Applications of Artificial Intelligence*, AAAI-02/IAAI-02, pp. 853-860, Edmonton, Alberta, Canada, AAAI Press/The MIT Press. http://lac.gmu.edu/publications/data/2002/dddacga.pdf

Tecuci, G., Boicu, M., Marcu, D., Stanescu, B., Boicu, C., Comello, J. (2002b). Training and Using Disciple Agents: A Case Study in the Military Center of Gravity Analysis Domain, in *AI Magazine,* 24(4), 51-68, AAAI Press, CA. http://lac.gmu.edu/publications/2002/TecuciG_Disciple_COG_IAAI.pdf

Tecuci G., Boicu M., Ayers C., and Cammons D. (2005a). Personal Cognitive Assistants for Military Intelligence Analysis: Mixed-Initiative Learning, Tutoring, and Problem Solving, in *Proceedings of the First International Conference on Intelligence Analysis*, McLean, VA, 2-6 May.

http://lac.gmu.edu/publications/data/2005/Tecuci-Disciple-LTA.pdf

Tecuci, G., Boicu, M., Boicu, C., Marcu, D., Stanescu, B., Barbulescu, M. (2005b). The Disciple-RKF learning and reasoning agent. *Computational Intelligence* **21(4)** 462–479. http://lac.gmu.edu/publications/2005/TecuciG_Disciple_RKF_CI.pdf

Tecuci, G., Boicu, M., Cox, M. T. (2007a). Seven Aspects of Mixed-initiative Reasoning: An Introduction to the Special Issue on Mixed-initiative Assistants. *AI Magazine* **28(2),** 11–18. http://www.aaai.org/ojs/index.php/aimagazine/issue/view/174/showToc

Tecuci, G., Boicu, M., Marcu, D., Boicu, C., Barbulescu, M., Ayers, C., Cammons, D. (2007b). Cognitive Assistants for Analysts, *Journal of Intelligence Community Research and Development (JICRD)*, 2007. Also published in John Auger, William Wimbish (eds.), *Proteus Futures Digest: A Compilation of Selected Works Derived from the 2006 Proteus Workshop*, Joint publication of the National Intelligence University, Office of the Director of National Intelligence, and U.S. Army War College Center for Strategic Leadership, 2007, pp. 303-329. http://lac.gmu.edu/publications/2007/TecuciG_Cognitive_Assistants.pdf

Tecuci, G., Marcu, D., Boicu, M., Le, V. (2007c). Mixed-Initiative Assumption-Based Reasoning for Complex Decision-Making, *Studies in Informatics and Control*, 16(4), December.

Tecuci, G., Boicu, M., Marcu, D., Boicu, C., Barbulescu, M. (2008a). Disciple-LTA: Learning, Tutoring and Analytic Assistance, *Journal of Intelligence Community Research and Development* (JICRD), July. http://lac.gmu.edu/publications/2008/Disciple-LTA08.pdf

Tecuci, G., Boicu, M., and Comello, J. (2008b). *Agent-Assisted Center of Gravity Analysis*, CD with Disciple-COG and Lecture Notes used in courses at the U.S. Army War College and Air War College, GMU Press, ISBN 978-0-615-23812-8. http://lac.gmu.edu/cog-book/

Tecuci, G., Boicu, M., Marcu, D., Barbulescu, M., Boicu, C., Le, V., Hajduk, T. (2008c). Teaching Virtual Experts for Multi-Domain Collaborative Planning, *Journal of Software*, Volume 3, Number 3, pp. 38-59, March. http://lac.gmu.edu/publications/2008/TecuciG_Disciple_VE_JS.pdf

Tecuci, G., Schum, D. A., Boicu, M., Marcu, D., Hamilton, B., Wible, B. (2010). Teaching Intelligence Analysis with TIACRITIS, *American Intelligence Journal,* Vol. 28, No. 2, December. http://lac.gmu.edu/publications/2010/Tiacritis-AIJ.pdf

Tecuci, G., Marcu, D., Boicu, M., Schum, D.A., Russell K. (2011a). Computational Theory and Cognitive Assistant for Intelligence Analysis, in *Proceedings of the Sixth International Conference on Semantic Technologies for Intelligence, Defense, and Security – STIDS 2011*, pp. 68-75, Fairfax, VA, 16-18 November. http://ceur-ws.org/Vol-808/STIDS2011_CR_T9_TecuciEtAl.pdf

Tecuci, G., Schum, D. A., Boicu, M., Marcu, D. (2011b). *Introduction to Intelligence Analysis: A Hands-on Approach with TIACRITIS*, 220 pages, Learning Agents Center, George Mason University, First edition 2010, Second edition 2011. http://lac.gmu.edu/publications/2011/TecuciG_Introduction_Intelligence_Analysis.pdf

Tecuci, G., Schum, D. A., Marcu, D., Boicu, M. (2013a). Recognizing and Countering Biases in Intelligence Analysis with TIACRITIS, in *Proceedings of the Eighth International Conference on Semantic Technologies for Intelligence, Defense, and Security* – STIDS, Fairfax, VA, 13-14 November. http://ceur-ws.org/Vol-1097/STIDS2013_T04_TecuciEtAl.pdf

Tecuci, G., Boicu, M., Marcu, D., Schum, D. (2013b). How Learning Enables Intelligence Analysts to Rapidly Develop Practical Cognitive Assistants, in *Proceedings of the 12th International Conference on Machine Learning and Applications (ICMLA'13)*, Miami, Florida, 4-7 December. http://lac.gmu.edu/publications/2013/LAC-ICMLA-13.pdf

Tecuci, G., Schum, D. A., Marcu, D., Boicu, M. (2014). Computational Approach and Cognitive Assistant for Evidence-Based Reasoning in Intelligence Analysis, *International Journal of Intelligent Defence Support Systems*, Vol. 5, No. 2, pp. 146 – 172. http://lac.gmu.edu/publications/2014/Disciple-CD-IJIDSS.pdf.

Tecuci, G., Marcu, D., Boicu, M., Schum, D. A. (2015). COGENT: Cognitive Agent for Cogent Analysis, in *Proceedings of the 2015 AAAI Fall Symposium "Cognitive Assistance in Government and Public Sector Applications"*, Arlington, VA, 12-14 November.

Tecuci, G., Schum, D. A., Marcu, D., Boicu, M. (2016a). *Intelligence Analysis as Discovery of Evidence, Hypotheses, and Arguments: Connecting the Dots*. New York, NY: Cambridge University Press.

Tecuci, G., Marcu, D., Boicu, M., Schum, D. A. (2016b). *Knowledge Engineering: Building Cognitive Assistants for Evidence-based Reasoning*. New York, NY: Cambridge University Press.

Tecuci, G., Kaiser, L., Marcu, D., Uttamsingh, C., Boicu, M. (2018). Evidence-based Reasoning in Intelligence Analysis: Structured Methodology and System, *Computing in Science and Engineering,* 20(6) 9-21, November/December.

Tecuci, G., Meckl, S., Marcu, D., Boicu, M. (2019). Instructable Cognitive Agents for Autonomous Evidence-Based Reasoning, *Advances in Cognitive Systems*, Vol. 8, 2019. Also in *Proceedings of the Seventh Annual Conference on Advances in Cognitive Systems*, Technical Report Number COLAB^2-TR-4, pp.183-204, August 2-5, 2019, Massachusetts Institute of Technology, Cambridge, MA.

Tecuci, G., Marcu, D., Boicu, M., Kaiser, L. (2020c). Instructable Cognitive Agent to Perform

Sensemaking in Intelligence, Surveillance and Reconnaissance. T*he Eight Annual Conference on Advances in Cognitive Systems*, August 8-10, 2020.

Toulmin, S. (1964). *The Uses of Argument.* Cambridge University Press, paperback edition, Cambridge UK.

Turoff, M. (2007). Design of Interactive Systems, in *Emergency Management Information Systems Tutoria*l, The Hawaii International Conference on System Sciences, HICSS-40, Hawaii, 3 January.

Walton, D. (2005). *Abductive Reasoning*. The University of Alabama Press.

Wigmore, J. H. (1913). The Problem of Proof. *Illinois Law Review*, Vol. 8, No. 2, 1913, 77-103.

Wigmore, J. H. (1937). The Science of Judicial Proof: As Given by Logic, Psychology, and General Experience, and Illustrated in Judicial Trials. 3rd Edition, Little, Brown & Co., Boston, MA, 1937.

van Gelder, T. J. (2007). The Rationale for Rationale, *Law, Probability and Risk*, 6, pp. 23-42.

W3C (2015). Semantic Web, http://www.w3.org/standards/semanticweb/

Zadeh, L. (1983). The Role of Fuzzy Logic in the Management of Uncertainty in Expert Systems, *Fuzzy Sets and Systems*, Vol.11, pp. 199 - 227.

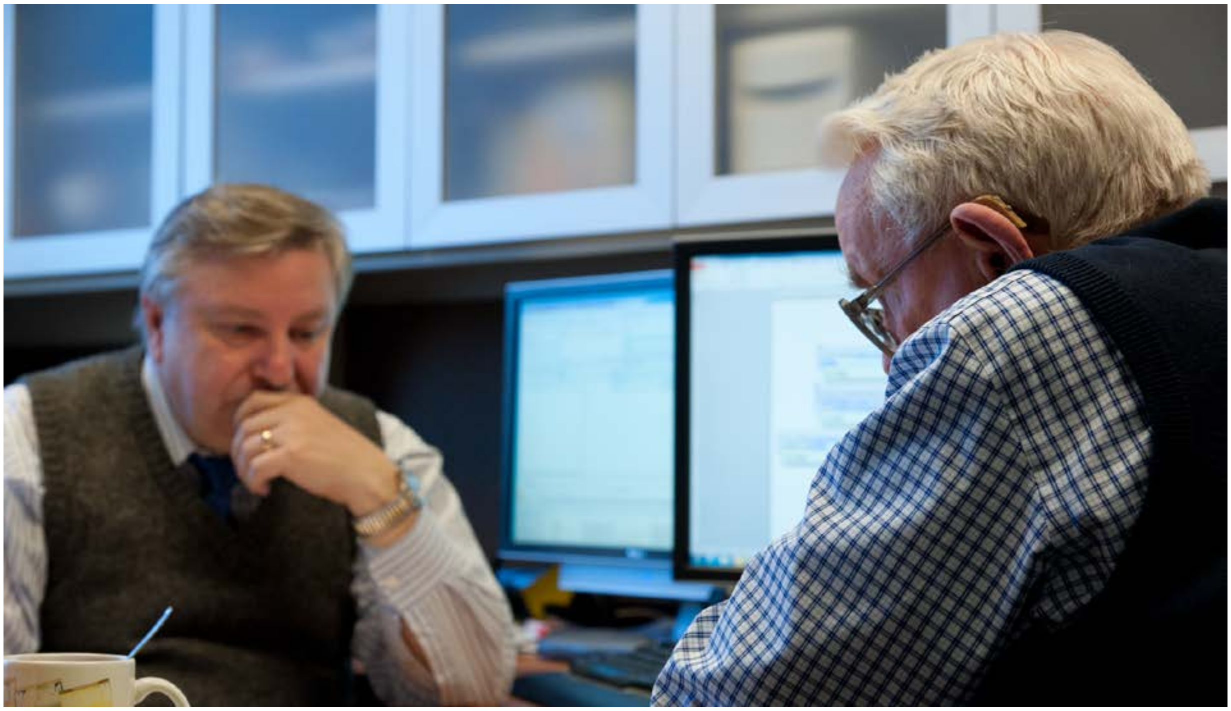

Gheorghe Tecuci is Professor of Computer Science and Director of the Learning Agents Center in the School of Computing of George Mason University, Member of the Romanian Academy, and former Chair of Artificial Intelligence in the Center for Strategic Leadership of the U.S. Army War College.

David A. Schum was Professor of Systems Engineering and Operations Research in the School of Enginnering, Professor of Law in the School of Law, Chief Scientist of the Learning Agents Center at George Mason University, and Honorary Professor of Evidence Science at University College London.