1 Intelligence Analysis: "Connecting the Dots"

1.1 HOW EASY IS IT TO CONNECT THE DOTS?

We have included a frequently used metaphor in our book's title: "Connecting the Dots." This metaphor seems appropriate in characterizing the evidential and inferential matters discussed in this book. The metaphor may have 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, 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 on television and radio, 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. We 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 *numbered* 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 (unnumbered) 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.*

https://doi.org/10.1017/CBO9781316388488.002 Published online by Cambridge University Press

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

1.1.1 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 based on the observance of trifles." A related problem here is that most items of intelligence evidence 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 Chapter 4, we give special attention to the problem of what qualifies as an evidential dot. *Not all data or items of information we have will ever become evidence in an analysis task.*

Example 1.1.

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 1.1. 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.

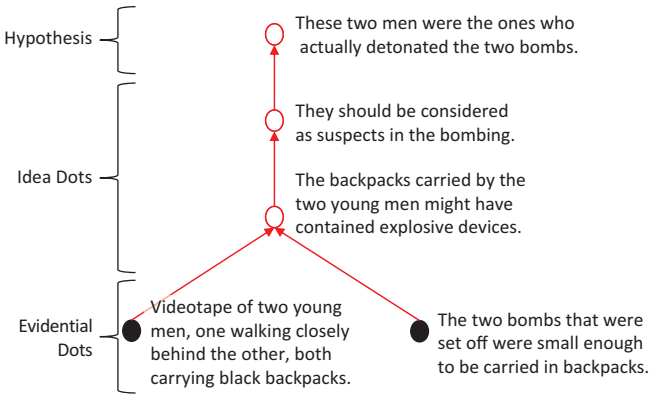

Figure 1.1. Types of dots to be connected: evidence, ideas, and hypotheses.

https://doi.org/10.1017/CBO9781316388488.002 Published online by Cambridge University Press

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, which 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 refers 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?

1.1.2 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. Just because this source says that this event occurred does not entail that it did occur. *So 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* represents the actual occurrence of event E.
- E^*_i represents the reported occurrence of event E from source *I*.

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 *believability* of source *I*. There are some real challenges here in discussing the believability of source *I*. Chapter 6 of this book is devoted to the task of assessing the believability of our sources of intelligence evidence. As we will see, the Disciple-CD system already knows much about this crucial task.

But there are even distinctions to be made in what we have called *evidential dots*. Some of these dots arise from objects we obtain or from sensors that supply us with records or images of various sorts. So one major kind of evidential dot involves what we can call *tangible evidence* that we can observe for ourselves to see what events it may reveal. In many other cases, we have no such tangible evidence but must rely upon the reports of human sources who allegedly have made observations of events of interest to us. Their reports to us come in the form of *testimonial evidence* or assertions about what they have observed. Therefore, an evidential dot E^*_i can be one of the following types:

- *Tangible evidence* such as objects of various kinds, or sensor records like those obtained by signals intelligence (SIGINT), imagery intelligence (IMINT), measurement and signature intelligence (MASINT), and other possible sources.
- *Testimonial evidence* obtained from human sources, or human intelligence (HUMINT).

The origin of one of the greatest challenges in assessing the *believability* of evidence is that we must ask different questions about the sources of tangible evidence than those we ask about the sources of testimonial evidence. Stated another way, the believability attributes of tangible evidence are different from the believability attributes of testimonial evidence.

https://doi.org/10.1017/CBO9781316388488.002 Published online by Cambridge University Press

Example 1.2.

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 Boyleston Street in Boston, Massachusetts, 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 would certainly arise in the minds of defense attorneys during the 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 Boyleston Street in Boston, Massachusetts? It may have been recorded on a different street in Boston, or perhaps on a street in a different city.

But there is another difficulty that is not always recognized that can cause endless trouble. While, in the case of tangible evidence, believability and credibility may be considered as equivalent terms, human sources of evidence have another characteristic apart from credibility; this characteristic involves their *competence*. As we discuss in Section 6.4, the credibility and competence characteristics of human sources must not be confused; to do so invites *inferential catastrophes*, as we will illustrate. The questions required to assess human source competence are different from those required to assess human source credibility. Competence requires answers to questions concerning the source's actual *access* to, and *understanding* of, the evidence he or she reports. Credibility assessment for a testimonial source requires answers to questions concerning the *veracity*, *objectivity*, and *observational sensitivity or accuracy* of the source. The Disciple-CD system knows what credibility-related questions to ask of tangible evidence and the competence and credibility-related questions to ask of HUMINT sources. We have much more to say about the forms and combinations of evidence in Chapters 6, 7, and 8 of this book.

There is no better way of illustrating the importance of evidence believability 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 believability of source *I* (whatever this source might be). In Figure 1.2 is a simple

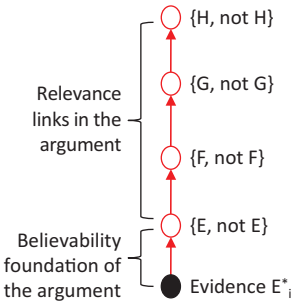

Figure 1.2. The believability foundation for an argument.

https://doi.org/10.1017/CBO9781316388488.002 Published online by Cambridge University Press

example illustrating this believability 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 whether or not hypothesis H is true. As shown, the very first stage in this argument concerns an inference about whether or not event E actually occurred. This is precisely where we consider whatever evidence we may have about the believability of source *I*. We may have considerable uncertainty about whether or not event E occurred. All subsequent links in this argument concern the *relevance* of event E to hypothesis H. As we noted in Figure 1.1, these relevance links connect the *idea dots* we discussed. As Figure 1.2 shows, each idea dot is a source of uncertainty associated with the logical connection between whether or not event E did occur and whether or not H is true.

1.1.3 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-related 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 1.3, 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), *deduction*, and *induction* (probabilistic reasoning). These forms of reasoning provide the bases for our idea dots.

1.1.4 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 1.3.

From European sources came word that terrorists of Middle Eastern origin would make new attempts to destroy the World Trade Center, this time

https://doi.org/10.1017/CBO9781316388488.002 Published online by Cambridge University Press

using airliners. Many threats are received every day, most of which come to nothing. However, from several civilian flying schools in the United States came word (to the FBI) that persons from the Middle East were taking flying lessons, paying for them in cash, and wanting to learn only how to steer and navigate heavy aircraft but not how to make takeoffs and landings in these aircraft. By itself, this information, though admittedly strange, may not have seemed very important. But, *taken together*, these two items of information might have caused even an Inspector Lestrade (the rather incompetent police investigator in Sherlock Holmes stories) 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 volume. 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 consider having some number N of evidential dots. We ask the question: How many combinations C of two or more evidential dots are there when we have N evidential dots? The answer is given by the following expression: $C = 2^N - (N + 1)$. 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 one hundred and would increase as time passes. Remember that we live in a nonstationary 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 possibly examine all possible evidential dot combinations* even when N is quite small.

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

https://doi.org/10.1017/CBO9781316388488.002 Published online by Cambridge University Press

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. Figure 1.3 is an abstract example; a tragic real-life example is what happened on September 11, 2001.

In Figure 1.3, there are four numbered evidential dots. The numbers might indicate the order in which we obtained them. In part (a) of the figure, we show an instance where these four dots have been examined separately or independently, in which case they tell us nothing interesting. Then someone notices that, taken together, these four dots combine to suggest a new hypothesis H_k that no one has thought about before, as shown in part (b) 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 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.

It might be said that the act of looking through everything in the hope of finding something is the equivalent of giving yourself a prefrontal lobotomy, meaning that you are ignoring any imaginative capability you naturally have concerning which evidential dot combinations to look for in your analytic problem area. 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 2.3.

1.1.5 How to Connect Evidential Dots to Hypotheses?

As discussed in Section 4.2, all evidence has three major credentials or properties: *relevance*, *believability* or *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

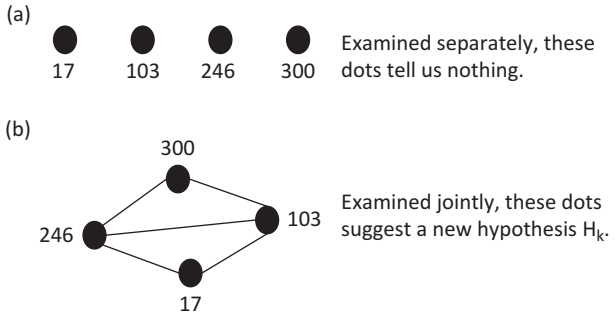

Figure 1.3. Evidential synergism.

https://doi.org/10.1017/CBO9781316388488.002 Published online by Cambridge University Press

previously, *believability* answers the question, "Can we believe what this evidence is telling us?" The 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 1.2, where we show a connection between evidence E^*_i and hypothesis H. An analyst has an item of information from source $\mathcal{I}$ 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 or not. 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 to hypothesis H.

In forming this argument, the analyst wisely begins with the believability foundation for this whole argument: Did event E really occur just because source $\mathcal{I}$ says it did? Also notice in Figure 1.2 that we have indicated the uncertainty associated with each idea dot in this argument. For example, the analyst only infers from E that F might have occurred, and so we note that we must consider F and notF as possibilities. The same is true for the other idea dots G and H.

There are several important things to note about relevance arguments; the first concerns their defense. Suppose the argument in Figure 1.2 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 valuable 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 considered only 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*.

From years of experience teaching law students to construct defensible and persuasive arguments from evidence, we have found that most of them often experience difficulty in

https://doi.org/10.1017/CBO9781316388488.002 Published online by Cambridge University Press

constructing arguments from single items of evidence; they quickly become overwhelmed when they are confronted with argument construction involving masses of evidence. But they gain much assistance in such tasks 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. We will encounter Wigmore's work in several places in our discussions, and you will see that the Disciple-CD system employs elements of Wigmore's methods of argument construction.

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.

1.1.6 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* are our arguments 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* of our arguments based on the combined evidence we have considered. *Direction* refers to the hypothesis we believe our evidence favors most. *Force* 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 1.2 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 over not-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 *believability* of source *I*, and those associated with

https://doi.org/10.1017/CBO9781316388488.002 Published online by Cambridge University Press

links in the analyst's *relevance* argument. So, the force of evidence E^*_i on hypotheses H and not-H depends on how much uncertainty exists in this entire argument involving each one of its believability and relevance links. The interesting message here is that the evidence force or weight credential depends on its other two credentials: believability and relevance.

In the simple example just discussed, there are four major origins of uncertainty, one associated with believability 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 believability 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 10, we consider four systems of probability in our work. 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*, the *Baconian*, and the *Fuzzy* probability systems. But we devote considerable attention to a combination of the Baconian and the 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 Disciple-CD 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 weight or force of evidence are 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 believability of sources of evidence or about the strength of relevance linkages. These disagreements can be resolved only when arguments are made carefully and are openly revealed so that they can be tested by colleagues. A major mission of the Disciple-CD 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 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; that is, 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 Disciple-CD system is designed to allow you and others to evaluate critically the arguments you have constructed.

1.1.7 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 potential 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 previously 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 our 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 who 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 the Disciple-CD process 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?

https://doi.org/10.1017/CBO9781316388488.002 Published online by Cambridge University Press

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.

1.2 IMAGINATIVE REASONING IN INTELLIGENCE ANALYSIS

1.2.1 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 *all of us* 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 Disciple-CD is based on the idea that you are naturally required to exercise your imaginations 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* are the exercises of our imaginations. 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 Disciple-CD system we have developed can only assist you in various ways, and so much 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. For example,

https://doi.org/10.1017/CBO9781316388488.002 Published online by Cambridge University Press

if you knew a fair amount about the behavior of various animal species, you might be able to generate very useful metaphors for characterizing the behavior of terrorists.

One of the most difficult problems we have faced in our work on Disciple-CD is assisting you to construct defensible and persuasive arguments from a *mass* of evidence supporting or challenging 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 we mentioned previously in this chapter, 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 Disciple-CD 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 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* works or courses 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 Disciple-CD system provides assistance in performing these complex tasks.

1.2.2 What Ingredients of Analysis Are to Be Generated by Imaginative Thought?

If we place such a premium on your imaginative reasoning, we ought to be able to tell you precisely what elements of intelligence analyses need to be generated or discovered and, if possible, how these activities might best be assisted. Figure 1.4 describes the major ingredients of intelligence analysis that result from *imaginative reasoning* coupled with *critical reasoning*.

https://doi.org/10.1017/CBO9781316388488.002 Published online by Cambridge University Press

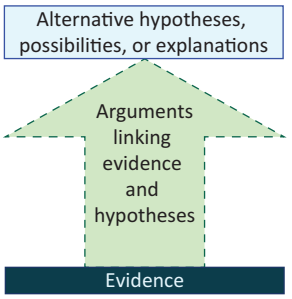

Figure 1.4. Major ingredients of intelligence analysis.

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 volume, you must also establish the relevance, believability, 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, 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.

1.2.3 Generating Main Hypotheses to Be Defended by Evidence and Argument

In any intelligence analysis, you will have to draw some kind of a conclusion. The *possible conclusions* you might draw can be in the form of *main hypotheses*. In most cases, these 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. Generally, our main 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:

https://doi.org/10.1017/CBO9781316388488.002 Published online by Cambridge University Press

Example 1.4.

(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.

Example 1.5.

(Hypotheses concerning an event that may be happening "now") You might have reason to suspect that Country Z is still holding prisoners of war (POWs) taken years ago during a conflict we had with it. Your suspicion here forms one hypothesis:

H_5: Country Z 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 "not H_5":

H_6: Country Z is not holding any of our POWs.

Example 1.6.

(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 "not H_7":

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 the conflict, how will it proceed, how long will it last, and who will win.

All of these examples concern events/situations that *might have happened*, are *now possibly happening*, or *might happen in the future*. 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, evidence and arguments, come in.

1.2.4 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 or not. Here is an example of its importance.

https://doi.org/10.1017/CBO9781316388488.002 Published online by Cambridge University Press

Take any of the three situations just mentioned in Section 1.2.3 concerning hypotheses about either past, current, or future events:

H_1: Person X was the one involved in this incident.
H_5: Country Z 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 now 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 are some examples of evidence we might find concerning the preceding 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 Z shows us a dog-tag he says was given to him by a resident of Z. On this tag is the name of a soldier who has been missing since our conflict with Country Z ended.
For H_7: We might obtain evidence bearing upon the state of military preparedness of either country.

1.2.5 Generating Arguments Linking Evidence and Hypotheses

The third major ingredient of intelligence analysis concerns the arguments we must construct in defense of the relevance, believability or 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 1.5 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_1 (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 link between the evidence and our hypothesis.

https://doi.org/10.1017/CBO9781316388488.002 Published online by Cambridge University Press

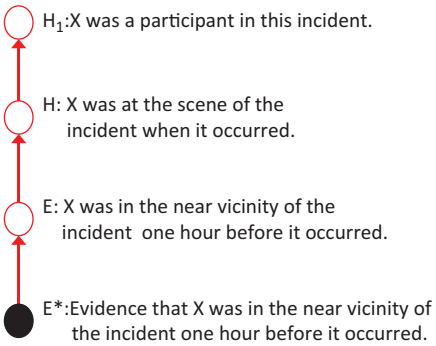

Figure 1.5. Sample argument.

However, all the argument in Figure 1.5 shows is that X *might have been* a participant in the incident. Remember from Section 1.2.3 that we were considering three other hypotheses, in addition to this one. Therefore, what we would like to know is which of the four is most likely. This would require the analysis of all these hypotheses, and not just based on a single item of evidence. It would also require an assessment of how likely each hypothesis is, based on the relevance, the believability, and the inferential force of evidence.

The next section introduces a systematic approach to intelligence analysis that is based on the scientific method and is supported by the Disciple-CD system.

1.3 INTELLIGENCE ANALYSIS AS DISCOVERY OF EVIDENCE, HYPOTHESES, AND ARGUMENTS

1.3.1 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 nonstationary world, involving collaborative processes of evidence in search of hypotheses, hypotheses in search of evidence, and evidentiary testing of hypotheses, as represented in Figure 1.6.

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, as will be illustrated with the use of the Disciple-CD cognitive assistant.

Through *abductive reasoning* (Peirce, 1992 [1898]; 1995 [1901]; Schum, 2001b) (which shows that something is *possibly* true), the analyst and Disciple-CD generate alternative hypotheses that explain their observations (see the left-hand side of Figure 1.6). Through *deductive reasoning* (which shows that something is *necessarily* true), they use these hypotheses to generate new lines of inquiry and discover new evidence (see the middle of Figure 1.6). And through *inductive reasoning* (which shows that something is *probably* true), they test each of these hypotheses with the discovered evidence and select the most likely one (see the right-hand side of Figure 1.6).

The following sections illustrate this systematic approach to intelligence analysis by using a specific example of anticipatory analysis where evidence about a canister of cesium-137 missing from a company leads to anticipating the fact that a dirty bomb will

https://doi.org/10.1017/CBO9781316388488.002 Published online by Cambridge University Press

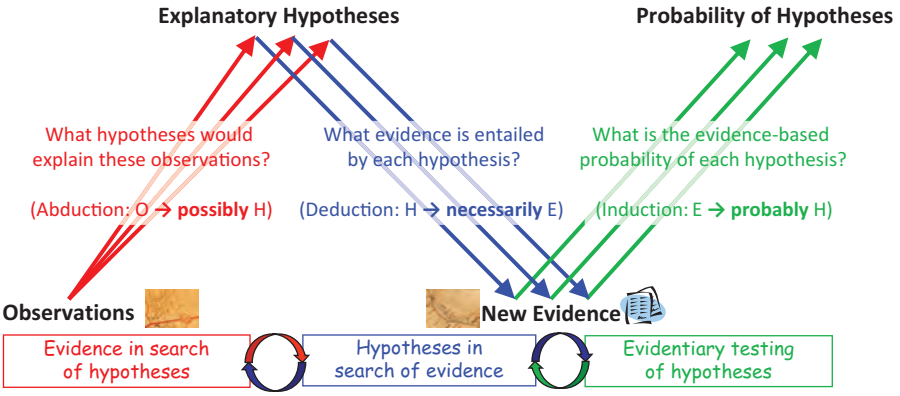

Figure 1.6. Framework of the computational theory of intelligence analysis.

be set off in the Washington, D.C., area. At the same time, this example will introduce the main concepts related to evidence and inference, which will be detailed and experimented throughout the rest of this book.

1.3.2 Evidence in Search of Hypotheses

Consider that you are an intelligence analyst and you read in today's *Washington Post* an article 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 United States, 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 XYZ Company in Maryland just three days ago. The XYZ Company manufactures devices for sterilizing medical equipment and uses cesium-137 in these devices along with other radioactive materials. This piece arouses your curiosity because of your concern about terrorists planting dirty bombs in our cities. The question is, "What hypotheses would explain this observation?" You experience a flash of insight that a dirty bomb may be set off in the Washington, D.C., area (see Figure 1.7).

However, no matter how imaginative or important this hypothesis is, no one will take it seriously unless you are able to justify it. So you develop the chain of abductive inferences (Peirce, 1992 [1898]; 1995 [1901]; Schum 2001b) shown in Table 1.1 and in Figure 1.8.

The chain of inferences from Table 1.1 and Figure 1.8 shows clearly the possibility that a dirty bomb will be set off in the Washington, D.C., area. Can you then conclude that this will actually happen? No, because there are many other hypotheses that may explain this evidence, as shown in Figure 1.9 and discussed in the following text.

Just because there is evidence that the cesium-137 canister is missing does not mean that it is indeed missing. At issue here is the believability of Willard, the source of this information. What if Willard is mistaken or deceptive? Thus, an alternative hypothesis is that the cesium-137 canister is not missing.

But let us assume that the cesium-137 canister is indeed missing. Then it is possible that it was stolen. But it is also possible that it was misplaced, or maybe it was used in a project at the XYZ Company without being checked out from the warehouse.

https://doi.org/10.1017/CBO9781316388488.002 Published online by Cambridge University Press

Figure 1.7. Hypothesis generation through imaginative reasoning.

Table 1.1 Abductive Reasoning Steps Justifying a Hypothesis

There is evidence that the cesium-137 canister is missing (E*).
Therefore it is possible that the cesium-137 canister is indeed missing (H_1).
Therefore it is possible that the cesium-137 canister was stolen (H_2).
Therefore it is possible that the cesium-137 canister was stolen by someone associated with a terrorist organization (H_3).
Therefore it is possible that the terrorist organization will use the cesium-137 canister to construct a dirty bomb (H_4).
Therefore it is possible that the dirty bomb will be set off in the Washington, D.C., area (H_5).

However, let us assume that the cesium-137 canister was indeed stolen. It is then possible that it might have been stolen by a terrorist organization, but it is also possible that it might have been stolen by a competitor or by an employee, and so on.

This is the process of *evidence in search of hypotheses*, shown in the left-hand side of Figure 1.6. We cannot conclude that a dirty bomb will be set off in the Washington, D.C., area (i.e., hypothesis H_5) until we consider all the alternative hypotheses and show that those on the chain from E* to 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 1.6, we put each hypothesis at work to guide us in the collection of additional evidence. This process is discussed in the next section.

1.3.3 Hypotheses in Search of Evidence

Let us first consider the hypothesis "H_1: missing" from near the bottom of Figure 1.9, shown as "H_1: cesium-137 canister is missing from the warehouse" in the top-left of

https://doi.org/10.1017/CBO9781316388488.002 Published online by Cambridge University Press

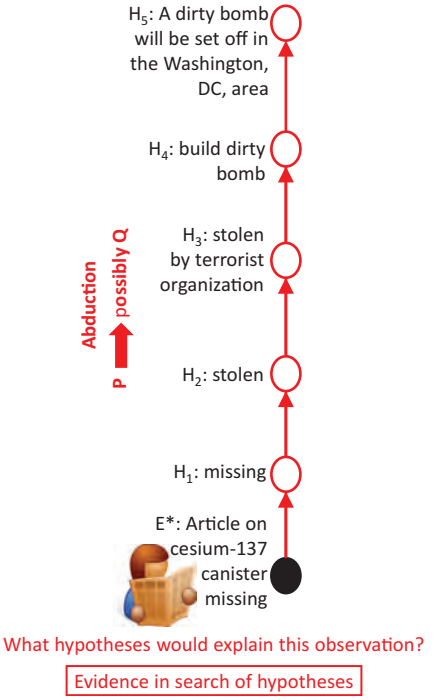

Figure 1.8. Justification of the generated hypothesis.

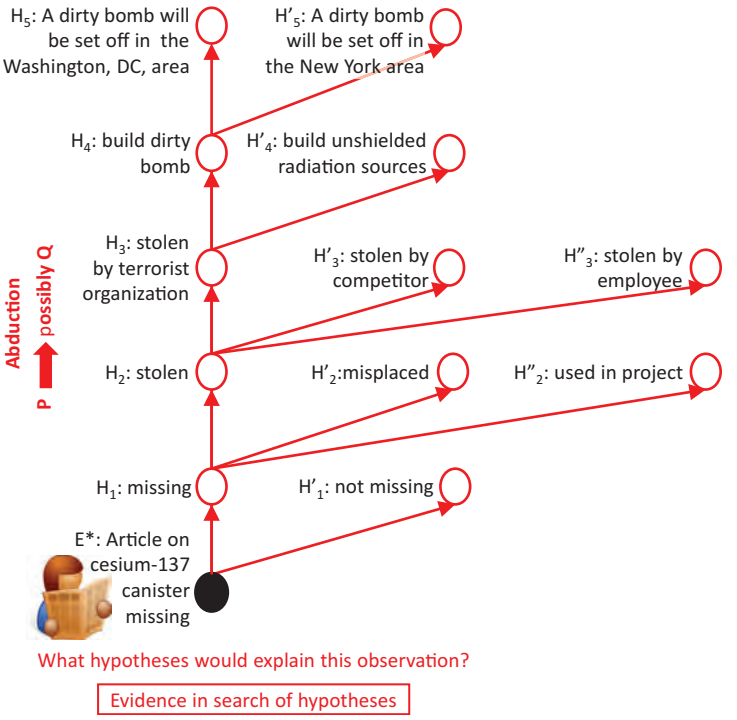

Figure 1.9. Competing hypotheses explaining an item of evidence.

https://doi.org/10.1017/CBO9781316388488.002 Published online by Cambridge University Press

Figure 1.10. 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 checked it out.

This suggests the decomposition of the hypothesis H_1 into three simpler hypotheses, as shown in the left part of Figure 1.10. 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. That is, by putting hypothesis H_1 to work, you were guided to perform the collection tasks from Table 1.2, represented in Figure 1.10 by the gray circles.

Guided by the evidence collection tasks in Table 1.2, you contact Ralph, the supervisor of the XYZ warehouse, who provides the information shown in Table 1.3 and in Figure 1.10.

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 Willard in his *Washington*

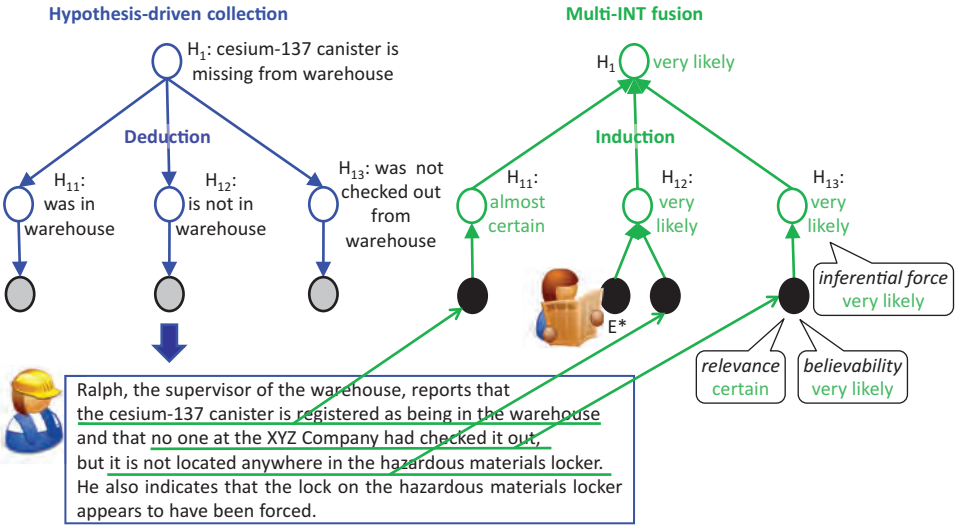

Figure 1.10. Hypothesis-driven evidence collection and hypothesis testing.

Table 1.2 Evidence Collection Tasks Obtained from the Analysis in Figure 1.10

Collection Task1: Look for evidence that the cesium-137 canister was in the XYZ warehouse before being reported as missing.
Collection Task2: Look for evidence that the cesium-137 canister is no longer in the XYZ warehouse.
Collection Task3: Look for evidence that the cesium-137 canister was not checked out from the XYZ warehouse.

https://doi.org/10.1017/CBO9781316388488.002 Published online by Cambridge University Press

Table 1.3 Information Obtained through the Collection Tasks in Table 1.2

INFO-002-Ralph: Ralph, the supervisor of the warehouse, reports that the cesium-137 canister is registered as being in the warehouse and that no one at the XYZ 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.

Table 1.4. Dots or Items of Evidence Obtained from Willard and Ralph

E001-Willard: Willard's report in the *Washington Post* that a canister containing cesium-137 was missing from the XYZ warehouse in Baltimore, MD.
E002-Ralph: Ralph's testimony that the cesium-137 canister is registered as being in the XYZ warehouse.
E003-Ralph: Ralph's testimony that no one at the XYZ Company had checked out the cesium-137 canister.
E004-Ralph: Ralph's testimony that the canister is not located anywhere in the hazardous materials locker.
E005-Ralph: Ralph's testimony that the lock on the hazardous materials locker appears to have been forced.

Post article. We parse it to extract the relevant information represented as E001-Willard in Table 1.4. Similarly, Ralph's testimony from Table 1.3 provides us with several dots or items of evidence that are relevant to assessing the hypotheses from Figure 1.10. These items of evidence are represented in Table 1.4.

This is the process of *hypothesis in search of evidence* that guides us in collecting new evidence. The next step now is to assess the *probability or likeliness* of hypothesis H_1 based on the collected evidence, as represented in the right-hand side of Figure 1.6 and discussed in the next section.

1.3.4 Evidentiary Testing of Hypotheses

Having identified evidence relevant to the hypotheses in Figure 1.10, the next step is to use it in order to assess these hypotheses. The assessments of the hypotheses will be done by using probabilities that are expressed in words rather than in numbers. In particular, we will use the ordered symbolic probability scale from Table 1.5. This is based on a combination of ideas from the Baconian and Fuzzy probability systems. As in the Baconian system, "no support" for a hypothesis means that we have no basis to consider that the hypothesis might be true. However, we may later find evidence that may make us believe that the hypothesis is "very likely," for instance.

To assess the hypotheses, we first need to attach each item of evidence to the hypothesis to which it is relevant, as shown in the right-hand side of Figure 1.10. Then we need to establish the *relevance* and the *believability* of each item of evidence, which will result in the *inferential force* of that item of evidence on the corresponding hypothesis, as illustrated at the right-hand side of Figure 1.10 and explained in the following.

https://doi.org/10.1017/CBO9781316388488.002 Published online by Cambridge University Press

Table 1.5 Ordered Symbolic Probability Scale

no support < likely < very likely < almost certain < certain

So let us consider the hypothesis "H_{13}: cesium-137 canister was not checked-out from the warehouse" and the item of evidence "E003-Ralph: Ralph's testimony that no one at the XYZ Company had checked out the cesium-137 canister."

Relevance answers the question, "So what? How does E003-Ralph bear on the hypothesis H_{13} that we are trying to prove or disprove?" If we believe what E003-Ralph is telling us, then H_{13} is "certain."

Believability answers the question, "To what extent can we believe what E003-Ralph is telling us?" Let us assume this to be "very likely."

Inferential force or weight answers the question, "How strong is E003-Ralph in favoring H_{13}?" Obviously, an item of evidence that is not relevant to the considered hypothesis will have no inferential force on it and will not convince us that the hypothesis is true. An item of evidence that is not believable will have no inferential force either. Only an item of evidence that is both very relevant and very believable will convince us that the hypothesis is true. In general, the inferential force of an item of evidence (such as E003-Ralph) on a hypothesis (such as H_{13}) is the minimum of its relevance and its believability. We can therefore conclude that, based on E003-Ralph, the probability of the hypothesis H_{13} is "very likely" (i.e., the minimum of "certain" and "very likely"), as shown in Figure 1.10.

Notice in Figure 1.10 that there are two items of evidence that are relevant to the hypothesis H_{12}. In this case, the probability of H_{12} is the result of the combined (maximum) inferential force of these two items of evidence.

Once we have the assessments of the hypotheses H_{11}, H_{12}, and H_{13}, the assessment of the hypothesis H_1 is obtained as their minimum, because these three subhypotheses are necessary and sufficient conditions. Therefore, all need to be true in order for H_1 to be true, and H_1 is as weak as its weakest component.

Thus, as shown at the top-right side of Figure 1.10, we conclude that it is "very likely" that the cesium-137 canister is missing from the warehouse.

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

Figure 1.11 summarizes the preceding analysis, which is an illustration of the general framework from Figure 1.6.

Now that we have concluded "H_1: missing," we repeat this process for the upper hypotheses (i.e., H_2: stolen; H'_2: misplaced; and H''_2: used in project), as will be discussed in the next section.

1.3.5 Completing the Analysis

Let us first consider the hypothesis "H_2: stolen." We need to put this hypothesis to work to guide us in collecting relevant evidence for its analysis. During our investigation of the security camera of the warehouse, we discover a video segment showing a person loading a container into a U-Haul panel truck. This new item of evidence, together with Ralph's testimony that the lock on the hazardous materials locker appears to have been

https://doi.org/10.1017/CBO9781316388488.002 Published online by Cambridge University Press

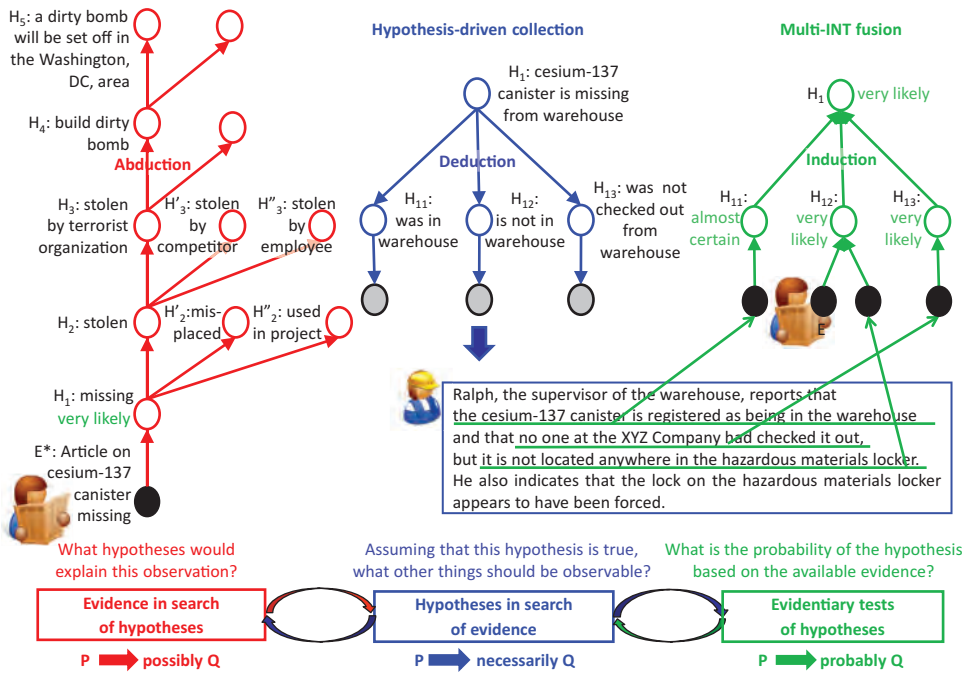

Figure 1.11. An illustration of the general framework from Figure 1.6.

forced (E005-Ralph in Table 1.4), suggests the following scenario of how the cesium-137 might have been stolen (see Figure 1.12): *The truck entered the company, the canister was stolen from the locker, the canister was loaded into the truck, and the truck left with the canister.*

Such scenarios have enormous heuristic value in advancing the investigation because they consist of mixtures of what is taken to be factual and what is conjectural. Conjecture is necessary in order to fill in natural gaps left by the absence of existing evidence. Each such conjecture, however, opens up new avenues of investigation, and the discovery of additional evidence, if the scenario turns out to be true. For instance, the first hypothesized action from the scenario ("Truck entered company") leads us to check the record of the security guard, which shows that a panel truck bearing Maryland license plate number MDC-578 was in the XYZ parking area the day before the discovery that the cesium-137 canister was missing.

The second hypothesized action in the scenario (i.e., "cesium-137 canister stolen from locker") is further decomposed into two hypotheses. The first one was already analyzed, "It is very likely that the cesium-137 canister is missing from the warehouse." The second subhypothesis ("Warehouse locker was forced") is supported both by Ralph's testimony (i.e., E005-Ralph in Table 1.4) and by the professional locksmith, Clyde, who was asked to examine it (E007-Clyde: Professional locksmith Clyde testimony that the lock has been forced, but it was a clumsy job).

Fusing all the discovered evidence, Disciple-CD concludes that it is very likely that the cesium-137 canister was stolen with the MDC-678 truck.

We repeat the same process for the other two competing hypotheses, H'_2: misplaced, and H''_2: used in project. However, we find no evidence that the cesium-137 canister might have been misplaced. Moreover, we find disfavoring evidence for the second

https://doi.org/10.1017/CBO9781316388488.002 Published online by Cambridge University Press

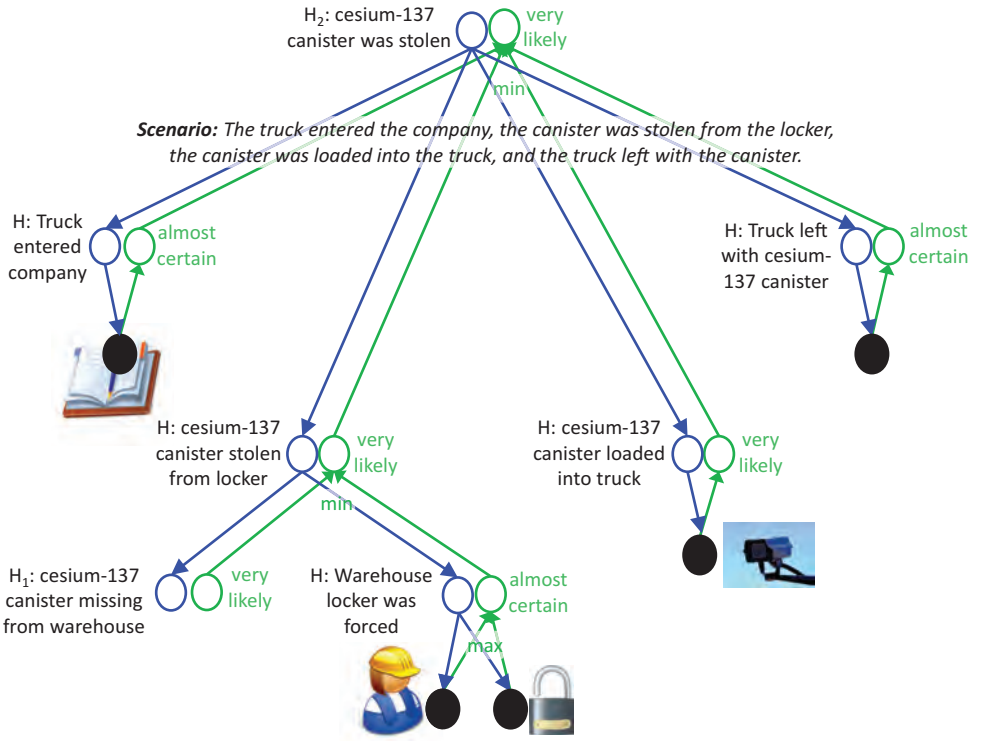

Figure 1.12. Another example of hypothesis-driven evidence collection and hypothesis testing.

competing hypothesis: Grace, the Vice President for Operations at XYZ, tells us that no one at the XYZ Company had checked out the canister for work on any project.

Thus we conclude that the cesium-137 canister was stolen and we continue our analysis with investigating the next level up of competing hypotheses: H_3: stolen by terrorist organization; H'_3: stolen by competitor; and H''_3: stolen by employee. Of course, at any point, the discovery of new information may lead us to refine our hypotheses, add new hypotheses, or eliminate existing hypotheses.

This example is not as simple as it may be inferred from this presentation. It is the methodology that guides you and makes it look simple. Many things can and will indeed go wrong. But the computational theory of intelligence analysis and Disciple-CD provide you the means to deal with any problems. Based on evidence, you come up with some hypotheses, but then you cannot find evidence to support any of them. So you need to come up with other hypotheses, and you should always consider alternative hypotheses. The deduction-based decomposition approach guides you on how to look for evidence, but your knowledge and imagination also play a crucial role. As illustrated here, we imagined a scenario where the cesium-137 canister was stolen with a truck. But let us now assume that we did not find supporting evidence for this scenario. Should we conclude that the cesium-137 canister was not stolen? No, because this was just one scenario. If we can prove it, we have an assessment of our hypothesis. However, if we cannot prove it, there still may be another scenario explaining how the cesium-137 canister might have been stolen. Maybe the cesium canister was stolen by someone working at the XYZ Company. Maybe it was stolen by Ralph, the administrator of the warehouse. The important thing is that each such scenario opens a new line of investigation and a new way to prove the hypothesis.

https://doi.org/10.1017/CBO9781316388488.002 Published online by Cambridge University Press

The next chapters of this book include exercises for completing this analysis that will further illustrate the synergistic integration of an analyst's imagination with a computer's critical reasoning. Having established that cesium-137 canister was stolen, we would further like to determine by whom and for what purpose. If it is for constructing and setting off a dirty bomb, we would like to know who will do this, where in the Washington, D.C., area the bomb will be set off, precisely when this action will happen, what form of dirty bomb will be used, and how powerful it will be. These are very hard questions that the computational theory of intelligence analysis presented in this book (as well as its current implementation in Disciple-CD) will help to answer.

One major challenge in performing such an analysis is the development of argumentation structures. An advantage of using an advanced analytic tool such as Disciple-CD is that it can learn reasoning patterns from the analyst to greatly facilitate and improve the analysis of similar hypotheses, as will be shown in the next chapters of this book.

In conclusion, the computational theory of intelligence analysis presented in this volume, as well as its current implementation in Disciple-CD, provides a framework for integrating the art and science of intelligence analysis to cope with its astonishing complexity.

However, while the computational theory and Disciple-CD guide you through the intelligence analysis steps, and also automates many of them, it requires you to continuously exercise your imagination. Therefore, in Chapter 2, 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 such as Disciple-CD is the approach we advocate for coping with the astonishing complexity of "connecting the dots."

1.4 REVIEW QUESTIONS

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

1.2. 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 main hypothesis about possible suppliers of these IEDs. What are some of these other hypotheses?

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

1.4. Consider the hypothesis that Al Qaeda Group A from Yemen was the one involved in the terrorist incident. What evidence we might find concerning this hypothesis?

1.5. 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?

1.6. 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

https://doi.org/10.1017/CBO9781316388488.002 Published online by Cambridge University Press

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?

1.7. 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?

1.8. Defendant Dave is accused of shooting a victim, Vic. When Dave was arrested sometime after the shooting, he was carrying a 32-caliber 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. Construct a simple chain of reasoning that connects Frank's report to the hypothesis that it was Dave who shot Vic.

1.9. Consider the situation from Question 1.8. 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?

1.10. Consider again the situation from Questions 1.8 and 1.9. In order to prove the hypothesis that it was Dave who shot Vic, we need additional evidence. As discussed in Section 1.3.3, we need to put this hypothesis to work to guide us in collecting new evidence. Decompose this hypothesis into simpler hypotheses, as was illustrated by the blue trees in Figures 1.11 and 1.12.

1.11. Our investigation described in Questions 1.8, 1.9, and 1.10, 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 Dave was carrying his weapon, and with the ballistics evidence puts additional heat on Dave. Extend the decomposition tree from Question 1.10 with assessments of the probability of the hypotheses, as was illustrated by the green trees in Figures 1.11 and 1.12. In Section 4.3, we will discuss more rigorous methods for making such probabilistic assessments. In this exercise, just use your common sense.

1.12. A car bomb was set off in front of a power substation in Washington, D.C., on November 25. 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 Budget Car Rental Agency. From information provided by Budget, it was learned that the car was last rented on November 24 by a man named M. Construct an argument from this evidence to the hypothesis that Person M was involved in this car-bombing incident.

1.13. Consider the situation from Question 1.12 and the corresponding argument. Suppose that we have determined that evidence E* is believable and therefore we think that M indeed rented a car on November 24. We want now to assess F, whether M drove the car on November 25. For this we need additional evidence. As discussed in Section 1.3.3, we need to put this hypothesis to work to guide us in collecting new evidence. Decompose this hypothesis into simpler hypotheses, as was illustrated by the blue trees in Figures 1.11 and 1.12.