# 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.*

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.

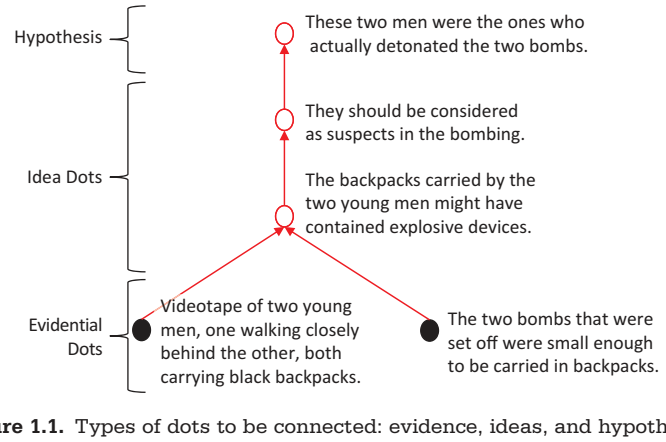


**Figure 1.1.** Types of dots to be connected: evidence, ideas, and hypotheses.

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.

**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

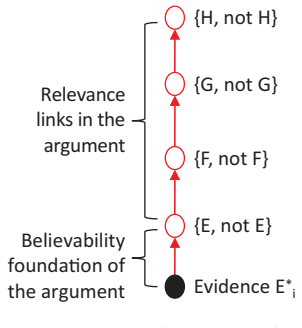


**Figure 1.2.** The believability foundation for an argument.

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

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

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

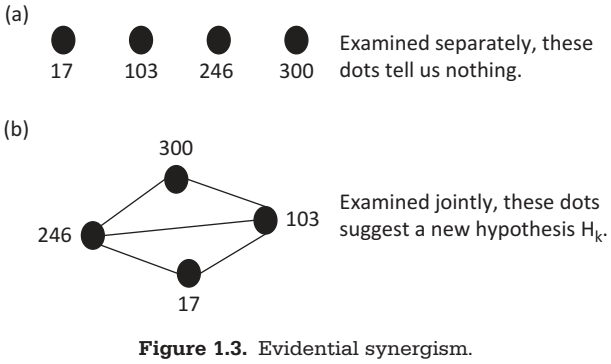


**Figure 1.3.** Evidential synergism.

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 $I$ 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 $I$ 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

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

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