To appear in the Proceedings of the 1989 Joint Director of Laboratories C^2 Symposium, Ft. McNair, VA, June 27-29, 1989.

A DATABASE TOOL TO SUPPORT PROBABILISTIC ASSUMPTION-BASED REASONING IN INTELLIGENCE ANALYSIS

Marvin S. Cohen

Decision Science Consortium, Inc.
Reston, VA 22091

ABSTRACT

The Self-Reconciling Evidential Database (SED) is a tool for intelligence analysts that combines a numerical uncertainty calculus with a process of higher-order reasoning about knowledge and assumptions. SED includes (1) a natural representation of evidential arguments in terms of a normal or first-blush reaction to the evidence plus a set of exception conditions, (2) a modeling technique that drastically reduces the number of assessments required to build complex arguments, and (3) a process of resolving conflict among competing arguments by examining and revising assumptions that led to the conflict rather than by statistically aggregating.

1. INTRODUCTION

The Problem

It is far easier to diagnose the reasons for an intelligence failure after the fact than it is to prevent one beforehand. Success or failure seems to hinge on analysis--noticing significant data in a background of noise, assessing their reliability, or finding a pattern that fills in gaps and resolves inconsistencies--as often as it does on the collection of data *per se* (cf., Laqueur, 1985; Burrows, 1986). Yet there are no easy prescriptions for these tasks:

- An analyst should avoid "biases" but must also draw effectively upon knowledge of the topic and area. That knowledge (if it is useful) will certainly predispose the analyst toward some hypotheses and away from others.
- Almost any data may mean something other than what they seem, due to deception. Sensitivity to the possibility of deception, however, can lead to disregard of genuine evidence.
- Involvement with policy makers may, on occasions, lead to interpretative errors--e.g., a "Cassandra" attitude (worst-case) or the opposite, "Pollyanna." Yet isolation from policy makers may lead to irrelevance and/or gaps in coverage.

The answer, it is easy to say, lies in *balance*: between attention to theory and respect for evidence; between extensive substantive knowledge and being ready, if necessary to question the assumptions embedded in it; and finally, between divergent and convergent modes of thought--generating and taking seriously alternative possibilities, even comparing and contrasting alternative models and types of analysis, and yet in the end offering a reasonable (and reasonably definitive) conclusion.

The problem, of course, is how to achieve such balance in practical terms. Few would claim that currently available tools supply all the help that is needed. Specialized techniques (e.g., critical-indicators analysis, throw-weight analysis, "crate-ology") do not address the general problem of combining evidence and analyses of diverse types. General-purpose tools (e.g., database systems, spreadsheets, hypercard), though useful, have little to offer that bears explicitly on the distinctive problems of inference. The most promising source of help may lie in technologies for handling uncertainty that have been introduced by statisticians and by expert system builders. Yet these suffer from a variety of drawbacks:

- The meaning of numerical assessments is often unclear, and numerical representations of inferential arguments often seem unnatural.
- An enormous number of assessments is required even in simple problems.
- Standard inference methods respond inadequately both to the challenge of stimulating alternative points of view ("divergence") and to the requirement of resolving them in a meaningful fashion ("convergence"). Computerized systems are not "intelligent" enough to sustain the kind of balance that the analyst must achieve.

Until all three of these problems are addressed, computerized aids for intelligence analysis are likely to be too confusing, too incomplete, and too inflexible.

A New Approach: Basic Concepts

The present report describes a system that addresses these problems directly. SED (Self-Reconciling Evidential Database) brings together aspects of two approaches: (1) symbolic techniques for structuring arguments and for the adoption, utilization, and revision of assumptions; and (2) mathematical techniques for combining and propagating the impact of evidence. The result, we hope, is not just a hybrid, but a deeper synthesis: a system that is both compatible with the way analysts would naturally approach a problem and at the same time likely to yield improvements. In brief, SED has the following features:

- a natural approach to argument construction that includes both an initial "automatic" response to evidence and a capability for drawing on more detailed and flexible models where appropriate;
- a method for creating complex numerical arguments that avoids the usual combinatorial explosion and requires only a small number of simple assessments; and

- **a capability not only to *use* arguments in reasoning, but to reason *about* the arguments and to revise them in light of their performance. While it has been customary to regard numerical calculi and assumption-based reasoning as competing methods for handling uncertainty, SED associates numerical arguments with the assumptions upon which they depend; conflict among different pieces of evidence is resolved not by blind statistical integration but by examination of the assumptions that led to the conflict.**

We will briefly describe each of these features in turn. The current SED prototype operates on an IBM PC/AT desktop computer. It utilizes the most recent version of an inference system called the Non-Monotonic Probabilist (Cohen 1986; Cohen, Laskey, and Ulvila, 1987), which combines aspects of both numerical and non-numerical approaches to uncertainty. NMP is implemented in Golden Common LISP by means of the Belief Maintenance System described by Laskey and Lehner (in press). A more detailed description of SED may be found in Cohen, Laskey, Vane, McIntyre, and Sak (1989). A discussion of different concepts of uncertainty and a theoretical rationale for SED may be found in Cohen, Laskey, and Ulvila, 1987.

2. BUILDING ARGUMENTS

At the highest level, SED organizes information by issues, i.e., topics, questions about those topics, and potential answers: e.g.,

ISSUE	TOPIC	QUESTION	ANSWERS
#1	Krasnoyarsk radar	What is its function?	Local defense Early warning Space tracking Other non-ABM
#2	Krasnoyarsk radar	Will the Soviets agree to dismantle it?	Yes No
#3	Soviet supersonic aircraft	How many have been delivered to Latin America?	0 1-10 51-100 100>
#4	Columbian heroin	What will be its foreign exchange value (in current US $) in 5 years	<$1M $2M-$10M $11M-$20M $20M-$100M $100M>

At the lowest level, SED organizes information by reports, i.e., evidence from satellites, informants, open sources, etc. Arguments, which link reports to issues, and issues to other issues, are the heart of SED. Each argument supports a particular position on an issue: e.g.,

TOPIC	QUESTION
Krasnoyarsk radar	What is its function?

ANSWERS

Local defense	
Early warning	*
Space tracking	*
Other Non-ABM	*
Support =	1.0

Here, the evidence demonstrates that the radar's function is *not* local defense, i.e., it is either early warning, space tracking, or some other non-ABM purpose; but the available evidence is unable to discriminate further among these possibilities.

In SED, the analyst is encouraged to state the *reasons* why a given conclusion might (or might not) follow from a particular piece of evidence--not simply a number measuring the degree to which the conclusion is associated with that evidence. Belief regarding an issue is always determined by one or more arguments.

A key feature of SED's approach is the phasing of argument construction to fit the natural stages of an analyst's reasoning: i.e., a "first-blush" or "normal" reaction to the evidence (which we call a "core position") is followed by specification of a set of possible disrupting factors. For example, photographic evidence that there are no significant military bases or other assets near the Krasnoyarsk radar would *normally* suggest that its function is not to support a local ABM defense, since there are no assets to protect. But this inference fails if (i) assets are planned, but not yet built, (ii) assets consist of natural resources or some other non-man-made feature, (iii) assets are camouflaged or buried, (iv) the function of existing structures has been concealed, (v) the photo analysis was badly done, etc. Typically, these exception conditions are assumed false in the absence of direct evidence one way or the other, until and unless the "normal" interpretation of the evidence runs into trouble (i.e., conflicts with the position supported by some other line of reasoning). SED thus focuses attention on an evolving understanding of the qualitative meaning and reliability of evidence, as opposed to cut-and-dried numerical assessments of evidence strength.

The construction of an argument in SED falls into natural phases:

- Step (1) is an initial face-value or "normal" interpretation of the evidence. It consists merely of specifying the evidence and a core position on the focal issue that seems to follow from it. For example, the first-blush meaning of the satellite photograph showing no assets near Krasnoyarsk is that the function of the radar is not local defense:

POSITION — CORE — Argument 1 of 2

	Krasnoyarsk radar	What is its function?	
Local defense			I
Early warning	*		
Space tracking	*		
Other non-ABM	*		
Core Support	1.0		II

PREMISE — 1 of 10

	Satellite Photo	Are there assets near Krasnoyarsk?	III
Yes			
No	*		

CORE

The ARGUMENTS screen, as illustrated here, is divided into three parts:

I. The focal topic, question, and answers;

II. Numerical measures for subsets of answers to the focal issue; and

III. Topic, question and answers for a premise.

- Step (2), involves fleshing out the core argument with a set of background premises. Background premises are necessary for the normal linkage between the evidence and the core position, even though they may have little or no relevance to the focal issue taken by themselves. For example, if there were a Soviet plan to build assets near Krasnoyarsk (e.g., a large military base), then the failure to observe current assets would lose its significance. Thus, we have as a premise in the core argument, the proposition that no new assets are planned:

```
—[POSITION]—CORE————Argument 1 of 2—

              Krasnoyarsk    What is its
              radar          function?

  Local defense
  Early warning     *
  Space tracking    *
  Other non-ABM     *

  Core Support      1.0

—[PREMISE]——————————————————2 of 10—

              Krasnoyarsk    Are assets
              radar          planned but
                             not built?

  Yes
  No                *

________________CORE________________________
```

Divergent Reasoning. In intelligence analysis, as in any inferential activity, there is sometimes a tendency to overlook potential weaknesses or sources of uncertainty in a favored hypothesis. In fact, experimental data, with experienced intelligence officers performing realistic intelligence tasks, suggest that apparently disconfirming evidence may be disregarded or even construed as supporting an initial hypothesis (Tolcott, et al., 1987). SED counteracts this tendency by focusing attention on the ways in which an argument could go wrong.

There are a variety of questions the analyst can ask himself to stimulate generation of exception conditions. The simplest is: "Under what conditions would this argument be valid?" or "What else must be the case for this position to actually follow from this evidence?" A more powerful method is a technique that we call Conflict Resolution (Cohen, 1989; Phillips in IPL/AMRD, 1982). The analyst is asked to suppose hypothetically, that the core position is *not* true and explains how that could be. An effective trick is to imagine that an infallible crystal ball says the core position is false even though the evidence is true. Typically, the analyst will then be able to generate an explanation: e.g., the core position could be false even if the evidence is true, if Q_1 is the case. The crystal ball now tells him that the core position is false and the evidence is true, but Q_1 is also false! As a result, the analyst devises a new explanation, Q_2. Again, the crystal ball tells him Q_2 is false; and so on. The analyst is thus prompted to act as his own Devil's Advocate, exposing hidden assumptions and exploring alternative points of view. The result is typically a long (and sometimes surprising) series of qualifications on the original argument: e.g., report R means position P *unless* qualification-1, qualification-2, etc.

Experiments with this technique in interviews with Army intelligence officers show that it produces a rich harvest of unexpected information. It was not unusual, for example, to obtain numerous additional argument premises by means of the "crystal ball" after more direct questioning of an analyst had run completely dry. In one instance, after assessing the probability of a conclusion as 1.0, an analyst was able (by means of the crystal ball) to generate 8 different exception conditions with an average assessed probability of .31.

In generating exception conditions, analysts must rely less and less on automatic responses, or easily accessible knowledge, and begin to open up "compartments" of knowledge that are not part of their ordinary reaction to the situation. They must become increasingly detailed in their examination of the causal or analytical processes that link evidence and conclusion, if they are to continue adding to the list of exception conditions in which those processes break down. A further stimulus in this process, therefore, is for the analyst to make explicit (perhaps in graphic form) the causal or analytical models underlying an argument. The crystal ball technique can be applied in turn to each *state* of the causal or analytical process. Moreover, charts of this sort can evolve into generic models that underlie a variety of related inferences.

- Step (3) simply adds an account of what happens to the position supported by the argument when a background premise is false:

```
—[POSITION]—CORE——EXCEPTIONS—Argument 1 of 2—

            Krasnoyarsk                What is its
            radar                      function?

  Local defense            0
  Early warning    *       0
  Space tracking   *       0
  Other non-ABM    *       0

  Core Support     1.0

—[PREMISE]————————————————————————2 of 10—

            Krasnoyarsk                Are assets
            radar                      planned but
                                       not built?

  Yes                      *
  No               *

_______________CORE___EXCEPTIONS_______________
```

The star next to "Yes" corresponds to the exception condition: new assets *are* planned. Above it in the same column, circles represent the impact of that exception on the position supported by the argument. In this case, circles are next to all four possible answers. If new assets were planned, the function of local defense could no longer be excluded, and the evidence could no longer discriminate among any of the hypotheses.

3. ECONOMY OF ASSESSMENTS

Qualitative Judgments

In a SED argument, the core position is true if all the premises are true. But what position on the focal issue is supported if one or more of the premises are false? On the face of it, this would seem to place an inordinate assessment burden on the analyst. For an argument with n premises, there are on the order of 2^n combinations of truth and falsity of the premises (if they are binary) for each of which a position on the focal issue would have to be specified. In a Bayesian model that conditions one variable on multiple other variables, a probability must be assessed for every value of the first variable conditional on every possible *combination* of values of the other variables (e.g., Pearl, 1986). Problems can sometimes be structured so as to insulate some variables from the influence of some others, but substantial gains in economy are by no means guaranteed (cf., remarks by Schum, 1980). One way to reduce the assessment burden, of course, is to reduce the number of variables that are included. The exponential growth in required assessments is perhaps a major reason why most approaches to inference do not actively encourage, as SED does, the process of making background variables or exception conditions explicit. As a result, however, the reasons for uncertainty are less well understood, and issues that may become crucial at a later point (e.g., to resolve conflict) are simply averaged out of the analysis.

A key feature of SED is the *compactness* of its representation, and the resulting ease of assessment. SED makes adding background variables virtually painless, even in the absence of elaborate structures. It does so by exploiting the idea that an exception condition has only *one* impact on a given argument--reducing its precision--and that such impact can often be regarded as independent of the impact of other exception conditions. Thus, for each background premise, the only requirement is to specify which answers to the focal issue could no longer be discriminated from one another if the premise were false. This is done simply and qualitatively by placing *O*'s next to the appropriate subset of answers. SED takes these assessments, together with the core position, and automatically calculates the position supported by each combination of truth and falsity of the premises. To do so, it simply takes the union of the core position with the subsets of answers associated with the false premises.

For example, suppose the analyst has an argument to the effect that the function of Krasnoyarsk radar is space tracking based on satellite photographs of the radar equipment. Among the premises of that argument might be assumptions about the state of Soviet technology and the technical choices that Soviet engineers would make to solve various problems. In particular, suppose one background premise is to the effect that this type of radar would not be used by the Soviets for purposes of early warning of missile attack. If this premise were wrong, the argument would no longer be able to discriminate space tracking from early warning. Another premise might be that the observed equipment is real radar and not a mock-up placed there for purposes of deception. If this premise were false, the argument could no longer discriminate space tracking from other non-ABM. To the extent that both premises might be false, the argument fails to discriminate among all three possibilities: early warning, space tracking, and other non-ABM.

When there are n premises in an argument, SED requires only n + 1 assessments: the core position plus an exception condition for each premise. If all combinations of answers to premise topic/questions had specific significance, the analyst could use SED to create 2^n arguments. But that is the worst case in SED, whereas it is the only case in traditional conditioning models, such as influence diagrams (Shachter, 1986), Bayesian causal nets (Pearl, 1986), and Bayesian hierarchial inference (Schum, 1980). The key difference is in the basic units of analysis. A topic/question is a variable that can take various subsets of answers as values. The basic atom of analysis in SED is the relationship between specific *values* of variables: i.e., a concrete scenario or sequence of events. By contrast, the atom of analysis in traditional conditioning models is the relationship among *variables*.

A rather simple generalization of the present approach preserves the linear relationship of assessments to premises when the impact of a premise is more complex and/or depends on the impacts of other premises. We can: (1) allow an exception condition to operate on the results of applying previous exception conditions in a temporal sequence; and (2) specify the impact of an exception condition more generally: instead of a subset of answers within which discrimination can no longer take place, we can use a rule that substitutes one answer or subset of answers for another.

These extensions provide a very economical tool for representing certain quite general evidential arguments. For example, a standard sequence of events involved in learning about an event or situation from a human source is the following (cf., Schum, 1989):

Event E_1 → occurs	Perception → by observer of event E_1	Belief by → observer that E_1 occurred	Overt report by observer that E_1 occurred

As noted by Schum (1989), each state (perception, belief, testimony) is subject to exception conditions that include both confusion and bias. The core position, based on the report of E_1, is that E_1 occurred. But the observer may have misspoken or be lying; he may honestly believe that he saw something different from what he actually saw, because of what he wishes had happened or because he doesn't remember accurately; he may have misperceived the event due to perceptual biases, poor observational conditions, or limited perceptual capacities.

The interaction of exception conditions in examples of this sort can be effectively represented simply by ordering them in a temporal sequence. The position supported by any combination of exception conditions can be found by working backward along the causal sequence from the evidence (e.g., the report) to the ground truth situation. The process starts with what has been reported (= the core position) and asks at each step how it could have been generated from the previous step in the causal sequence under the given set of exception conditions: e.g., What beliefs could have led to the report? What perceptions could have led to *those* beliefs? and, What true situations could have led to *those* perceptions? More generally, the process starts with Result = the core position and transforms Result at each step according to the appropriate exception condition rule at that step. When multiple exception conditions are temporally unordered (at the same step), Result becomes the union of their impacts and the previous Result. Result at the end of the sequence is the supported position for that combination of conditions.

It can be speculated that any valid example of knowledge involves a causal connection of some sort between one's evidence and the conclusions one wishes to draw (cf., Shope, 1983; Nozick, 1981). SED (in this extended version) exploits such causal connections in a very pragmatic way--to provide an economical representation of rather complex evidential arguments. n premises can still be accommodated by n + 1 assessments, if we add a specification of their causal order.

<u>Adding Numbers</u>

Assessments (other than 0 and 1.0) are not required in SED to build an argument (or indeed an entire structure of interconnected arguments). It may, however, be useful to describe gradations of support that issues obtain via the arguments that bear on them. The philosophy of SED is to keep direct numerical assessments simple and to build relatively more complex numerical models on their basis.

Simple numerical judgment is required only for those issues that are at the "edge" of the inference network, i.e., issues which serve as premises in arguments, but not as focal issues for other premises. The analyst need only provide a number between 0 and 1.0 to indicate where he believes the truth lies: e.g.,

POSITION —CORE— Argument 1 of 2

	Krasnoyarsk radar	Are assets planned but not built
Yes		*
No	*	*
Core Support	.3 (1)	.7 (2)

(1) Assessed by analyst
(2) Supplied automatically by SED

The analyst in this example has assessed a 30% chance that his current knowledge proves there is no planned construction of assets in the vicinity of Krasnoyarsk. If the analyst provides no further inputs, SED will automatically allocate the remaining 70% support to the set of all answers, (Yes, No); i.e., there is a 70% chance that the analyst's knowledge is inconclusive on this issue.

This assessment is a very simple "belief function" (Shafer, 1976). A belief function is a measure of evidential support that assigns belief to *subsets* of answers rather than (as in Bayesian probability theory) to the answers themselves. As in probability theory, however, the sum of support for all the subsets must equal 1. As we have already seen, belief functions are useful for representing ignorance: assigning support to subsets with more than one answer means that the evidence fails to discriminate among those answers. Support assigned to the subset containing all possible answers (e.g., in this case, (Yes, No)) signifies the chance that the evidence is completely inconclusive. By contrast, a standard probability approach requires that all the probability somehow be allocated among the specific answers.

Assessments with regard to premises enable SED to generate the degree of support implied by arguments for the issue of main concern. For example, if there were no other premises in the argument based on failure to observe nearby assets, the analyst would get the following revised position:

POSITION —REVISED— Argument 1 of 2

	Krasnoyarsk radar	What is its function
Local defense		*
Early warning	*	*
Space tracking	*	*
Other non-ABM	*	*
Support	.3 (1)	.7 (1)

(1) Supplied automatically by SED

The core position of the argument (that the function of the radar is anything but local defense) is supported to the degree that the premise (no assets planned) is supported. To the extent that the premise is false or unknown, the argument can no longer discriminate local defense from the other possibilities.

In more complicated cases, where multiple premises are assigned varying degrees of support, SED computes (1) the position on the focal issue supported by each combination of truth and falsity of the premises, and (2) the aggregate degree of support for that position. The result may be a more complex belief function.

There is an affinity between SED's logical structures and Shafer-Dempster belief functions, since a belief function quantifies the chance that given evidence proves or fails to prove a hypothesis. Belief functions are, therefore, based on underlying (typically implicit) sets of judgments regarding the reliability of the link between evidence and hypothesis (Shafer, 1981b); SED requires that these judgments be made explicit as premises. In the process, SED breaks the assessment process down into sample components and clarifies the meaning of a belief function representation.

4. HIGHER-ORDER REASONING ABOUT CONFLICT AND ASSUMPTIONS

When two pieces of evidence or lines of reasoning appear to have conflicting implications, standard normative models statistically aggregate the numerical measures of their strength (e.g., by Bayes' Rule, Dempster's Rule, fuzzy logic, etc.). For example, suppose an analyst has (1) the photographic evidence alluded to above (that no significant assets have been seen near Krasnoyarsk); and suppose he assigned a high degree of numerical strength based on this evidence to the position that Krasnoyarsk is not intended for local ABM defense. Now suppose (2) a covert human source, highly placed in the Soviet military hierarchy, reports that Krasnoyarsk *is* being built for purposes of local defense. Given his previous experience with this source, the analyst assigns the same high level of strength based on the new evidence to the conclusion that Krasnoyarsk *is* intended for local defense. In numerical systems, these two pieces of evidence will simply cancel one another out, leaving equal amounts of belief in both possibilities. An analyst, by contrast, is more likely to wonder why two highly regarded sources are telling different stories.

He will look for an *explanation* of the conflict and, if he can, try to reduce it. SED supports that process. SED uses conflict as a symptom that something is wrong with one or more assumptions that led to the conflict (e.g., one or more sensors, models, human sources, etc. are not as reliable as supposed), and implements a process of higher-order reasoning that attempts to reduce conflict by reasoning about the assumptions or by collecting further data. Conflict, in short, is an opportunity to *learn* (e.g., are there possible undiscovered assets near Krasnoyarsk? Is there evidence of camouflage? How trustworthy is the informant? How credible are *his* sources? etc.)--not to blindly aggregate. The result may be valuable information for future use, and often, a more definitive picture of the problem at hand.

Assumptions

Knowledge requires assumptions. An analyst will be justified in believing nothing at all unless he is prepared to act *as if* other things were true. Even in cases of reasonable certainty, e.g., when two reliable and independent sources confirm a conclusion, there is the *possibility* of error (satellite photographic evidence can be fooled; a human informant may be misled). When sources do not agree, the dependence on assumptions merely becomes more salient. No analyst has the time or resources to rule out ahead of time all possible exceptions to a conclusion (and exceptions to those exceptions, etc.). In short, although he may have knowledge or evidence regarding some of the premises of an argument, such knowledge will never be complete or completely certain.

SED permits such assumptions to be adopted and utilized. However, SED makes a distinction (though only a matter of degree) between assumptions and firm belief. An assumption in SED is a belief that is:

(1) constrained by (though it goes beyond) what is more firmly known, and

(2) subject to retraction when and if it conflicts with new evidence or with lines of reasoning supported by other assumptions.

Could a rational decision maker get along without assumptions in this sense? To do so, he would have to deny (1) that any of his numerical judgments of belief are more firmly based than others, and (2) that he would ever retract such judgments in case of unexpected conflict with other lines of reasoning. In an ideal universe, where judgments reflect the totality of relevant knowledge, such claims may be plausible. In the real world, they are not.

The two definitions of assumption (going beyond firm belief, and subject to retraction in case of conflict) correspond to two complementary ways analysts may choose to assess their assumptions by using SED:

(1) "Bottom-up," by starting with a firm assignment of belief based on knowledge. This form of assignment may be too imprecise to support an argument which the analyst wishes to make. Thus, the analyst may use assumptions to reallocate belief that was committed to a set of possibilities to a proper *subset* of those possibilities.

(2) "Top-down," by starting with overall belief and specifying how much of it is firm and how much he would be willing to retract in case of conflict with other arguments. The analyst specifies how much of the belief in a set of possibilities he would transfer to a less precise superset of possibilities in case of conflict.

As an example of (1), the analyst may feel that the argument based on the failure to observe assets near Krasnoyarsk should carry more weight. In traditional systems, there is no way to reconcile the two judgments: (i) uncertainty about whether future assets are planned near Krasnoyarsk and (ii) reasonable confidence in the argument that the absence of present assets rules out local defense. The analyst would be compelled either to exaggerate his knowledge about the former or to relinquish his confidence (and his ability to act) on the latter.

SED solves this problem by making a distinction between what is firmly known about a proposition and the impact it has on a current argument. Its impact can be increased provisionally over what is strictly warranted by firm belief. The 70% support that remained uncommitted with respect to the premise defines an area within which the analyst is free to make assumptions. He may allocate all or part of it, by assumption, either to Yes or to No, by specifying a number between 0 and 1.0 for "% Assumed": e.g.,

POSITION —CORE— Argument 1 of 2—

	Krasnoyarsk radar	Are planned but not built?
Yes		*
No	*	*
Core Support	.3[1]	.7[2]
% Assumed	1.0[1]	
Final Support	1.0[2]	

(1) Assessed by analyst
(2) Supplied automatically by SED

In this example the analyst assumed no assets were planned. Final support of 1.0 is equal to the core support of .3 plus 100% of .7. SED will now generate a more decisive position for the argument:

POSITION —REVISED— Argument 1 of 2—

	Krasnoyarsk radar	What is its function?
Local defense		
Early warning	*	
Space tracking	*	
Other non-ABM	*	
Support	1.0	

The demarcation between knowledge and assumption is not absolute and fixed. Firmness of knowledge is a matter of degree: assumptions need not be entirely without evidential warrant; conversely, any belief might be retracted under some circumstances and thus have to be regarded as an assumption. The location of the boundary between "firm belief" and "assumption" is thus a matter of judgment for the problem at hand. Nevertheless, the distinction is a real one: there are beliefs the analyst is likely to hold onto come what may, and other beliefs that he is more likely to relinquish in the face of unanticipated conflict. The ability to draw such a boundary, even if it is itself a provisional one, is a powerful tool for capturing crucial aspects of reasoning about evidence.

Resolving Conflict

Two very different approaches to conflicting evidence have been adopted by students of inference. In logic-based systems, if it is possible to derive a contradiction from a set of statements, then one or more of the statements must be false. Suppose, for example, we start from the following beliefs:

Argument #1. If Source A reports anything, it is true.
Source A reports R.
R implies S.

Argument #2. If Source B reports anything, it is true.
Source B reports Q.
Q implies –S.

From these two arguments, we could infer an impossibility: the truth of both S and –S. To remove the inconsistency, at least one of the beliefs responsible for it must be revised. We *know* we are wrong about at least one of the following: the credibility of Source A or B, what they reported, or the implications of what they reported for S, –S.

A quite different approach has been adopted in systems that quantify and combine *degrees* of belief, like probability theory, fuzzy logic, or Shafer-Dempster theory. Suppose we believed:

Argument #1. Support (If Source A reports anything, it is true) = .99
Support (Source A reports R) = .99
Support (R implies S) = .99

Argument #2. Support (If Source B reports anything, it is true) = .99
Support (Source B reports Q) = .99
Support (Q implies –S) = .99

Although it may follow that we have very strong evidence for S and very strong evidence for –S, there is no logical contradiction. Even strong evidence may be imperfectly correlated with hypotheses. Legitimate evidential arguments may, therefore, point in different directions as long as each argument falls short of conclusive proof. Thus, it is *conceivable* that all our original beliefs were correct: both Source A and Source B are highly credible; A reported R; B reported Q; the former is strong evidence for S; and the latter is strong evidence for –S. The more pertinent question is whether it is still *plausible*, in light of this conflict, to regard all these beliefs as true.

The first approach to conflicting evidence is *epistemic*: conflict is regarded as a symptom of faulty beliefs and is used as an opportunity to correct them--by explicitly identifying potentially erroneous steps in the conflicting arguments. The second approach may be loosely referred to as *stochastic*: conflict among imperfect arguments is *expected* to occur by chance some portion of the time, and it is dealt with not by changing the arguments, but by statistically aggregating them when they both apply.

Each approach has virtues: On the one hand, the "stochastic" view, unlike the epistemic, permits gradations of belief; moreover, belief revision in epistemic systems is often arbitrary since there is no principled way to select one culprit from among the many beliefs responsible for a contradiction (cf. McDermott and Doyle, 1980). On the other hand, the stochastic approach is likely to "resolve" conflict in ways that are unconvincing and that fail to extract permanent lessons that might improve future inferences. Resolutions of conflict by stochastic methods are typically either too bland or too definitive. In the example above, since arguments #1 and #2 are equally strong in support of S and –S respectively, the conclusion is equal support for S and –S. If both arguments had been 100% certain, there would have been no determinate answer at all. The stochastic approach is even more likely to produce overly definitive results, as in the following hypothetical cases:

- Argument #1 strongly supports hypothesis S_1 but allows a very small chance that S_2 is correct; argument #2 strongly supports hypothesis S_3 but allows a very small chance that S_2 is correct. Statistical aggregation (Bayes' Rule, Dempster's Rule, etc.) results in 100% belief in S_2, which both sources regarded as highly unlikely (cf., Zadeh, 1984).

- Argument #1 strongly supports S and argument #2 strongly supports –S, but the degrees of support are not quite symmetrical, e.g., 99 to 2 in favor of S for argument #1, 99 to 4 in favor of –S for argument #2. The result: a 2 to 1 preponderance of belief in favor of S.

- According to argument #1, –S is impossible; according to argument #2, –S is favored 10^{10} to 1. The result: 100% belief in S.

For most people, these conclusions will seem a bit premature. Not surprisingly, therefore, the initial response to conflicting arguments is epistemic, rather than stochastic. Even when conflicting arguments have been expressed numerically, people look for *reasons* for the conflict: Did I overestimate the accuracy or honesty of one or both sources? (e.g., Should I reduce my belief in Source A's credibility from .99 to something lower)? Was I wrong in my understanding of what they said? Do my conclusions really follow from my understanding of what they said? The result, hopefully, is both a more convincing resolution of the conflict and an enhanced store of permanent knowledge.

In SED, numerical measures and an epistemic response to conflict are complementary rather than mutually exclusive. Conflict resolution is carried out by higher-order processes that *reason about* quantitative uncertainty models; conversely, numerical measures from those models provide guidance for decisions about adopting and revising assumptions. The result is a generalization of the epistemic approach, in which belief is graded, conflict is a matter of degree rather than all-or-none, and assumption revision is intelligently directed at those beliefs that are most likely to be in error.

Returning to the Krasnoyarsk radar example, suppose we have the two arguments: (1) against local defense, based on observation of no assets and with the assumption that no future assets are planned; and (2) for local defense, based on the report of a covert Soviet source. The analyst's firm belief that the source is reliable is .4 but he chooses to allocate 50% of the remaining .6 support to the *assumption* that the source is reliable. We can represent the combination of these two arguments, with their respective assumptions, in the following way:

ARGUMENT 1

	Source is reliable (.4)	Assume: source is reliable (.3)	Don't know about source reliability (.3)
No planned assets (.3)	Conflict .3 x .4 = .12	Conflict .3 x .3 = .09	Not local defense .3 x .3 = .09
Assume: no planned assets (.7)	Conflict .7 x .4 = .28	Conflict .7 x .3 = .21	Not local defense .7 x .3 = .21

ARGUMENT 2

Each cell represents the supported position and the degree of support that is implied by the relevant combination of circumstances regarding the premises of the two arguments. In the standard belief function approach, the result of combining these two arguments would be 100% belief in "not local defense." Conflict, since it represents an impossible state of affairs, is disregarded and the remaining cells are normalized.

In SED, however, the total amount of conflict between two arguments is interpreted as evidence that beliefs contained in those arguments are mistaken. This is a straightforward generalization of the logical strategy of showing a belief to be false by deriving a contradiction from it. Let T be the conjunction of beliefs in arguments 1 and 2. If T implies p and $\neg p$, then $\neg T$. In SED, T implies a quantitative weight on p and $\neg p$, corresponding to the *chance* that the beliefs in T imply a contradiction. In the example, the total amount of conflict is .12 + .28 + .09 + .21 = .7. That weight can be taken as the chance that conflict proves at least one of the members of T to be false. If the conflict measure were smaller, the two arguments could perhaps be left as they are and conflict resolved stochastically (in effect, by dropping the impossible states of affairs from the calculations). When the measure is large, however, it may be wiser to take a closer look at the contents of T.

In order to resolve conflict, the analyst needs to focus his scrutiny on those assumptions that seem to bear the most responsibility. Thus, SED's CONFLICT screen provides a rough decomposition of conflict into components that are attributable to separate assumptions:

CONFLICT	TOPIC	QUESTION
.70	Krasnoyarsk radar	What is it's function

CONTRIBUTION TO CONFLICT	TOPIC	QUESTION
.49	Krasnoyarsk radar	Are assets planned but not built?
.30	Soviet source	Is he reliable?

The degree of conflict attributable to the assumption that no future assets are planned is .28 + .21 = .49. The degree of conflict attributable to the assumption about the reliability of the observer is .21 + .09 = .30. These numbers might lead the analyst to drop the assumption regarding future assets (thereby reducing total conflict to .7 - .49 = .21). Alternatively, he might seek additional data to confirm or deny either or both of the assumptions.

Notice that the two assumptions in the above example were not independent in their impact. After dropping the first assumption, the contribution of the second assumption to conflict would be reduced from .30 to .09, since part of the total conflict (.21) was jointly determined. Each measure of an assumption's contribution to conflict is thus a sort of upper bound, conditional on retaining both the other assumptions and the firm beliefs that it clashes with. Assumptions have by definition a higher prior likelihood of being in error than firm beliefs and are thus more likely to be retracted. Hence the measure of an assumption's contribution to conflict is less ambiguous when the assumption clashes only with firm beliefs and not with other assumptions; in that case, it is more readily interpretable as the chance that conflict proves the assumption wrong. Indeed, if firm beliefs could *never* be withdrawn, SED might focus exclusively on the conflict attributable to assumptions (in this example, .28 + .21 +.09 = .58). Only this portion of the conflict would be treated epistemically; conflict due to firm beliefs alone (.12) would always be handled stochastically.

For SED, however, the boundary between assumptions and firm beliefs is itself subject to review. A large measure of conflict, if there were no assumptions or no assumptions clearly identifiable as culprits, might very appropriately lead an analyst to re-examine the relevant "firm beliefs." He might then convert a firm belief into an assumption by using the top-down method (i.e., using the ARGUMENTS screen to specify what portion of the total belief was firm). Alternatively, he might add exception conditions to the argument expressing a firm belief (as with the "crystal ball" technique). He might then return to CONFLICT to observe the potential effect on conflict of dropping the newly defined assumptions. Conflict resolution is thus an occasion for the continued elicitation and refinement of the analyst's beliefs.

Conflict can help an analyst search deeply through a network of beliefs for a potential culprit, and revisions may be made at any level. In particular, conflict resolution may be a valuable tool for detecting deception. Although direct evidence of deception is possible (e.g., overheard communications, observation of no activity at a dummy facility), more often than not evidence for deception is available only indirectly in the form of evidential conflict. SED is a uniquely appropriate tool for assessing the possibility and the scope of deception activities.

SED embeds numerical uncertainty representations within a process of higher-order reasoning. Is such a higher-order process really necessary? Could the functions of conflict resolution be accomplished instead *within* a standard numerical calculus? The answer is: in principle, yes: in practice, no. To simulate the effect of conflict resolution with a numerical calculus, it would be necessary to explicitly represent all the situations in which conflict could arise and decide on a resolution ahead of time. We would need a vast number of exception conditions specifying which other sources and arguments would override a given argument, e.g.,

- Source A is reliable when he reports R *unless* source B reports Q and source C reports T and source D reports U... or source E reports V and source F reports W... or ...

In a numerical framework (e.g., Bayesian or Shaferian), a huge set of conditional assessments would be required, linking the elements of every line of reasoning to the elements of all other possible lines of reasoning. The price of such a strategy comes not only in the sheer quantity of inputs and computational intractability, but also in a loss of naturalness and modularity.

In order to remain tractable, numerical inference models typically treat hypotheses about diverse information sources or lines of reasoning as if they were independent. The result is a stochastic approach to conflict that fails to extract the real significance of conflict when it occurs. SED achieves the best of both worlds: It enables the analyst to bring to bear the conclusions of one argument on the evaluation of the other without sacrificing the modularity of the different lines of reasoning.

5. CONCLUSION

While the difficulties of collecting intelligence data are well understood, the difficulties in analyzing and interpreting those data are often overlooked. There is a growing awareness, however, that the success of the overall intelligence enterprise depends crucially on those processes which occur *after* the data have been collected. The present report has described a tool which is designed to make the intelligence analyst's task easier and more successful.

SED embodies promising technical solutions to all three of the problems we identified in the first section. It clarifies the meaning of numerical assessments by emphasizing qualitative models of how evidence is linked to conclusions; it requires only simple numbers reflecting different ways that such evidential links could be broken. More complex numerical models are then automatically generated. SED wards off a combinatorial explosion of assessments in two ways: (1) by introducing a simple method for deriving the impact of multiple factors on a conclusion from assessments of their separate impacts, and (2) by providing for non-independence of different lines of reasoning through a higher-order process of conflict resolution. As a result, SED encourages, rather than discourages, users to introduce new factors into an analysis; i.e., to make the reasons for uncertainty explicit. Even when there is no direct knowledge regarding such factors, SED permits users to introduce them and, if they wish, to adopt assumptions. Finally, SED does not assume that a problem has been solved simply because a numerical model has been created. It focuses on the processes that intelligently create and revise such models. When two or more arguments point in different directions, SED does not sweep conflict under the rug by statistically aggregating them. It supports the analyst in a process of re-examining and modifying beliefs and assumptions that contributed to the conflict.

In most computerized aids that quantify uncertainty, inference is equated with an essentially linear process, in which a model or "knowledge base" is built, numerical inputs assessed, and outputs generated:

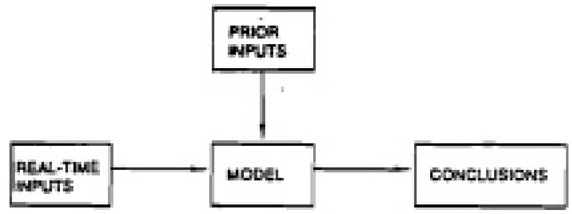

Such an approach may ensure consistency of inputs and outputs with respect to a set of axioms, e.g., probability theory; the problem is, more than one set of inputs and outputs, with vastly different implications for a decision, will be equally acceptable from a strictly *formal* point of view. Automation of uncertainty handling thus omits the *thinking* processes by means of which an analyst selects *one* consistent set of beliefs out of all those that are possible. Actual probabilistic reasoning is typically highly iterative: the results of one line of reasoning are compared with the results of other lines of reasoning (or with direct judgment); if there is a discrepancy, the inputs, parameters, and even the structure of the model or knowledge base may be revised:

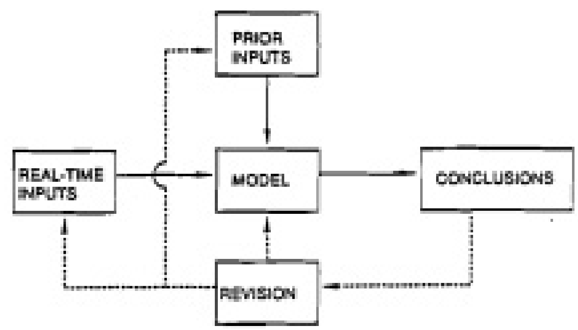

SED provides direct support for the intelligent construction and modification of inference models in the light of experience with their application. In effect, SED redefines "reasoning": it is no longer the blind *application* of an uncertainty model, but its creation and maintenance.

REFERENCES

Burrows, W.E. *Deep black.* New York, NY: Random House, 1986.

Cohen, M.S. An expert system framework for non-monotonic reasoning about probabilistic assumptions. In J.F. Lemmer and L.N. Kanal (Eds.) *Uncertainty in Artificial Intelligence* Amsterdam: North Holland Publishing Co., 1986.

Cohen, M.S. *Conflict resolution as a knowledge elicitation device.* Draft Technical Report. Reston, VA: Decision Science Consortium, Inc., 1989.

Cohen, M.S., Laskey, K.B., and Ulvila, J.W. *The management of uncertainty in intelligence data: A self-reconciling evidential database* (Technical Report 87-8). Falls Church, VA: Decision Science Consortium, Inc., June 1987.

Cohen, M.S., Laskey, K.B., Vane III, R.R., McIntyre, J.R., and Sak, S.G., *Divergent and convergent feasoning in intelligence analysis: A self-reconciling evidential database* (Technical Report 88-14. Reston, VA: Decision Science Consortium, Inc., 1989.

Phillips, L., IPL/AMRD. Applications of inference theory and analysis of artificial intelligence to problems in intelligence analysis. *Intelligence Production Laboratory Conference Proceedings*. Rosslyn, VA: Analytic Methodology Research Division, Office of Research and Development, Central Intelligence Agency, November, 1982.

Laqueur, W. *A world of secrets*. New York, NY: Basic Books, Inc., 1985.

McDermott, D., and Doyle, J. Non-monotonic logic I. *Artificial Intelligence*, 1980, *13*, 41-72.

Nozick, R. *Philosophical explanations*. Cambridge, MA: Harvard University Press, 1981.

Pearl, J. Fusion, propagation and structuring in belief networks. *Artificial Intelligence*, 1986, *29*(3), 241-288.

Schum, D.A. Current developments in research on cascaded inference processes. In T.S. Wallsten (Ed.), *Cognitive process in choice and decision behavior*. Hillsdale, NJ: Lawrence Erlbaum Associates, 1980.

Schum, D.A. Knowledge, probability, and credibility. *Journal of Behavioral Decision Making*, *2*, 39-62, 1989.

Shachter, R.D. Evaluating inference diagrams. *Operations Research* 34 (1986) 871-882.

Shafer, G. *A mathematical theory of evidence*. Princeton, NJ: Princeton University Press, 1976.

Shafer, G. Jeffrey's rule of conditioning. *Philosophy of Science*, *48*, 337-362, 1981b.

Shope, R. *The analysis of knowing: A decade of research*. N.J., Princeton University Press, 1983.

Tolcott, M.A., Marvin, F.F., and Lehner, P.E. *Effects of Early Decisions on Later Judgments in an Evolving Situation* (Technical Report 87-10). Reston, VA: Decision Science Consortium, Inc., July 1987.

Zadeh, L.A. Review of Shafer's, A mathematical theory of evidence. *AI Magazine*, 1984, *5*(3), 81-83.