CONSISTENCY AND PLAUSIBLE REASONING

J.R. Quinlan *
Rand Corporation
Santa Monica CA 90406

ABSTRACT

The usual approach to plausible reasoning is to associate a validity measure with each fact or rule, and to compute from these a validity measure for any deduction that is made. This approach is shown to be inappropriate for some classes of problems, particularly those in which the evidence is not internally consistent. Two current plausible reasoning architectures are summarized and each applied to the same small task. An analysis of the performance of these systems reveals deficiencies in each case. The paper then outlines a new approach based on the discovery of consistent subsets of the given evidence. This system can be used either in isolation or in conjunction with a validity-propagating architecture. Comparative results from implementations of all three systems are presented.

1. INTRODUCTION

Research in expert systems is concerned with how to represent and reproduce the problem-solving skills that experts exhibit in their respective domains. *One* of the most basic of these skills is the ability to put two and two together—to draw reasoned conclusions that supplement direct observations. This poses a difficulty because our models of reasoning are derived from the deduction mechanisms of logic, but investigators have noted that expert reasoning cannot be understood in terms of such precise schema. Logic deals with an idealized world in which facts are known with certainty and rules of inference allow other facts to be deduced with equal certainty. Experts, on the other hand, are usually required to form judgments based on evidence that may be subject to errors of measurement or interpretation. Again, experts can function in environments containing inconsistent or contradictory "facts", but such environments are useless in the logical sense.

Systems for expert-like reasoning typically work with unquantified propositions and relations among them. Propositions may be facts ("the car won't start"), hypotheses ("the trouble is in the ignition system"), findings ("the distributor cap is defective"), or any concept in the domain that is relevant to the expert's problem-solving behaviour. Relations connecting subsets of the propositions are usually expressed as logical definitions ("A is the conjunction of B, C, and D") or inferential links ("If A is true, then so is B"). The feature that distinguishes uncertain inference from the familiar propositional calculus is the qualified nature of knowledge about both the relations and propositions. Instead of being either true or false, propositions have associated with them some continuous measure of their validity. Inferential relations also have a validity measure that weakens the connection between antecedent and consequent; the relation "If A, then B" may support a less-than-categorical affirmation of B even when A is known with certainty.

A useful way of viewing this formalism is as an inference net (Duda, Gaschnig, and Hart, 1979) in which propositions are nodes and the relations among propositions become the links of the network. Whenever the validity measure of a node is changed, such as by the arrival of new evidence, this information propagates along the links to related nodes and may cause changes to their validity measures in turn. The secondary changes propagate in the same way so that the evidence responsible for the initial change may be reflected in altered validity measures of many propositions.

The general inference net framework does not address the important questions of how validity is to be represented and how the propagation above is to be carried out. Some approaches measure the validity of a proposition as its posterior probability given all the evidence to hand, computed via Bayes' Theorem and ancillary assumptions. Others use probability intervals rather than point probabilities as a measure of validity, relying on more general schemes of updating such as the Dempster-Shafer theory of evidence (Garvey, Lowrance, and Fischler, 1981). It is not uncommon for the form of the inference net to be restricted; for example, (Pearl, 1982) requires that it be a tree. Many systems treat the links representing relations as directional, so that the relation "If A, then B" allows updating of B's validity when A is known to be true but does not allow A's validity to be altered if B is found to be false. A review and critique of the more common approaches can be found in (Quinlan, 1983).

* Present address: School of Computing Sciences, NSW Institute of Technology, Sydney 2007 Australia

Expert systems embodying mechanisms for uncertain inference have achieved notable successes, as documented in recent reports (Buchanan, 1982; Campbell, Hoilister, Duda, and Hart, 1982) on two pioneering efforts, MYCIN and Prospector. Nevertheless, there appear to be applications requiring an uncertain inference capability that are not handled well by any current system. The characteristics of these applications are discussed later, but the gist of the difficulty and the proposed solution can be obtained from the following example.

Consider the task of a fictional detective investigating a case in which (as usual) there are many apparent contradictions in the evidence that he unearths. How is he to proceed? Current approaches to plausible inference would have him weigh evidence for and against each hypothesis, considering the hypothesis confirmed to the extent that the balance of evidence supports it. But any mystery buff knows that this approach differs from the one Poirot would adopt, and might even lead to the anomalous situation in which the balance of evidence individually supports propositions A and B, but where A and B cannot both have occurred. This paper suggests an alternative method of forming conclusions that our detective, would find more familiar. Instead of making deductions from contradictory information, we divide the evidence into two classes, items to be believed and items to be disregarded, so that all the evidence in the former category is consistent and "makes sense". Where there are many possible divisions we use some model to weigh the validity, not of individual propositions, but of the <u>division itself</u>. For example, a division that would require our detective to disregard significantly more data than another might be judged to be less valuable.

In the following sections we examine a seemingly simple uncertain inference problem. Taking Prospector as an example of a directed Baycsian architecture, we show that the problem must be redrafted to meet Prospector's requirements and that there are difficulties interpreting the results. We then describe INFERNO, a non-directed non-Bayesian architecture, and show that it is also less than satisfactory for this task. This leads to a discussion of Ponderosa, a new system that performs uncertain inference by evidence division rather than by propagation of validity.

2. DESCRIPTION OF THE TRIAL APPLICATION

The setting for this application of uncertain inference is a model of the interactions among five econometric indicators. We are given several assertions concerning both general relationships among the indicators and predictions about what will happen in the near future. The goal is to draw meaningful inferences from these assertions so as to arrive at a composite picture of what will happen to all the indicators.

Table TI contains the ten assertions that define the model. Numbers in brackets following assertions are validity measures in the range 0 to 1; two such numbers following an assertion correspond to the "if" and "only if" cases respectively. Since we have not defined what we mean by "validity", the precise interpretation of these numbers is open. A proposition or relation with validity 1 is correct without qualification and one with validity 0 is false, but any of the different meanings of a middle-ground validity that are used in current systems wil] be acceptable.

Table TI: Assertions Defining the Model

AI Stocks will fall (.55)

A2 Either taxes will not be raised or both stocks will fall and interest rates will fall (.85)

A3 Either taxes will be raised or interest rates will not fall (.9)

A4 Interest rates won't fall (.75)

A5 Either taxes will-be raised or there will be a high deficit (.85)

AG Bonds will rise or interest rates will fall if, and only if, stocks fall or taxes are not raised (.6, .85)

A7 Stocks will fall if, and only if, bonds rise and taxes are raised (.7, .8)

A8 If interest rates fall, either stocks will not fall or bonds won't rise (.95)

A9 Interest rates will not fall if there is a high deficit (.95)

A10 If there is a high deficit, stock:; will fall (.8)

The application maps directly into the inference net formalism. There arc five basic propositions corresponding to the indicators of primary concern,

stocks will fall	(abbreviated stocks-)
interest rates will fall	(interest-)
taxes will be raised	(taxes+)
bonds will rise	(bonds+)
there will be a high deficit	(deficit-*)

We have also several composite propositions stated as logical combinations of these propositions, such as "bonds will rise or interest rates will fall", that must be defined by logical relations. Assertions AI through A5 each provides evidence in the form of a validity for one of the basic or composite propositions, while each of A6 through A10 becomes one or two inferential relations.

Despite the simplicity of this model, it may not be immediately apparent that the information in the assertions is inconsistent. AI and A7 jointly support the inference that taxes will be raised, while assertions A2 and A4 together suggest that taxes will not be raised. In the logical sense, therefore, this collection of assertions is of no value because anything at all can be inferred from it via the tautology A implies (~A implies B). However, it seems that most plausible reasoning tasks involve inconsistent information so that the example is not an unfair one.

3. PROSPECTOR

Prospector (Duda, Hart, and Nilsson, 1976) is a general-purpose architecture for uncertain inference that has been used with several geological models and whose basic approach has been taken up by other systems such as AL/X (Reiter, 1981). It is therefore representative of a well-developed school of thought about uncertain inference.

3.1 Overview of Prospector

Prospector takes the validity of a proposition to be its posterior probability given the evidence at hand. Let H be some proposition about which inferences are to be drawn and E another proposition. Bayes' theorem gives the posterior probability (or likelihood) of H given E as

$$P(H|E) = P(E|H) \times P(H) / P(E)$$

where P(E) and P(H) are prior probabilities, and similarly

$$P(\sim H|E) = P(E|\sim H) \times P(\sim H) / P(E)$$

Assuming that the latter is non-zero, we can divide the first equation by the second to obtain

$$O(H|E) = O(H) \times [P(E|H) / P(E|\sim H)]$$

which may be stated as, the posterior odds of H is its prior odds multiplied by a factor (called X) that characterizes the sufficiency of E as a predictor of H. A similar analysis can be performed replacing E by E in the above, and the corresponding factor λ characterizes the necessity of E if H is to hold.

This formalism is insufficient by itself to determine what should happen to the odds of H when several propositions El, E2, ... are relevant to it, or when the E's are known with less than certainty. The approach taken in Prospector is to make two additional assumptions (conditional independence and interpolation) that allow the posterior odds of proposition H to be computed as the product of its prior odds and effective multiplying factors obtained for each Ei.

Inferential links from one proposition to another can thus be implmented by choosing appropriate values for the factors X and λ. The posterior odds of logical combinations of propositions is computed from those of the components, e.g., if A is the conjunction of Bl, B2, ... , the odds of A is the minimum of the odds of any Bi. Each relation can cause the odds of only one proposition to be altered directly; inferential relations "If E, then H" as before affect only H, and logical relations as above affect only A. Accordingly, the links representing relations are thought of as directed into the affected proposition. Prospector requires that there be no cycles in the inference net and allows observed probabilities to be given only for "evidence" propositions that have no links directed into them.

3.2 Applying Prospector to the Model

Several difficulties arise when we attempt to use the Prospector architecture for the model. The more serious of these are consequences of Prospector's tacit assumption that propositions can be arranged in a hierarchy with inference chains flowing smoothly from raw evidence through to conclusions.

Consider, for example, the proposition "stocks will fall". This appears as evidence in Al but as a conclusion in several other assertions. Again, assertions A2, A3, and A5 provide validities for logical combinations of propositions, but Prospector contains no mechanism that would allow evidence to bear directly on such composite propositions. Similar problems arise from A6, A7, and A8, where logical combinations are on the receiving end of the inferential links.

The steps taken to reformulate the example are as follows: (1) The two propositions "stocks will fall" and "interest rates will fall" that appear both as evidence and as potential conclusions are represented each by two nodes in the net. The first is a conventional evidence node with a very strong inferential link to the second copy that is also the recipient of other inferential links. (2) Assertions such as A2 of the form "A or B" are represented functionally as the pair of inference relations "If A is false, then B" and "If is B is false, then A". (3) Complex assertions are broken down into more primitive relations that have a single proposition as the inference. For example, A6, of the form "A or B if, and only if, C or D" becomes the set of relations (and their symmetric counterparts)

```
If (A or B) and ~C, then D
If (A or B) and ~D, then C
If ~A and ~B, then ~C
If ~A and ~B, then ~D
```

(4) Finally, all prior probabilities are taken by default as 0.5 since the example does not specify other values, and the strengths of the multipliers λ and λ' are determined so that, if the relation "If A, then B" has validity V, the posterior probability of proposition B given A is also V.

Even with these changes, the reformulated example violates a Prospector prohibition on cycles in the net. These arise from strong interconnections among the five indicators, however, and there seems to be no way of eliminating them. Rather than abandon the enterprise, we will generalize the Prospector algorithm to allow computation of posterior probabilities by relaxation, terminating when changes are small so that cycles will not cause infinite loops.

A Prospector-like system embodying this modification was used to obtain the results shown in Table T2. These results are deficient in at least two respects. (1) They give no hint that the assertions from which the model was derived are inconsistent. The assumptions that Prospector makes will never produce an overconstrained system, so any collection of evidence and relations will lead to a solution. (2) The statement of a result as a probability is fine when there is only one result of interest, but can lead to problems in cases such as this when we need a simultaneous reading on several hypotheses. Suppose that the model builder wished to predict the most likely future state from the 2 possible in terms of the five indicators. Converting the probabilities to categorical form by thresholding as in Table T2 would lead to the conclusion that

stocks will fall;
interest rates will not fall;
taxes will not be raised;
bonds will rise; and
there will be a high deficit

These conclusions jointly violate the "only if" part of assertion A7! Thus mapping from probalistic to categorical results for several variables may produce conclusions that do not fit with the evidence.

Table T2: Results from a Prospector-Style System

Proposition	Posterior Probability	Categorical Interpretation
stocks-	.64	T
interest-	.08	F
taxes+	.27	F
bonds+	.59	T
deficit+	.66	T

In summary, in order to run our example on Prospector we had to make significant alterations to the formulation of the model and to modify Prospector as well; even so, the results we obtained were deficient. For all these reasons it would seem that Prospector is not well-suited to this application.

Konolige (1982) has developed an alternative scheme for Bayesian inference that finesses many of the difficulties above. This scheme, based on information theory, allows non-directional propagation among Local Event Groups, each characterized by a complete probability distribution. The model could be run without modification, but the deficiencies regarding inconsistencies and categorical interpretation would seem to remain.

4. INFERNO

INFERNO (Quinlan, 1983) is another inference network system that was designed around four ideas:

1. General systems for uncertain inference are better off without assumptions such as conditional independence whose universal validity is suspect (Pednault, Zucker, and Muresan, 1981).
2. On the other hand, it should be possible to assert that particular groups of propositions exhibit relationships such as independence.
3. There should be no restrictions on the direction of information flow in the network. (This was the cause of much of Prospector's difficulty with the model.)
4. The consistency of the data should be checked and the system should be able to advise on alternative methods of rectifying inconsistencies

The effect of these requirements has been to lead away from Prospector-style formalisms.

4.1 Description of INFERNO

The first difference comes in the way that the validity of a proposition is represented. Instead of a single point probability, INFERNO uses probability bounds; a proposition A is characterized by a lower bound t(A) on the probability P(A) of A and a lower bound f(A) on P(~A). This approach has two advantages. The uncertainty of our knowledge about A is apparent, being just the difference between t(A) and 1 - f(A). Second, the values of t(A) and f(A) are derived from evidence tending to support and to deny A respectively, and these values are retained and propagated separately.

To achieve the non-directed propagation of inferences as in point (3) above, INFERNO follows WAND (Hayes-Roth, 1981) in viewing relations as constraints on the respective validities of collections of propositions. Changing a probability bound of any proposition in the collection may require some other bound to be altered to preserve the constraint. For example, one form of inferential relation, written as

A enables B with strength X

is intended to capture the (uncertain) relation "If A, then B". This relation has two associated constraints:

$$t(B) \geq t(A) \cdot X$$
$$f(A) \geq 1 - (1 - f(B)) / X$$

and thus can cause t(B) to be increased when t(A) is increased, or f(A) to be increased when f(B) is increased. Logical connections among propositions are handled in the same manner. The relation defining A as the conjunction of Bl, B2, ... , Bn gives four constraints: for all Bi,

$$t(A) \geq 1 - \Sigma_i(1 - t(Bi))$$
$$f(A) \geq f(Bi)$$
$$t(Bi) \geq t(A)$$
$$f(Bi) \geq f(A) - \Sigma_{j \neq i}(1 - t(Bj))$$

These and all other INFERNO constraints can be derived from simple probability identities without other assumptions.

This representation supports a probabilistic concept of consistency. If t(A) + (f(A) > 1 for some proposition A, the information about A is inconsistent and one or both of the bounds must be incorrect. Since the propagation constraints are provably correct, the inconsistency can only arise from contradictions implicit in the information given to the system. INFERNO can suggest ways to alter the data so as to make it consistent. Such a suggestion, called a rectification, identifies one or more assertions whose given validity must be reduced and/or inferential relations that must be weakened. INFERNO can generate the best n of the possible rectifications, ranking them under the assumption that those involving the least alteration of the original data are more likely to be acceptable.

4.2 Applying INFERNO to the Model

When we wished to apply Prospector to the model we first had to reformulate it to conform to Prospector's architectural restrictions. INFERNO does not impose any .such restrictions and the mode can be run in its original form.

INFERNO immediately finds the set of assertions to be inconsistent. Analysis of the various interdependencies then leads it to propose four alternative rectifications, each of which will correct, all inconsistencies. Each rectification consists of a single change:

Reduce the validity of assertion A4 to .71
Reduce the validity of assertion A2, to .81
Reduce the validity of assertion Al to .5
Weaken the only-if strength of assertion A7 to .727

This analysis is intended to permit the user to review selected fragments of the data with an eye to making it consistent before trusting conclusions based on it.

Let us suppose the user, after reflection, decides that assertion Al is inapplicable in this case and should be completely disregarded rather than just having a lower validity. The consistent set of probability bounds that INFERNO obtains from A2 through A10 is shown in Table T3. In general it is more difficult to place a categorical interpretation on INFERNO's ranges than it was in the case of Prospector's point probabilities, but in this instance the mapping to (T,?,F) seems reasonable. Notice, though, that the categorical interpretation again violates a relatively strong relation (A10) predicting that stocks will fall if there is a high deficit!

To summarize: INFERNO avoids three of the four difficulties that Prospector experienced with the model. It allows assertions and inferences about logical combinations of propositions and is not put out by cycles in the net. It also makes apparent any inconsistencies in the data and provides helpful aids to reviewing it. However, an attempt to place categorical interpretations on the results can once more lead to conclusions that are not consistent with the data.

Table T3: Results from INFERNO

Proposition	Probability Range	Categorical Interpretation
stocks-	.36 - .5	F
interest-	.138 - .25	F
taxes+	.288 - .4	F
bonds+	.288 - 1	?
deficit+	.45 - .625	T

5. PONDEROSA

Ponderosa represents a departure from current plausible inference systems because, although it still deals with uncertain assertions and relations, it docs not attempt to propagate validity measures of any kind. Instead, it follows the approach of trying to identify internally consistent subsets of the information given to it. The merit of any such division is then established as a function of the validities of assertions that were not included.

5.1 Description of the Approach

Each assertion in the model can be viewed as a well-formed formula (wff) of the propositiona] calculus with a validity measure attached, or, in the case of the "if and only if" assertions, a pair of such formulas. Let C be a subset of the wffs, where we disregard for the moment each wff's validity measure. C is consistent if there is no wff that can be both proved and disproved from C* A subset is maximally consistent if it is consistent but the addition of any other wff from the original set will make it inconsistent.

Suppose now that the original set of wffs has been divided into a maximally consistent subset C and remainder R = {R1,R2, ..., Rnj and let V(Ri) be the validity measure of Ri. One way of assessing the situation would be to accept the wffs in C together with all their (consistent) inferences and to ignore the wffs in R as being either erroneous (e.g., resulting from faulty observation) or default assertions that do not apply in this case. How plausible is this division? If it is to be correct, each individual Ri must be incorrect or inapplicable. The probability that this division is correct is then the probability of the conjunction

$$P(\sim R1 \& \sim R2 \& \ldots \& \sim Rn)$$

If we again treat validity measures as probabilities and use the identity

$$1 - (P(\sim A) + P(\sim B)) \leq P(A\&B) \leq P(A),P(B)$$

we obtain the probability P(C,R) that the divisio is correct as

$$1 - \Sigma_i \ V(Ri) \leq P(C,R) \leq \min_i \ (1 - V(Ri))$$

Since we are identifying validity measures with probabilities, P(C,R) represents the validity

* This notion of consistency is stronger than the one used for INFERNO in which it is permissible to infer both A and ~A so long as the sum of the upper bounds of P(A) and P(~A) does not exceed 1.

of the division of the original set of wffs into C and R.

The number of potential splits of a set of wffs into a maximally consistent subset and a remainder grows exponentially with the size of the set. The validity measures attached to propositions, however, provide methods of reducing the computational load. First, we are clearly uninterested in any division whose validity is zero. If any wff in R has a validity of 1, the inequality above gives a zero upper bound on the validity of that division. Consequently, we need consider only divisions in which all categorical assertions are included in the consistent subset C. Second, we do not wish to swamp the user with all possible divisions, but rather to generate and display only the best N of them for some small, fixed N.

5.2 Algorithm for Finding Divisions

We now give a method for finding the best N maximally consistent subsets of the wffs. This is necessary to demonstrate that Ponderosa's approach is computationally feasible, but casual readers might prefer to move directly to section 5.3.

Each proposition A is broken into two *findings* "A is true" and "A is false". Associated with each finding is a collection of *justifications* for the finding, where a null justification indicates that there is no reason to believe the finding. Each justification for the finding is either that the finding is an explicit assertion given to the system, or that the finding is an inference from a relation and one or more other findings with non-null justifications. For instance, the finding "B is false" and the relation "A implies B" together justify the finding "A is false", and the logical relation "X is the disjunction of A and B" together with both these findings justifies "X is false".

Every datum is either a relation or a given finding and all findings depend ultimately on the data. Ponderosa keeps with each finding a *removal plan* in the form of a collection of sets of data, the idea being that *all* justifications for this finding would evaporate if, and only if, any one of these sets of data were removed. The algorithms below depend on the observation that a removal plan is isomorphic to a logical expression in disjunctive normal form. Let us map each datum D to the predicate "D is excluded" and the removal plan

{ {D11,D12,...}, {D21,D22,... }, ... }

to the logical expression

(D11 is excluded and D12 is excluded and ...)
or (D21 is excluded and D22 is excluded and ...)
or ...

Then the expression is true if, and only if, one of the sets of data making up the removal plan has been discarded, in which case the plan is satisfied.

The computation of removal plans keeps pace with the propagation of inferences. Initially the only findings with justifications are those that appear in the data, and the removal plan for such a finding is {{itself}}. Suppose now that a new justification for finding F has been inferred from a relation R and findings {si}. This justification could be removed if either R or any of the S's could be removed, as given by the plan (in disjunctive form)

X = { R} v removal plan(S1)
v removal plan(S2)
v ...

But previous justifications may have been found for F and removal of F would require removal of them as well. In this case, the new removal plan for F becomes the conjunction of the old removal plan and X.

When the data are inconsistent there will be one or more propositions { Ai} that can be both proved and disproved, i.e., one or more pairs of findings "Ai is true" and "Ai is false", both with non-null justifications and removal plans. Clearly, the data would become consistent if, and only if, one of each such pair of findings could be removed. When put into disjunctive normal form, the removal plan obtained as the conjunction over i of

removal plan ("Ai is true")
v removal plan ("Ai is false")

is then just the set of remainders corresponding to all possible maximally consistent sets.

The final problem is to find the best divisions without computing all of them. The validity of a division is known only as a range, but divisions can be ranked approximately by comparing the midpoints of their ranges. Ponderosa computes the overall removal plan above in a depth-first fashion so that, if a partial remainder is generated that is already more implausible than the best N complete remainders found so far, all possible remainders containing the partial one are skipped.

5.3 Applying Ponderosa to the Model

As was the case with INFERNO, Ponderosa contains no restrictions that would require the model to be reformulated. Once again the information in assertions A1 through A10 is found to be inconsistent and Ponderosa generates the six possible divisions of the corresponding wffs into a maximally consistent subset and a remainder. The six remainders are displayed in Table T4 together with the bounds on the validity of the divisions and the midpoints of these ranges. Notice that, whereas INFERNO would accept the weakening of just "stocks will fall" as sufficient to remedy the inconsistencies, Ponderosa uses a stronger definition of consistency and finds that removal of assertion A1 alone is not enough.

Table T4: Ponderosa Remainders

Remainder		Validity Low	Mid	High
A7b	stocks- only if bonds+ & taxes+	.2	.2	.2
A2	~taxes+ v stocks- & interest-	.15	.15	.15
A1 A4	stocks- ~interest-	0	.125	.25
A1 A10	stocks- if deficit+ then stocks-	0	.1	.2
A1 A5	stocks- taxes+ v deficit+	0	.075	.15
A4 A8	~interest- if interest- then ~stocks- v ~bonds+	0	.025	.05

Ponderosa does not automatically select the "best" or any other maximally consistent .subset as being correct. Its function stops with pointinq out to the user the possibilities that exist for making his information consistent, using the validity rankinq only as a filter and heuristic guide. The user's specialist knowledge may place a value on various subsets of the information that differs from this simple plausibility model. In this instance, let us suppose that, the fourth remainder (assertions Al and AlO) is selected as the least, valuable of the six. When these assertions are deleted, we have a maximally consistent subset of the data whose implications for the five indicators appear in Table T5.

Table T5: Consistent Inferences

Proposition	Validity
stocks-	F
interest-	F
taxes+	F
bonds+	T
deficit+	T

6. CONCLUSION

This paper has focused on a class of plausible reasoning tasks with three characteristics: inconsistent data, non-hierarchical interaction of concepts, and the need to obtain simultaneous readings on several hypotheses. A simple model with these attributes was used to demonstrate that existing systems for inexact inference are not suited to this kind of task. We first examined Prospector as the quintessential example of a Bayesian system and showed that both the model and Prospector itself would have to be altered to get any results at all. Even then, the inconsistency inherent in the given model was not made evident and a straightforward interpretation of the results turned out to be at variance with the model. INFERNO, a more tolerant non-Bayesian system, fared better in that the model did not have to be changed and its inconsistencies were discovered, but once more the attempt to wring a categorical interpretation from the results produced an anomaly. Ponderosa was introduced as a system to perform uncertain inference by finding maximal consistent subsets of the model, leading to results that are always categorical and that agree with whatever reduced model is used.

There are clearly other classes of plausible reasoning tasks to which Ponderosa is unsuited. If all the data is consistent or if there is a single proposition about which information is sought, the probability-bounding approach of INFERNO gives a better appraisal of the confidence with which the results can be accepted. This suggests an interesting possibility for combining the talents of Ponderosa and INFERNO. First, Ponderosa would be used to find whether the data is categorically consistent and, if not, to help the user choose a maximally consistent subset of it. INFERNO could then be run with this subset to supplement Ponderosa's categorical inferences with probability bounds. For instance, in the previous section we selected a maximally consistent subset A2 through A9 of the assertions in Table Tl. The analysis of this subset with INFERNO is shown in Table TG. It now becomes apparent that, while categorical inferences from the subset justify both the predictions that bonds will rise and that there will be a high deficit, the former conclusion has weaker probability bounds as a consequence of its derivation from less valid assertions.

Table T6: Combining INFERNO and Ponderosa

Proposition	Categorical Validity	Probability Bounds
stocks-	F	0 - .5
interest-	F	.25 - .25
taxes+	F	.15 - .4
bonds+	T	.11 - 1
deficit+	T	.45 - .79

Ponderosa has been implemented in Pascal and C for a VAX 11/780 minicomputer, based on a similar implementation of INFERNO. The prototype has been applied only to small tasks with less than 100 relations and propositions, and on these it is fast enough to be useful but considerably slower than INFERNO. For comparison, where INFERNO required just over a second to run the model, including finding rectifications, Ponderosa needed about 6 seconds.

ACKNOWLEDGEMENTS

I am indebted to Donald Michie of the University of Edinburgh and to Norman Shapiro of Rand for their insights and suggestions.

[1] Buchanan, B.G., "New Research on Expert Systems", Machine Intelligence 10 (J.E. Hayes, D. Michie, and Y-H. Pao, eds.), London: Ellis Horwood, 1982.

[2] Campbell, A.N., V. Hollister, R.O. Duda, and P.E. Hart, "Recognition of a Hidden Mineral Deposit by an Artificial Intelligence Program" Science, Vol. 217, September 1982, pp. 927-929.

[3] Duda, R.O., P.E. Hart, and Nils Nilsson, "Subjective Bayesian Methods for Rule-Based Inference Systems", Technical Note 124, Artificial Intelligence Center, SRI International, 1976.

[4] Duda, R.O., J. Gaschnig, and P.E. Hart, "Model Design in the Prospector Consultant System for Mineral Exploration", Expert Systems in the Micro Electronic Age (D. Michie, ed.), Edinburgh University Press, 1979.

[5] Garvey, T.D., J.D. Lowrance, and M.A. Fischler, "An Inference Technique for Integrating Knowledge from Disparate Sources", Proc. 7th International Joint Conf. Artificial Intelligence, Vancouver, 1981, pp. 319-325.

[6] Hayes-Roth, F., "Probabilistic Dependencies in a System for Truth Maintenance and Belief Revision", working paper, Rand Corporation, 1981.

[7] Konolige, K., "An Information-Theoretic Approach to Subjective Bayesian Inference in Rule-Based Systems", Draft, SRI International, 1982.

[8] Pearl, J., "Distributed Bayesian Processing for Belief Maintenance in Hierarchical Inference Systems", Cognitive Systems Laboratory, UCLA, Report UCLA-ENG-CSL-82-11, 1982.

[9] Pednault, E.P.D., S.W. Zucker, and L.V. Muresan, "On the Independence Assumption Underlying Subjective Bayesian Inference", Artificial Intelligence, Vo. 16, 1981, pp. 213-222.

[10] Quinlan, J.R., "INFERNO: A Cautious Approach to Uncertain Inference", Computer Journal, Vol. 26, 1983.

[11] Reiter, J., "AL/X: An Inference System for Probabilistic Reasoning", M.Sc. Thesis, Department of Computer Science, University of Illinois at Urbana-Champaign, 1981.