

# “Murdered by persons unknown” – Speculative Reasoning in Law and Logic

Jeroen Keppens<sup>1</sup>, Burkhard Schafer<sup>2</sup>

<sup>1</sup> *University of Wales Aberystwyth, Department of Computer Science, Aberystwyth, Ceredigion, SY23 3DB*

<sup>2</sup> *University of Edinburgh, School of Law, Joseph Bell Centre, Edinburgh EH8 9YL*

**Abstract.** This paper discusses some features of evidentiary reasoning in law, more specifically the introduction of “alternative suspects” by defence solicitors, within the framework of a first principle (or model) based decision support system for crime investigation. We introduce some observations from legal reasoning and legal doctrine to argue that the existing approach may not be able to capture some important distinctions and concepts of legal reasoning. We then introduce some ideas from dynamic logic and update semantics, in particular the concept of “pegs”, as good candidates to add the necessary expressive power to model based reasoning systems.

## 1 Introduction

In the late 80s, a string of high profile miscarriages of justice shook the foundations of the British legal system. In 1991, the Runciman Commission was established to examine the effectiveness of the criminal justice system in all its aspects. In its wake, a significant body of knowledge has been produced analysing the potential for errors in criminal investigations and prosecutions [1]. We have argued elsewhere [2] that the use of model based reasoning in crime investigations could address some of the most persistent problems in police investigations and criminal prosecutions that these studies identify. While the reader is referred to the original paper for the technical details, we will start by giving a short overview of the motivation behind the use of this approach, outline the main technical features and introduce a short example that will serve as a basis for the discussion later. We will then introduce some empirical data from legal reasoning and legal doctrine, more specifically the Scots law concept of “incrimination” and different formulation of the burden of proof in common law systems. We will argue that legal reasoning seems to assume some pertinent conceptual distinctions that cannot be represented in the system that we developed so far. Moreover, we will show that the problem relates to the interpretation of quantification, which means that all formal systems that use classical (referential) semantics are equally affected. In the concluding part, we argue that the use of “pegs”, a device used in dynamic logic and update semantics, could be a candidate for a possible solution to this problem and a number of related issues in legal reasoning.

## 2 Model Based Decision Support and Crime Analysis

### 2.1 Model Based Reasoning: why

One recurrent theme in analyses of miscarriages of justice is the problem of premature case theories. Instead of establishing in a neutral fashion what has happened, police officers tend to decide at a very early stage of an investigation on the most likely suspects, and from then

on investigate *against* them. The use of such “case theories” is probably inevitable [4]. The problem is therefore not the fact that case theories are used at all, but rather the restricted scope of alternatives that is considered. [5]. This problem is reinforced by the professional culture of the police service. Work is done properly, and a case solved, if a suspect gets convicted. This orientation towards positive results favours an “inductivist” ethos, where those pieces of evidence that points towards the guilt of the main suspects are seen as more valuable than those that would “falsify” the leading hypothesis. While the police service might pay lip service to a falsificationist model of rationality (“asking witnesses to come forward to eliminate them from the inquiry”) existing reward structures make it difficult to implement this in practice. Irving and Dunningham address possible solutions to this problem [6]. They argue for the need to improve officers’ reasoning and decision-making by challenging the “common sense” about criminals and crimes and the detective craft’s “working rules about causation, about suspicion and guilt, about patterns of behaviour and behavioural signatures.”

Our system addresses these two points by combining “backward chaining” abductivist reasoning with a “forward chaining” model that is based on the idea of indirect proof. Presented with all the evidence collected at a given point, the model based reasoner will first develop a range of alternative scenarios that causally explain the facts as established so far, by selecting generic reusable scenario fragments from the knowledge base and instantiating and composing them. This encourages creative speculation and the questioning of common sense assumptions. In the second step, it generates those (as yet uncollected) pieces of evidence that could differentiate between these different theories through a process of falsification. This fosters critical thinking away from the prevailing “inductivist” ethos.

A quick illustration explains what this means. Imagine a police officer arriving at a potential scene of crime. He notices a person, identified to him as the home owner, on the floor of a second floor flat, with injuries consistent to hits with a blunt instrument. The window of the room is broken, and outside a step ladder is found. Our officer, to make sense of the scenario as described above, will arrange (probably pre-linguistically) the features of the scene in coherent whole or Gestalt. In the same way as we cannot but see a forest when there are many trees, he will at a very early stage “see” a scenario in which a burglar entered with the ladder through the window, was approached by the home owner and killed him with a blunt instrument. This whole “picture” or “story” is influenced by typical associations, e.g. burglar with ladder. In our example, the evidence includes the broken window, the dead body, the wounds on the body and the ladder, and the preferred hypothesis is one of a burglary gone bad. The system should remind the officer for instance that on the basis of this evidence, it is also (though not necessarily equally) possible that the dead person did some Do-It-Yourself in his flat, suffered a heart attack while on the ladder, fell from the ladder to the ground, and the ladder fell through the window. This involves several “Gestaltswitches”: the ability to see the ground as a “blunt instrument”, the window as an opening that let things go out as well as in, and the entire scenario from one of crime to one of domestic accident. It should then identify those pieces of evidence that could rule out one of the theories, e.g. fingerprints (or lack of them) on the ladder and on the hammer that is considered the murder weapon, or fibres from clothing on the broken glass.

## *2.2 Model Based Reasoning: how*

We are giving here only a very abbreviated account of the formal mechanism which we have developed to distinguish murder from accidental death and suicide, using a (by now considerable) database of knowledge of homicide & suicide scenarios.

We use a model based reasoning technique, derived from the existing technology of com-

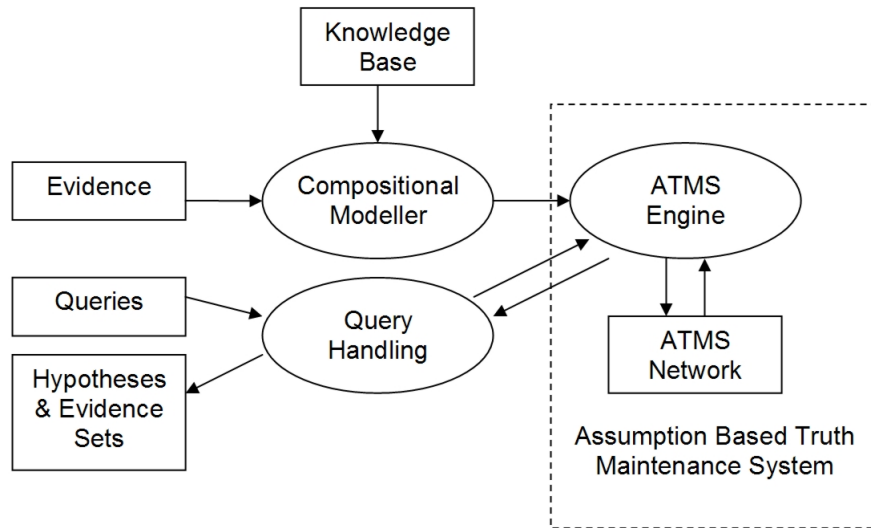


Figure 1: Basic Architecture

positional modelling, to automatically generate crime scenarios from the available evidence [7,8]. Consistent with existing work on reasoning about evidence the method presented here employs abductive reasoning. That is, crime scenarios are modelled as the causes of evidence and they are inferred based on the given set of evidence which they may have produced.

Figure 1 shows the basic architecture of the proposed model based reasoning DSS. The central component of this architecture is an *Assumption Based Truth Maintenance system* (ATMS). An ATMS enables a problem solver with an inference engine to reason about multiple possible worlds or situations. Each possible world describes a specific set of circumstances, a crime scenario in this particular application, under which certain events and states are true and other events and states are false. What is true in one possible world may be false in another. The task of the ATMS is to maintain what is true in each possible world.

An ATMS is mechanism that maintains how each piece of inferred information depends on presumed information and facts and how inconsistencies arise. For more details, the reader is referred to the original papers by de Kleer [9]. The goal of the decision support system (DSS) described in this paper is to find the set of hypotheses that follow from scenarios that support the entire set of available evidence. This set of hypotheses can be defined as:

where  $H$  is the set of all hypotheses (e.g. accident or murder, or any other important property of a crime scenario)  $S$  is the set of all consistent crime scenarios, and  $E$  is the set of all collected pieces of evidence.

In an ATMS, each piece of information of relevance to the problem solver is stored as a *node*. Some pieces of information are not known to be true and cannot be inferred from other pieces of information. In the ATMS, they are represented by a special type of node, called *assumption*. Inferences between pieces of information are maintained within the ATMS as *justifications* of the form  $n_1 \wedge \dots \wedge n_j \wedge \dots \neg n_k \dots - > n_m$  where the  $n_x$  are nodes (including assumptions) representing information the problems solver is interested in. The plausibility of these assumptions can then be determined through verification of nodes inferred from them. Continuing our “burglary gone bad” example (and bearing in mind that we are in an abductive environment), this could be understood as: the investigative hypothesis that it was an accident ( $n_m$ ) is justified by the presence of the victim’s fingerprints on the ladder ( $n_1$ ) and the absence of fingerprints of a third party ( $\neg n_k$ ). An ATMS can also store justifications called *nogoods* that lead to an inconsistency, i.e. justifications of the form  $n_1 \wedge \dots \wedge n_j \wedge \dots \neg n_k \dots - > \perp$ . The latter nogood implies that at least one of the statements in the antecedents must be false.

This accounts for the “critical” ability of our system: presented with two conflicting hypotheses, it will direct its user to collect evidence in a way that one of them is “justified” by a nogood, that is undefeated evidence that is incompatible with the investigative theory. The information maintained by the ATMS is generated by a *compositional modeller*. The compositional modeller employed herein has been devised to construct a space of plausible crime scenario by instantiating the knowledge base of scenario fragments and inconsistencies into an ATMS. The algorithm which we have developed expands on an existing composition modelling algorithm devised for the automated construction of ecological models [10].

As mentioned above, the compositional modeller generates the scenario space on to basis of a given set of evidence. To that end, it employs a knowledge base of scenario fragments. In our “burglary” example, the presence of large amounts of broken glass outside the house, implying that the window was broken from the inside, would be inconsistent with the accident hypothesis, but might lead to further hypotheses – the burglar broke the window in order to escape the house, or the home-owner fell head-first from the broken window (several hypotheses might include this node).

In this approach, it is presumed that the states and events constituting a scenario can be represented as predicates or relations. Naturally, states and events do not exist in isolation from one another. Certain states or events may be consequences of combinations of other states and events. For example, if a person is being assaulted and capable of self-defence, then (s)he will probably engage in some form of defensive action. Such knowledge is represented by *scenario fragments*. To illustrate the concept of scenario fragment, consider this example which we first give in formal notation, then in a verbal transcription:

if{ doctor(d), person(b), brain\_trauma(b) } assuming { cause-of-death(b,brain\_trauma), correct-diagnosis(d,cause-of-death(b)) } then { medical-report(d,cause-of-death(b), brain\_trauma) }

This scenario fragment states the following: given a person b, a doctor d and the fact that b suffered a brain trauma; and assuming that the cause of death of b is the brain trauma and that d makes a correct diagnosis of that cause of death; then a medical report must exist, written by d, stating that the cause of death of b is a brain trauma.

This scenario fragment can fulfil a dual purpose in an application. First and somewhat trivially, it ensures that the absence of a medical certificate is a reason to doubt that B died of a brain trauma – allowing in our example for the alternative explanation of “heart attack”. Secondly, it means that a medical report is not in any way different from say DNA evidence or a fingerprint: all are facts that are *explained by* certain assumptions. The medical report is *an observable consequence of* a state of affairs.

Once the compositional modeller has synthesised the scenario space, which is maintained by the ATMS, it can be analysed by the *query handling facility*. The approach taken here involves translating queries into formal ATMS nodes and justifications, thus enabling the existing ATMS label propagation to answer the queries of interest:

*Which hypotheses are supported by the available evidence?*

Every hypothesis that follows from a plausible scenario is supported by the available evidence. In our “burglary gone bad” example, there are two environments that support the available evidence. According to one, the house owner b suffered from a hit with a blunt instrument on his head, according to the other, he suffered heart failure, fell from the ladder and hit his head on the ground.

$E1 = \{ \text{high-cholesterol}(b), \text{accidental-coronary-blood-vessel-rupture}(b), \text{cause-of-death}(b, \text{heart-attack}), \text{correct-diagnosis}(d, \text{cause-of-death}(b)) \}$

$E2 = \{ \text{hammer-attack}(p, b), \text{brain-trauma-due-to-attack}(b), \text{cause-of-death}(b, \text{brain-trauma}), \text{correct-diagnosis}(d, \text{cause-of-death}(b)) \}$

In the possible world described by environment  $E1$ , accident(b) is true and in the one described by  $E2$  homicide(b) is true. Therefore, it follows that both hypotheses are supported by the available evidence.

*What additional pieces of evidence can be found if a certain scenario/hypothesis is true?*

All the states and events, including pieces of evidence, that are logical consequence states and events in plausible scenarios are generated in the forward chaining phase of the algorithm. Therefore, the initial state of the ATMS will contain nodes representing pieces of evidence that are produced in certain scenarios but were not collected in  $E$ . A piece of evidence  $e$  can be found under a given hypothesis  $h$  if a possible world exists that supports both the evidence and the hypothesis. Continuing with the ongoing example a piece of evidence  $e$  that consists of a medical report documenting high cholesterol in  $b$ , medical-report( $d$ , high-cholesterol( $b$ )), is generated under the environment:

$E3 =$  {high-cholesterol( $b$ ), accidental-coronary-blood-vessel-rupture( $b$ ), cause-of-death( $b$ , heart-attack)},  
correct-diagnosis( $d$ , cause-of-death( $b$ )) test( $d$ , test-cholesterol( $b$ ))}

This means simply that under the hypothesis of accident, this third piece of evidence, a report, may be found.

*What pieces or sets of additional evidence can differentiate between two hypotheses?*

Let  $h1$  and  $h2$  be two hypotheses, then any set of pieces of evidence  $E$  that can be found if  $h1$  is true, but are inconsistent with  $h2$ , can differentiate between the two hypotheses. For example, it follows from the above discussion that the piece of evidence medical-report( $d$ , heart attack( $b$ )) may help to differentiate between the two hypotheses, accident( $b$ ) and homicide( $b$ ). This information suggests to a police officer or a prosecution lawyer examining the case that ordering tests for symptoms of heart attack would be useful.

### 3 A Legal Puzzle

In developing alternative scenarios consistent with the evidence, the ATMS performs some of the scrutiny a good defence solicitor would subject the prosecution case to. A defence solicitor has broadly speaking two strategies available to him. First, he can question the factual correctness or the legal admissibility of evidence presented by the prosecution. Second, he can accept the evidence at face value and argue that alternative explanations for their presence are possible that do not incriminate his client. We are concerned here primarily with this second strategy. However, it is here that we encounter a certain ambiguity, an ambiguity explicitly recognised by the Scots law of evidence. The defence has in fact again two strategies available to it. The first can be dubbed the ‘‘Perry Mason Stratagem’’. Like the fictitious advocate, the defence can pursue its own investigation and ‘‘point to the real culprit’’. In Scots law, this is known as the special defence of incrimination [18], recently used (unsuccessfully) in the Lockerbie trial. This strategy has a number of psychological and legal advantages. The same reason that makes it the solution of choice for crime writers also works well with juries: no loose ends are left and the crime is avenged. Procedurally, it allows the defence to submit also other pieces of evidence. This corresponds to the ‘‘forward chaining’’ aspect of our ATMS: The party named by the defence will have interacted causally with the crime scene. This will have created evidence which can strengthen the defence case. This allows introduction of additional ‘‘suspect specific’’ evidence (such as alibi) evidence about other people, which otherwise might be ruled out as irrelevant. The defence of course need not prove the guilt of the other party; it only needs to establish it as a plausible alternative. This way of dealing with the evidence mirrors particularly closely the working of our ATMS.

In reality however, this strategy faces considerable obstacles. Defence solicitors normally don’t have the resources, time or training to engage in investigative activity of their own.

In developing alternative accounts of the evidence, they will therefore typically settle for something less. They will argue that a hypothetical "someone" might have been the real culprit, and that this alternative is not ruled out by the evidence.

A (fictitious) example of a cross examination can illustrate this point:

Police officer: When I heard that scream, I ran up the stairs to the room where the victim was. There I found the accused, with a bloodied knife in his hand

Defence solicitor: Is it not true that it took you more than three minutes to find the right room? In this time, the real murderer could have escaped through the window, couldn't he? So that when my client arrived at the scene shortly before you, he found the body of his wife lying there, and took the knife in an attempt to guard himself against whoever killed her?

At this point, the prosecution might accept this alternative "for the sake of the argument" when re-examining this witness:

Prosecution: According to your description of the scene of crime, is it not true that this "mysterious Mister X" would have had to jump down from the third floor, and then run across a busy street with blood all over his clothes?

The argument proposed here is that the use of "someone" in this strategy is not best represented as a classical existential quantifier that picks a specific object from the universe of discourse. This is particularly important to the prosecution, who may fear damage to their case from a statement that carries an ontological commitment to a "Mister X".

What happens in this example is categorically different from a defence of incrimination. In an incrimination defence, defence and prosecution disagree about the state of the world. In the dialogue by contrast, what is debated is not so much a state of affairs. Rather, what is exchanged is "discourse information". We can understand the defence argument as a meta-statement about the evidence: "It does not follow logically from the case as presented that my client is the person who committed the crime". Conversely, the prosecution is stating: "the alternative of the defence is inconsistent with the evidence". We note in the passing that this distinction finds its equivalent in another contested part of legal doctrine. At present, the law treats the following two expressions as synonymous for instance for the purpose of jury direction by the judge in his summing up:

1. You must be convinced of the guilt of the defendant beyond reasonable doubt
2. You must be convinced of the guilt of the defendant beyond *a* reasonable doubt

However, empirical studies have shown that juries that have received the first instruction are more likely to acquit than those that received the second. Here too, we can understand the first instruction as requiring merely discourse information: What was the quality of the case of the prosecution? A "not guilty" in the second case however carries an ontological commitment: to acquit; there has to be some object, an alternative story or an alternative suspect that functions as "the" doubt.

Evidentiary legal argumentation typically involves a combination of exchanges about "discourse information" (the rules of admissibility might be seen as a particularly prominent example) and discussion about the world. For this reason alone, it might be worthwhile to amend the formal system in a way that both types of reasoning can be represented *in their distinctiveness*. To conclude this passage, and to show the pervasiveness of the issue discussed here, another example is given that at first sight is unrelated to the introduction of

hypothetical suspects. It will turn out, however, that solutions discussed in natural language analysis for just this problem also provide a natural solution to the issue at hand. Consider the following scenario.<sup>1</sup> Several witnesses claim to have seen “someone” suspicious running away at the time of the crime.

Witness 1: Yeah, he was a black guy, with a red cap, running to his car. It was a Volvo

Witness 2 (wife of 1): It was not a man, it was a girl, and she wasn’t wearing a red cap but a red scarf. Also she wasn’t really black, more tanned. Oh, and it was a Vectra

This exchange is intelligible only if witness 2 has reasons to believe that the indefinite “a black guy” of statement 1 refers to an entity which witness 2 thinks to be a woman with a red scarf. None of the original attributes of this entity seem to be agreed upon. The police officer needs to be cautious: it might turn out after all that they saw two different people and our ATMS in turn needs to be able to cope with partially conflicting descriptions of what might or might not be the same object.

#### 4 Dynamic Logic and the Concept of Pegs

The linguistic phenomenon described in the last example has been at the centre of research in formal linguistics for quite some time and produced copious observations and theories [12, 13]. The question for linguists has been, how is it possible for a pronoun to be bound without being in the syntactic scope of its antecedent? The obvious answer is that the semantic scope of an expression may reach beyond its syntactic scope. Formalizing this insight was less straightforward, however. The first theory to implement this idea was Discourse Representation Theory (DRT) [14]. We want to draw the attention here to the “Dutch school” which, inspired by DRT, developed a family of formal systems which are particularly suited to address the issue not only of anaphoric binding, but more generally the “unspecific” use of quantifiers discussed here.<sup>2</sup> They also allow a meta-linguistic treatment of the kind of discourse information that we mentioned above, allowing for a more unified theory of evidentiary legal reasoning. Update semantics, dynamic logic and data semantics are varieties of this approach. Technically, they are extensions of standard model theoretical (Kripke) semantics, which should make their incorporation into our ATMS straightforward. However, they radically reinterpret the meaning that they give to Kripke models. Frank Veltman, in his influential paper on update semantics [15], coined the slogan that summarized the unifying assumption of these different formalisms: “You know the meaning of a sentence if you know the change it brings about in the information state of anyone who accepts the news conveyed by it”. Meaning thus is not (just) a relation between sentences and the world, but the potential to change a context, where contexts are identified with information states. The recursive definition of semantic interpretation is formulated in the language of updating information states with sentence in such a way that meaning becomes identifiable with update functions [11]. It is this move away from “truth and denotation” to “discourse and information” that ultimately allows incorporation of the above mentioned “discourse information”, previously treated in informal pragmatic theories, directly into the formal representation. Generally speaking, discourse information concerns matters relevant to the process of linguistic interpretation. In

<sup>1</sup>Modelled on an example given by Groenendijk, Stockhof and Veltman. See [11] p. 3

<sup>2</sup>a typical example of anaphoric binding are “donkey sentences” of the form: Every farmer who owns a donkey beats *it*. “Every” in the first NP binds anaphorically “it” in the second noun phrase. To model the semantic of these sentences, the quantifier has to bind beyond its syntactic scope. An intensive discussion of the linguistic data and current approaches to this problem with further references can be found in [19]

informative discourses, it is normally subservient to the primary goal of acquiring information about the world, though we would argue that legal discourses are an exception to this rule.

To deal with anaphoric binding, it became necessary for these systems to introduce as part of the discourse information a new type of objects, commonly called "pegs". "Pegs" are intermediate discourse entities. They are connected to the variables in the object language on the one hand, and with the objects in the models (our scenarios) on the other.

An information state as a whole consists not only of discourse information (in the form of a referent system), but also of information about the world, and of a link between the two types of information. An information state in this approach is regarded as a set of possibilities, an assumption that resonates well with the concept of "scenario space" in our ATMS, only that now the very meaning of pieces of evidence is seen to be determined by the possible scenarios in which it appears. Each possibility consists of a referent system; a possible world; and an assignment function which assigns some object from the domain of that world to each of the pegs present in the referent system. Information growth can take place in two ways: the referent system may be extended with new pegs, (re)associating variables with them and assigning them suitable objects; and/or certain possible assignments or possible worlds may be eliminated. This last possibility corresponds to the "forward chaining phase" of our ATMS.

The notion of update and exchange of information about the values of variables has some intuitive appeal. As the examples above show, it comes naturally to us to talk about 'indefinite' objects, objects of which we have only partial knowledge. We talk about these objects (the mysterious Mr X), they are ascribed (possibly conflicting) properties and people are informed about their existence. Crucially though, these 'objects' must not be understood as classical objects. From the perspective of agents with partial information the ultimate identity of such objects may be left unresolved. Nonetheless, they can be topics of information exchange. Fred Landman has developed the most explicit theory of such partial objects to date [16], even though the idea traces back to Karttunen's seminal paper [17]. According to Landman, they are things that don't have properties, but to which properties can nonetheless be ascribed and, similarly, things that don't have identity conditions, but that have identity conditions ascribed to them:

"the essence of partial information is that it cannot justify certain distinctions, and the decision about the identity of certain pegs is a prime example of that".

The proposal that we are making here is that the assumption of purely formal objects or pegs allows development of theories that explain the peculiar features of legal discourse noted above. While the defence of incrimination introduces a classical object into the discourse, the more abstract speculation about "possible other parties" introduces (merely) pegs. The prosecution, as we have seen, can in turn refer to these objects – without incurring a commitment to the existence of any corresponding person in the universe of discourse. Similarly, prosecution and defense stories can accommodate witness accounts of "the same" person even though this "person" is ascribed incompatible attributes.

A final example, taken again from [11], can illustrate the convergence between the legal issues of this paper and the intuitions about natural language that are formally expressible in dynamic logic. Groenendijk *et al* discuss sentences of this type:

Peter is guilty. He may not be guilty.

They observe that while inconsistent as the utterance of one speaker, it can be acceptable (provably so in their formalism) as a debate between two speakers. They argue that confronted with the claim  $p$ , one might well retort by pointing out that according to one's own



information the possibility of  $\neg p$  is not excluded. According to their analysis, this is not a case of straightforward disagreement about the world, as there are information states which support both utterances. In this case, the second speaker rather displays that she has little trust in what the first speaker said. She is “non-accommodating”. A different game, so Groenendijk, is being played here, in which exchanging information about the world apparently is not the first objective. This corresponds to our analysis regarding the role of defense arguments involving possible alternative perpetrators. They are not a contribution to world-information, but a statement about the trustworthiness of the prosecution case. Unlike in the form of discourse discussed by linguists, formal attribution of burden of proof in legal discourses means that these pegs may be persistent, and not as easily discarded after they served a temporary function of anaphoric binding. In adding pegs to our ATMS, speculation about alternative scenarios can get a much wider scope, and moves closer to natural language argumentation; in short, closer to the reality of legal argumentation.

## 5 Conclusion

Evidentiary reasoning in law is characterized by a dialogical exchange between parties with only partial, and possibly contradictory information. Speculation about possible alternative models supported by the established evidence plays a crucial role in investigation and prosecution. Rules of legal procedure ensure that this speculation might as be both about the state of the world (referential interpretation of quantification and “truth”) and about the legal discourse that established truth. Both forms of speculation require provisional reference to objects other than the classical objects of model theoretical semantics. While ontologically promiscuous, the introduction of formal objects without identity criteria to ATMS allows development of an understanding of evidentiary reasoning that is truer to the empirical linguistic data. This is not only desirable for the purpose of the design of decision support systems such as the one envisaged in the first part of this paper. Legal confusion over the precise nature of the incrimination defence and its relation to other types of defence arguments has already resulted in fatal misdirections of juries [18 p.30]. Generally, the nature of speculative legal reasoning, its scope and limits are a field of evidence law as yet under-researched. The proposed formal analysis might offer valuable insights in this field, while the linguistic field data compiled for the purpose of formal semantics provides a valuable yardstick against which legal intuitions and conceptual choices can be judged.

### Acknowledgment

The authors have been supported by the UK EPSRC grant GR/S63267

## References

- [1] S. Walker, *Justice in Error*. Blackstone, London 1993.
- [2] J.Keppens, J. Zeleznikow, J. A Model Based Reasoning Approach for Generating Plausible Crime Scenarios from Evidence. *Proceedings of the 9th International Conference on Artificial Intelligence and Law*, 51-59 2003
- [3] D. Dixon, Police Investigative procedures. In: C. Walker (ed.) *Miscarriages of justice*. Blackstone, London 1999
- [4] McConville, Sanders and Leng: *The Case for the Prosecution*. Routledge, London 1991
- [5] S. Greer, Miscarriages of criminal justice reconsidered. *Modern Law Review* **58** (1994) p.71
- [6] B. Irving, C. Dunningham, Human Factors in the quality control of CID investigations and a brief review of relevant police training. *Royal Commission on Criminal Justice Research Studies* **21** (1993)

- [7] B. Falkenhainer, K. Forbus, Compositional modelling: finding the right model for the job. *Artificial Intelligence* **51** (1991) p. 95-143
- [8] J. Keppens, Q. Shen, On compositional modelling. *Knowledge Engineering Review* **16** (2001) p. 157-200
- [9] J. de Kleer, An assumption-based TMS. *Artificial Intelligence* **28** (1986) pp 127-162
- [10] J. Keppens, Q. Shen, Compositional model repositories via dynamic constraint satisfaction with order-of-magnitude preferences, *Journal of Artificial Intelligence Research* **21** (2004) pp 499-550
- [11] J. Groenendijk, M. Stokhof & F. Veltman, Coreference and modality. In S.Lappin (ed.), *The handbook of contemporary semantic theory* , Oxford, Blackwell 1996 pp179-213
- [12] Gawron, J. M., Peters S.: *Anaphora and Quantification in Situation Semantics* . CSLI, Stanford 1990
- [13] B. Geurts, Donkey business. *Linguistics and philosophy* **25** (2002) pp129–156
- [14] H. Kamp, U. Reyle, *From discourse to logic* . Dordrecht, Kluwer 1993
- [15] F. Veltman, Defaults in update semantics, *Journal of Philosophical Logic* **25** (1996) pp. 221-261
- [16] F.W. Landman, *Towards a theory of information* . The status of partial objects in semantics. Foris, Dordrecht, 1986
- [17] L. Karttunen. Discourse referents. In J. McCawIcy, editor, *Syntax and Semantics* , Vol. 7: Notes from the Linguistic Underground. Academic Press, New York, 1976
- [18] D. Field, F. Raitt. *Evidence* , Edinburgh, W.Green, 1996
- [19] J. Koster, and E. Reuland (eds.) *Long-distance anaphora* . Cambridge: Cambridge University Press 1991