



Artificial intelligence and the evidentiary process: The challenges of formalism and computation

RONALD J. ALLEN

*John Henry Wigmore Professor, Northwestern University School of Law, 357 East Chicago Avenue,
Chicago, IL 60611, USA*

Abstract. The tension between rule and judgment is well known with respect to the meaning of substantive legal commands. The same conflict is present in fact finding. The law penetrates to virtually all aspects of human affairs; virtually any interaction can generate a legal conflict. Accurate fact finding about such disputes is a necessary condition for the appropriate application of substantive legal commands. Without accuracy in fact finding, the law is unpredictable, and thus individuals cannot efficiently accommodate their affairs to its commands. The need for accuracy and predictability in legal fact finding has generated a search for formal tools to apply to the task. Among the tools that have been examined are Bayes' Theorem and expected utility theory (Bayesian or statistical decision theory). These tools do not map well onto trials, which in turn has generated an examination of alternative approaches, in particular the story model and the relative plausibility theory. This paper discusses these issues in turn. It elaborates the basic structure of trials in the American tradition; examines the uneasy relationship between trials and such formalisms as Bayes' Theorem and expected utility theory; and introduces the relative plausibility theory as an explanation of the nature of juridical proof.

The domain of the law in western societies is the domain of human affairs. Every act that a person takes is surrounded with a web of regulation formed from the commands and dispensations of constitutions, statutes, ordinances, the common law, executive orders, and administrative rules (and I have probably left out a few sources). The web is multi-faceted both in its construction and in its implications. It creates rights against others, but also obligations to others. It expressly permits without requiring some actions, and permits but disfavors others. There are innumerable cross references, ranging from those that are internal to codes to others that are between separate bodies of law emanating from independent sovereigns, and there are endless exceptions. Some of law's commands are explicitly admonitory; some are ambiguous; some are linguistically, others conceptually, incoherent. In addition to the formal commands contained in the web, the web provides a means of resolving virtually any kind of conflict that can emerge as a result of human interaction. And, the most banal or domestic of acts can lead to significant legal questions in need of resolution. A mother feeds a child some cereal only to notice after the child has finished a rat's tail in the box. And so breakfast becomes the start of a law suit. A driver turns into the sun, reaches for his sunglasses and at that moment strikes a pedestrian. And so driving home at night initiates a legal dispute.

A neighbor parks her car in such a way that makes access to my garage difficult, and another neighbor plays loud music late at night, both of which could result in a legal dispute calling for an authoritative resolution.

Because any human action (and many nonactions, too) can generate a conflict in need of authoritative resolution upon which individual lives and fortunes may rest, there is the need to provide clarity in the law, to reduce the law's commands to a set of necessary and sufficient conditions so that rights and obligations are knowable in advance, thus reducing the potential for arbitrariness, yet the scope of the demand for authoritative articulations of rights and obligations vastly outstrips the human capacity to provide it. The world is too complicated to capture in rules. The best example from the United States is the effort of the Internal Revenue Service to reduce taxation to a set of necessary and sufficient conditions through a complex code and considerably more complex regulations. The effort has failed. As soon as amendments to either code or regulations are promulgated, the tax lawyers and accountants find ways around them, thus generating yet more amendments and additions, and so on. In one quite small area of the law, and with enormous resources put to the task, the aspiration to legal clarity has proven impossible to achieve. *Res ipsa loquitur*, as we say in the law.

What then do you do? You have no choice but to view the legal system in large part as an *ex post* process, rather than simply as an *ex ante* generator of definitive norms – and thus the growth of the common law, both the common law in pure terms and the common law driven by statutory interpretation. Fair minded individuals – judges, typically, but sometimes juries – are asked to use their native cognitive capacities and reach results in cases that no one anticipated would arise, largely through analogical reasoning. But, that analogical reasoning is always defeasible. One case about a rat tail in a cereal box may control the next case involving an adulterated substance, but there may be something about the next that distinguishes it from the first, and so on.¹ And so off goes the law in a never ending process of looking for useful analogies in an attempt to reduce, even though elimination is impossible, the problem of *ex post facto* rule making (a problem because of its denial of the central aspiration of law to lay out in advance a person's rights and obligations in order to allow one to accommodate to them).

Note that the complexity I am describing has two different components. One is standard computational complexity, and the formal complexity of the law is daunting indeed. But the conceptual problem presses considerably deeper. The problem is not just the meaning and relationships of presently existing authoritative commands. It is that many of what will turn out to be authoritative commands have not even been thought of yet, and could not be. They will emerge only as consequences of the facts of the disputes brought into the legal system in one form or another.

¹ For example, on 16 December 1997, CNN on line reproduced an Associate Press story about the conviction of Michael F. Zanakis for extortion and fraud. Zanakis allegedly took a rat's tail from the research laboratory where he worked, fried it, and placed it in a package of french fries that he purchased at McDonald's for his two year old son. He then demanded a five million dollar settlement.

This has left the law with no formal rational mechanisms for formulating emerging standards, explains why there is so much reference to discretion and judgment in the law, and why the concept of a “precedent” is fraught with ambiguity.

Although artificial intelligence is not my field, obviously work is being done that may shed light on the manner in which norms emerge in the legal system, such as by deepening our understanding of the nature of analogy and case based reasoning. This has the potential of helping those of us in the law to better understand what terms such as “discretion”, “judgment” and “precedent” may mean, although doing so is likely to be very difficult indeed, precisely as a consequence of the complexity of the human condition. The present special issue is designed to highlight the work being done on artificial intelligence and legal evidence, on the possible contributions of formal modeling to the understanding of the nature of juridical proof. I am glad to learn of this effort. Concentrating, in the words of the organizers of the special issue on “the application of paradigms from AI and formal reasoning to legal evidence”, could be of the utmost value to the legal process. The evidentiary process at trial, where the rules of evidence have their direct impact, is infused with even more ambiguity, uncertainty and openendedness than the law generating process described above. There are millions of disputes every year that lead to at least the first stage of litigation, virtually no two cases are identical in the evidence produced at trial, and subtle differences among cases can generate opposing inferences.

Moreover, the relevant information for decision at trial is constantly changing as a consequence of the natural learning processes of the humans who decide cases – judge or juror. Consider a person called to decide a case last year and another this year. Over the intervening year, the decision maker has had vast sensory inputs and opportunity for reflection, and thus brings a changed conceptual apparatus to decision in the second case. This conceptual point has been a primary impediment to the systematizing of legal proof, for any systematization apparently would have to anticipate how – and what – individuals learn in order to be useful. Nonetheless, the absence of systematized proof processes – of much at all that looks like “rules” – gives rise to identical concerns about whim and caprice that animate the concern over the inability of the law to specify its commands clearly in advance. Both seem to have the capacity to generate unpredictable, idiosyncratic decisions that are the antithesis to the rule of law. Any assistance in this regard that artificial intelligence can provide would be most welcome indeed.

There is a second, more localized, problem that has impeded the development of formalisms within the proof process, which is that trials in the American system have certain attributes that impede the mapping of one promising formalism – Bayes’ Theorem – onto the trial process, and that make the implications of other promising formalisms—such as expected utility theory – uncertain. The relationship between Bayes’ Theorem and trials is the focus of a recent special issue of Blackstone Press’s *International Journal of Evidence and Proof*. I was the editor for the issue and provided the lead article that attempted to demonstrate the difficulties in

mapping Bayes' Theorem onto trials. The primary difficulty is the incompatibility of the formal requirements of subjective Bayesianism and the structure of trials in the American tradition. This leads to a general point: Any effort to systematize the proof process must accommodate, or provide an alternative to, the actual structure and operation of trials. Thus, an understanding of the nature of trials is necessary for those working within artificial intelligence and law. Perhaps also an understanding of how the law has responded to the concerns I have articulated would be helpful, as well as an explanation of the fate that has befallen some of the formalizations that have been advanced as descriptive or prescriptive for the law. These points constitute the remainder of this paper.

1. The nature of trials and the meaning of "evidence"

The pre-trial condition of fact finders at trial is largely one of ignorance. This is especially true of jurors, less so of judges. Jurors will be largely ignorant of the applicable law and the material to be produced at trial. The judge will know, or have the capacity to determine, the relevant law, and may have some knowledge of the material likely to be produced, although the extent of that knowledge will vary dramatically from case to case. In civil cases with extensive pre-trial proceedings, the judge may be relatively well informed; in criminal cases, the judge will be about as ignorant as the typical juror.

Neither pre-trial instructions nor opening statements impart significant substantive legal knowledge.² Jurors are typically informed of the legal charge, but not the elements. In some instances they will have some knowledge of the elements from their conventional knowledge; in many others they will have essentially no knowledge whatsoever, such as in a civil suit involving the intricacies of the Tax Code. Opening statements impart some but little knowledge. They are typically brief overviews of what witnesses are likely to say. Although there is a story told, it is not one to which the party, or the fact finder, is bound. As one trial judge instructed the jury: "The final thing before we get under way with the evidence is that the attorneys have the opportunity to make opening statements to the jury and these are very brief outlines the attorneys could give you if they wish showing what they expect the evidence in the case is going to prove. Sometimes it happens, of course, that the evidence comes out differently from what they expect" (Allen et al. 1997, pp. 8–9).

A typically disjointed presentation of evidence begins in which no one tells the whole story in many instances. Typically, it is put together in bits and pieces. Cross-examination raises doubts by indirection:

² The relative absence of substantive legal knowledge from preliminary instructions and openings can be determined by reviewing trial transcripts. The points I make henceforth can be confirmed by reviewing two transcripts, *People v. Steele*, which is reproduced in Allen and Kuhns (1989, pp. 3–95), and *People vs Johnson*, reproduced in Allen et al. (1997, pp. 4–118).

Finally, an effective cross examination often succeeds through the use of implication and innuendo. It is not necessary, and it is often harmful, to ask a witness the “ultimate question”. Final argument is your opportunity to point out the relationship between facts, make characterizations, and draw conclusions based upon the accumulation of details (Lubet 1993, p. 60)

The salient point is that fact finders will not be in a position to evaluate evidence until the end of the trial.

At the end of the evidentiary process, closing arguments are made. It is only here that a clear and detailed picture of the parties’ various assertions is presented. The lawyers articulate their respective theory of what happened, and marshal the evidence on behalf of the theories. Not only what was said at trial, but what reasonably can be inferred from what was said, and indeed from the process as a whole, form the basis of closing argument. The implications, as the lawyer sees them, of the subtle questions that were asked on cross-examination are now advanced and examined. In essence, each party is claiming that its preferred version of what happened is better captured by the evidence as it was produced at trial.

Final instructions are given, and “deliberations” begin, whether collaborative in the case of juries or solitary in the case of judges. To “deliberate” usually refers to analyzing the evidence from the individual perspectives of the jurors in an effort to reach a consensus as to what happened. But, there is a preliminary process that unites the activity of judge as fact finder and juror. In both cases, the decision maker must reflect upon the evidence from his or her own standpoint and decide, at least preliminarily, what it means. With jurors, this individualized appraisal is then followed by group discussions in an effort to forge a consensus. At both levels of individualized consideration of the evidence and group deliberation, an enormous wealth of knowledge and beliefs will be brought to bear on the meaning of the evidence. Evidence does not announce its own implications. Those implications come from the interaction of what is formally referred to as “evidence” and belief structures of the fact finder. It is thus a mistake to think of evidence as what is produced at trial. It is instead the outcome of the process of appraising what is produced at trial by the fact finder, who in doing so invokes a large storehouse of “evidence” that is summed up in its beliefs (Allen 1994).

Every decision maker will have an idiosyncratic belief set precisely because no two humans have lived the same lives. Therefore, just as the law is significantly *ex post facto*, so too decisions on facts cannot, in large measure be known in advance. They emerge as consequences of the process of decision, which means that the implications of evidence is in many – indeed most – instances unknowable *a priori*. Thus, the attempt to apply formal reasoning methodologies to the “evidence” at trial risks misconceiving the process being modeled. It is a creative rather than a static one. In virtually all cases, no discrete implications flow from any particular evidentiary proffer. Consider:

1. Does the blood on the gate of Nicole Simpson's house imply that O.J. Simpson committed the crime, or that it was placed there by the police in a frame-up, and thus that he did not do it?
2. A witness is sweating and stammering. Is this evidence of lying or of a truthful witness intimidated by the process?
3. Four witnesses tell essentially the same story about what they saw. Does each story confirm the others, or do the four of them together confirm the witnesses are in a conspiracy to defeat the truth?

And so on.

"Evidence" in law, then, is not like "evidence" in science in the context of controlled experiments. Rather, "evidence" in law is much closer to the "evidence" marshaled in science on behalf of competing theories that is adduced in order to precipitate changes in belief. The differences between evidence at trial and evidence resulting from a controlled study are at least two fold. First, in controlled studies, the data are typically not problematic (there were observable traces in the bubble chamber or there were not), although the interpretation of the data often is. In the law, inconsistent primary data ("the light was red" – "no, it was green") are the norm, and replication is generally impossible. Second, scientific knowledge is typically organized in a hypothetico-deductive fashion. Even those advancing competing theoretical conceptions typically share with their competitors methodological and mathematical principles, agree on what counts as evidence, and are able to express their scope of disagreement comprehensibly.

Now we see the crucial distinction between factual inquiry in the legal process and scientific investigation. Science proceeds through the simplification of phenomena through controlling as many variables as possible. Judgments about ordinary events, by contrast, rarely can be recast as the results of controlled experiments. Too many variables are constantly and necessarily in play.

We thus reach a counter-intuitive but, I believe, true point, and one on which I would expect to find AI people in agreement: Lay knowledge is more, not less, complicated than scientific knowledge because it ranges over all human affairs.³ Faced with such wide ranging and intractable ambiguity, the law forgoes formalized theories of evidence and instead simply locates decision over what is evidence

³ For example, Rashkin and Yuelle (1989) have developed a method for analyzing interviews of children in alleged sex abuse cases. This of course is just one of many kinds of interviewing problems the law faces. According to Rashkin and Yuelle's work, the following criteria are relevant to this process: A. General Characteristics of the Allegation (1. Coherence, 2. Spontaneous Reproduction, 3. Sufficient Detail); B. Specific Contents of the Allegation (4. Contextual Embedding, 5. Descriptions of Interactions, 6. Reproduction of Conversation, 7. Unexpected Complications during the Incident); C. Peculiarities of the Content (8. Unusual Details, 9. Peripheral Details, 10. Accurately Reported Details Misunderstood, 11. Related External Associations, 12. Accounts Subjective Mental State, 13. Attribution of Perpetrator's Mental State); D. Motivation Related Contents (14. Spontaneous Corrections, 15. Admitting Lack of Memory, 16. Raising Doubts about One's Own Testimony, 17. Self Deprecation, 18. Pardoning the Perpetrator); E. Offense Specific Elements (19. Details Characteristic of the Act).

and its implications somewhere and more or less lives with the outcome. The proffered data become evidence if they influence a fact finder. Whether they do is determined by the sum total of that person's experiences at the moment of decision, experiences which will by that time include the advocates' efforts to enlighten the fact finder about the implications of the material produced at trial and all the other observations generated by the trial. As I have developed this concept elsewhere:

Thus the banality of the theory of juridical evidence. Evidence is not a set of things, as the conventional theory would have it; it is instead the process by which fact finders come to conclusions about the past. It is banal in the sense that it reduces to the proposition that a disinterested fact finder reconstructs the past based on all the observational inputs available at the moment of judging. But with its banality comes genius. The genius of the theory of juridical evidence lies in the twin recognition that there is no alternative except official orthodoxy on conventional matters and that the probability of gathering the necessary information for an accurate reconstruction of the past increases astronomically with the size of the fact finding body and not just with the size of the evidentiary proffer at trial. Each fact finder is in essence a solitary scientist constantly reducing conventional ambiguity. This type of ambiguity is treated differently from the scientific variant because resources would be wasted by attempting to reach explicit agreement on its contours, and so we do not organize assaults on it. But when a dispute arises involving it, we gather a small set of individuals, from one (a judge) to twelve (an historical jury), to pool its members' knowledge in order to make sense of the evidence adduced. Given the lack of a social justification for organizing conventional knowledge on a scientific model, the law achieves an analogous result by holding a convention of lay scientists and requiring that they deliberate long enough to reach agreement. (Allen 1994, pp. 627–628)

2. The legal regulation of proof: The rules of evidence

The convention of lay scientists is, however, a curious one. It is part active and part passive. It is quite active during the deliberative phase in the sense that each individual juror's own knowledge is brought to bear on the questions at hand, and may in fact be discussed openly by the jury, in that sense becoming evidence in the case. By contrast, it is quite passive during the formal evidence taking process. The tradition in the United States has been of virtually total passivity, without so much as questions from the jury being encouraged. This is undergoing very slow change with some states experimenting with a more active jury, but the norm remains passivity.

The information provided to the convention of lay scientists, or a trial judge at a bench trial, is largely, but not completely, that which the parties wish to produce. The parties are constrained in what they can produce by the rules of evidence,

which are the law's response to the concerns of idiosyncratic decision discussed at the beginning of this paper. There are, roughly speaking, three animating motivations to the rules of evidence: uncertainty is regulated through rules of decision in the forms of burdens of proof, inferences, and presumptions; appeals to sentiment, bias, and other kinds of influences that are potentially irrational from the point of view of accurate fact finding are regulated through relevancy and materiality requirements; and various policies are advanced independently of their affect on accurate fact finding. I discuss these three points in turn.

2.1. THE REGULATION OF UNCERTAINTY

Legal decision is always decision under uncertainty. It is widely if not universally believed that the law should determine how to decide in the face of uncertainty rather than let each fact finder decide idiosyncratically. Consequently, one typically finds a quite brief rule (or small set of rules) and a robust set of cases that regulate this uncertainty in various way. At the most general level, each substantive cause of action has an associated burden of persuasion that specifies the degree to which the fact finder must be persuaded in order to return a verdict (and a burden of producing evidence, an important point whose implications cannot be developed here because of space limitations). In civil cases, the general standard is that each element of the cause of action must be established by the plaintiff to a preponderance of the evidence, which is normally understood to mean anything greater than a 0.5 probability. For a whole host of reasons, modifications to this general rule exist, sometimes raising and sometimes lowering a plaintiff's burden of persuasion. In addition, some issues are deemed defenses, and the burden of persuasion on those is given to defendants. In criminal cases, because of constitutional command, the prosecution must prove every element of the offense beyond reasonable doubt. *In re Winship*, 397 U.S.358 (1970). Again, though, certain issues are deemed defenses, and their burden of persuasion is often allocated to defendants.

Uncertainty is also regulated through a robust set of instructions on inferences that explicitly commend particular inferences to the fact finder, judge or jury. This is done because of beliefs that have emerged slowly over time that fact finders consistently misappraise the implications of certain evidence, and either are more or less persuaded by it than they should be. The inference instructions are given to encourage the fact finder to reach what the system as a whole believes to be correct decisions. Examples are that persons normally intend the natural and probable consequences of their acts, and that possession of recently stolen goods gives rise to an inference of knowledge of their stolen character.

In addition to the direct allocation of burdens of persuasion and instructions on inferences at what might be called the global level, the law makes more refined adjustments to the proof process through the use of presumptions. These usually take the form of changing the standard burden of persuasion, or the giving of a nonstandard inference instruction, only if specific facts are found. For example,

usually the person suing for breach of a contract would have to show acceptance of an offer to enter into a contract. In some states, the mailing of a letter accepting a contract shifts the burden of persuasion to the other side to show that it did not arrive, and thus that no contract came into being. In others, the mailing of the letter can generate an instruction to the effect that there is good reason to believe a valid contract was created, thus commending an inference to the fact finder.

There are an enormous number of adjustments to the fact finding process by presumptions, but they are all driven by the same concern that, left to its own devices, the fact finder is more likely to make what the policy maker believes to be a mistake than if the presumption is employed. (Allen 1980, 1981) Thus, burdens of persuasion, inferences, and presumptions all respond directly to the concern of rational decision making. There are nonetheless difficulties, which I address below.

2.2. RELEVANCE AND MATERIALITY

The second set of limitations on the proof process come in the universal requirements that evidence offered at trial be relevant to prove what it is offered to prove, the relevancy requirement, and that the proposition the evidence is offered to prove be of consequence to the cause of action, the materiality requirement (Federal Rule of Evidence 401). Whether something is relevant is left to the good sense of the trial judge in the first instance, who is empowered to exclude evidence if unpersuaded of its relevancy, and in the second instance by jurors who, once the evidence is admitted, generally are empowered to make of it what they will. Whether a proposition is material is deductively derived from the controlling substantive law. Through these related requirements, the trial judge may exclude evidence that bears no logical relationship to the cause of action and thus the only effects of which can be to waste resources and generate erroneous conclusions. Under modern law these decisions are truly given to the discretion of the trial judge. Some precedents do arise that constrain that discretion somewhat, but not much. Thus, this form of regulation of the proof process depends crucially on the judgment of a human actor. It simply replaces the judgment of the parties with that of the judge.

Judges are also empowered to exclude relevant evidence if, in the words of Federal Rule of Evidence 403, “its probative value is substantially outweighed by the danger of unfair prejudice, confusion of the issues, or misleading the jury, or by considerations of undue delay, waste of time, or needless presentation of cumulative evidence”. This interesting provision combines a preference for admission in the “substantially outweighed” language with another recognition of discretionary power in the trial judge. The trial judge can exclude evidence on a belief that there is a substantial risk of evidence perverting rather than promoting the fact finding process. Again, some minimal case precedent arises over the meaning of the terms of the rule, but the rule for the most part simply allocates judgment over these issues to the judge.

The general principles of relevancy and materiality animate a number of specific provisions of the rules. The general requirement that real (i.e., nontestimonial) evidence be authenticated, which means shown to be what it purports to be, is an example (Federal Rules of Evidence, Article IX). The best evidence rule, which is that the original of a writing must be employed if its contents matter, or its absence excused, is another (Federal Rules of Evidence Article X). Aspects of the hearsay rule are clearly relevancy driven (Federal Rules of Evidence, Article VIII). Statements can be taken out of context, and thus be misleading, a risk that is heightened with hearsay, for example. A few rules specifically admit or exclude evidence because of its likely relevance or prejudicial effect (Federal Rules of Evidence 609, pp. 412–415). In modern evidence codes, however, these are remarkably few indeed.

2.3. POLICY CONSIDERATIONS

In modern evidence codifications, much more space is spent on adjustments to the process of proof for reasons of policy than for reasons of relevancy, and correspondingly less discretion is given to trial judges. Hearsay, which in part is relevancy driven, is also in part driven by policy considerations, in particular in the criminal arena where there is a deeply held belief in the importance of confrontation. A series of rules exclude evidence to advance certain policies, such as excluding recent repair in order to encourage repair, excluding offers in compromise in order to encourage compromise, excluding the prior sexual history of an alleged victim of sexual assault in order to encourage the bringing forth of such complaints, excluding evidence of insurance in order to remove “deep pocket” motivations from jurors, and so on. On occasion the law also provides for the admission of evidence to advance policy goals, the most recent and most lamentable being an amendment to the Federal Rules of Evidence to admit the sexual background of sexual predators apparently regardless of its relevancy or prejudicial effect (Federal Rules of Evidence, pp. 413–415).

A reasonably robust set of privileges exist whose primary purpose is to facilitate and encourage certain relationships. Confidential communications are protected between spouses, doctors and patients, psychotherapists and patients, and clergy and penitent. Various states have numerous other privileges, such as a journalist privilege, trade secrets, and a voting privilege.

3. The difficult relationships between formalizations and the legal process

Various formalizations have been advanced as descriptive of or prescriptive for legal decision making. The implications of two have been particularly well investigated, with somewhat disappointing results: expected utility theory as an explanation of burdens of persuasion and Bayes’ theorem as a theory of proof.

A third formal literature involving micro- economic analysis of the proof process is in its infancy, and I will not discuss it further here (Allen et al. 1990).

3.1. EXPECTED UTILITY THEORY

Burdens of persuasion appear quite readily to lend themselves to decision theoretic explanations. Actual truth is rarely known in any particular case, otherwise there would be little need for a trial, and a large number of cases need decision. It certainly violates no canon of common sense to view this as calling for maximizing expected utility. Moreover, simple calculation generates the expected utility maximizing rules. In a case involving a binary choice where the disutilities of wrongful verdicts are equal, decision should be for whomever the probabilities favor, which is the 0.5 rule of civil law. In cases involving more than two possible explanations, decision should be for the most probable (and here we see the first problem, for this is not the law). If disutilities of wrongful decisions are not equal, as in criminal cases where a wrongful conviction is considerably worse than a wrongful acquittal, the decision rule is adjusted to accommodate the difference. Simple as that (Kaplan and Kaye 1980; Kaye 1982).

The relationship between the world of mathematics and the world of human affairs, however, is not that simple, and there are problems. First, in its present version, expected utility theory does not in fact describe the law, although it may be used to criticize it. The law applies burdens of persuasion to elements, not to causes of action as a whole. Expected utility theory conflates the two as though the question asked at trial were liability or no liability. In a sense it is, but only after burdens of persuasion are applied to individual elements. Applying burdens of persuasion understood as probability measures to elements yield the well know paradoxes of proof. If, for example, two stochastically independent issues are each established to a 0.6 probability, the probability of them both being true is 0.36. Returning a verdict for a plaintiff in such cases is not going to maximize expected utility.

The puzzles press more deeply, however. The expected utility theorist may respond by criticizing the law and arguing that it is the conjunction of elements that should be found to a specific level. This, too, yields unacceptable consequences, by making the level of proof of specific elements turn on the fortuity of the number of elements in a cause of action. My favorite example is to compare theft to murder. Theft has considerably more elements than murder. Thus to convict for theft requires on average that intent to steal be established to a higher probability than intent to kill for a murder conviction. This strikes all legal observers as both unacceptable and absurd. In the concluding section, I advance one possible solution to this problem.

There is a second set of problems with expected utility theory: in its simplest manifestation, it ignores base rates, the accuracy of probability assessments, and that it is not the subjective expectation of judges and jurors that the legal system

wishes to maximize. More sophisticated versions, by contrast, have not given a plausible account of how these matters could in fact be taken into account in such a way as to increase the probability of furthering the objectives of the legal system. Both versions also neglects certain implications of subjective probabilities that I will discuss in the next subsection.

The legal system involves third party decision makers – judges and juries – implementing the wishes and commands of the sovereign people – or less grandly the policy makers, who are typically legislators and sometimes constitution makers. Thus the utility to be maximized is that of the policy makers, not that of the judges and juries, and the two could be widely disparate. For example, with no knowledge of base rates or the relative accuracy of probability assessments, the lessons of expected utility theory are quite straightforward for fact finders. However, the policy maker may think he is in possession of such knowledge, and that knowledge dramatically affects the expected outcome, from the policy maker's point of view. Take a simple example. If no factually liable defendants go to trial, the only kind of error possible is holding a defendant wrongfully liable (a false positive). Lowering the burden of persuasion can only increase the policy maker's expected utility, whatever it does to that of the fact finder. Analogously, policy makers may believe that probability assessments of fact finders are skewed in some fashion, generating the same problem.

No means of accommodating this point has been advanced. The proponents of expected utility theory within the law have simply asserted that beliefs about base rates and the probability assessments can themselves be taken into account in forming subjective probabilities. True enough, but it is difficult to see what programmatic implications this may have, for informing the fact finder of this knowledge would have unpredictable effects on the fact finder's appraisal of the evidence. Any particular fact finder may over- or underestimate the probabilities of liability, and information about the systemic knowledge may lead to widely disparate adjustments to accommodate that knowledge.

The law does not just stand passively by in these circumstances, however. Indeed, one of the points of presumptions and inferences is to adjust the fact finding process in light of such concerns. The adjustments, however, proceed cautiously with a healthy dose of skepticism that we know what we are doing, and a healthy respect for the independence of the fact finder. Of course, if the claim of expected utility theory is simply that it provides another vocabulary for the discussion of such problems, it is in no way objectionable, but neither is it in any way particularly helpful.

3.2. BAYES' THEOREM

Just as burdens of proof seem to lend themselves to decision theoretic analyses, so does the evidentiary process at trial seem to lend itself to a Bayesian interpretation. Indeed, it is quite natural to think of the evidentiary process as the updating of a

prior in light of new evidence. There are a number of difficulties, however, both pragmatic and conceptual.

The pragmatic difficulty is computational complexity. A huge and complicated data set is involved at most trials, even most "simple" trials. No computer, let alone any human, has the computational capacity to do the calculations necessary for the operation of Bayes' Theorem in a reasonable amount of time. The response of the Bayesians to this has been two fold. First, it is not their fault that the world is complicated. True enough, but it is and it is not obvious how Bayes' Theorem helps us navigate that world, given the conditions at trial. Second, they argue that nothing within Bayes' Theorem instructs on what the unit of analysis is. Thus, the fact finder can lump a bunch of stuff together and update his prior using the bunch of stuff as the datum of "new evidence" (Allen 1997). Maybe so, but all the intellectual work will have been done in the "bunching", and thus Bayes' Theorem in this manifestation is not terribly helpful or interesting.

The conceptual difficulty is deeper and even more intractable. It has two components, the first emerging directly from the nature of trials. Reflect upon the description of trials. Fact finders have no good sense of what is going on until the end of the trial at closing arguments. Moreover, they are not bound in any way by those arguments, and thus are free to generate their own theories of what happened. This has two implications. First, once the fact finder hears the various theories in closing, to operate Bayes' Theorem it must then assign the various theories probabilities. But those probabilities will be assigned in light of the evidence heard at trial, and thus all that evidence is what is called "old evidence", which simply means it has already been taken into account. Second, even following the assignment of initial probabilities, new theories can emerge during deliberations. The emergence of a new theory requires, for Bayes' Theorem, a reassignment of initial probabilities of all possible theories, and again the problem of old evidence rears up. The possible scope of Bayesian computations is thus exceedingly limited (Allen 1997).

The second aspect of the conceptual problem emerges directly from the assignment of initial probabilities. Those assignments are subjective and need respect only the conditions of consistency and summing to 1.0. That means that individuals can begin from radically different perspectives, and each, in Bayesian terms, will be operating equally rationally. Bayes' Theorem provides no method of adjudicating such differences. In other contexts, such as science, these differences can be marginalized by convergence theorems that demonstrate that over time and with enough new evidence the divergent initial starting points will washout and the result will converge on the truth. There is nothing even remotely analogous to this in the condition of trials. Jurors are like scientists reflecting on new theories for the first time; they are not like scientists who have generated substantial evidence over time designed to adjudicate between competing scientific theories. Without something to take the place of convergence theorems, the arguments about Bayes'

Theorem in the law are left with no obvious bridge between the subjective and the objective (Allen 1997).

There are additional formal problems. Savage's formalization of subjective probability includes the "sure thing" principle, which is the pivotal axiom that produces the interchangeability of subjective and objective probabilities (Savage, pp. 21–26). In words, this axiom says that if act A is preferred to act B under one set of assumptions about nature, then augmenting the set of assumptions should not cause a reversal of preferences, i.e., for B to be preferred to A. For example, if you are given a menu and you prefer chicken to turkey and then are told that the kitchen also serves duck, it is a violation of the sure thing principle to say, "Given that additional information, I will switch my order to turkey".⁴ Unfortunately, humans disobey this axiom all the time. Suppose you believe that turkey requires great care in preparation, and you ordered chicken because you are risk averse. Duck is very difficult to prepare, however. Having learned that duck is on the menu, you have greater trust in the chef and so switch from chicken to turkey. This violates the sure thing axiom. People regularly disobey this axiom because it requires the articulation of all logical propositions in a probability space, a daunting task even in a quite confined space and an impossible one when the probability space ranges over all human affairs. This is a particularly acute problem for group decision making. As Savage himself said, "It would not be strange, for example, if a banquet committee about to agree to buy chicken should, on being informed that goose is also available, finally compromise on duck" (Savage, p. 207).⁵

3.3. WHAT IS THE ALTERNATIVE?

The primary competitor to the Bayesian view of juridical proof at the moment is that legal fact finding is primarily a matter of the plausibility of the parties' explanations offered at trial, a theory which has come to be known as the relative plausibility theory. In civil cases, the fact finder is to identify the most plausible account of the relevant events, and in criminal cases the prosecution must provide an account of guilt and that there is no plausible account of innocence (Allen 1991, 1994). Although there are many nuances to the relative plausibility theory

⁴ The textual words are largely Albert Madansky's of the University of Chicago, for whose help with some of the more difficult aspects of subjective probability I am indebted.

⁵ There are certain complexities here. Empirical work has demonstrated that individuals violate the sure thing principle (Shafir 1994; Tversky and Shafir 1992). Whether the textual example is a real life example is more difficult, for it might instead be an example of a misspecification of the probability space. The result, from my perspective, is the same, regardless what the example is an example of. If it is a violation of the sure thing principle, then subjective probability axioms are violated. If is an example of a misspecified probability space, it demonstrates how at trial the probability space is constantly corrigible based on new information until the point of decision. That in turn means that there is no work for Bayes' Theorem to do until the point of decision, at which point the probability space is formed; but, at that point there is no work for Bayes' Theorem to do either, for all evidence will be old evidence.

(including that fact finders are not restricted to the stories advanced by parties), in general the structure of liability in this theory is provided by the formal elements, but that is distinguished from the proof process, which proceeds in a largely comparative fashion over the stories advanced by the parties. Once the most plausible account of the relevant events is determined, liability flows deductively from the formal structure of the law. The relative plausibility theory is largely analytic, but it bears a close relationship to the empirical work on jury decision making done by Pennington and Hastie (1991). It also bears a close relationship to the work done on hypothesis comparison through the use of connectionist approaches, such as in the work of Paul Thagard (1992).

Among the advantages of the relative plausibility theory are:

1. It explains what advocates actually do at trial.
2. It avoids the formal difficulties of Bayesianism, as it has so far been developed in the literature.
3. The paradoxes of proof are marginalized because they are distributed evenly over both sides of a dispute.
4. Computational complexity is largely eliminated as a problem because litigation focuses on the plausibility of coherent stories advanced by the parties rather than on discrete items of evidence.
5. The theory bears many interpretations, including a Bayesian one, although its primary ingredient looks more like a likelihood ratio than anything else. As a formal matter, the issue is the relative plausibility of the evidence given the stories rather than the probability of any particular story given the evidence.
6. It explains what fact finders actually do.
7. It explains many aspects of the rules of evidence. (Allen 1991, 1994)

The theory has some disadvantages, also, at least one of which may be amenable to artificial intelligence work. "Plausibility" is large undefined, although Pennington and Hastie have specified what is in fact taken as persuasive by jurors, and Paul Thagard's work shows that hypothesis comparison can occur consistently with Bayes' Theorem but without its attendant combinatorial explosion. The views of the artificial intelligence community could be quite important on these matters. Another disadvantage is that the theory is not formalized, and thus will not appeal to the more formalistic theorists. And finally, little can be said about the theory's implications for errors and utility. This may be a strength, however. For the reasons previously discussed, decision theoretic approaches have limited utility, which at least suggests that attention should focus instead on the fairness of procedures for resolving disputes. The relative plausibility theory captures this point quite nicely, for it essentially guarantees the adversaries the right to present what they each think most persuasive to the fact finder and asks the fact finder to determine which state of affairs is more plausible. This at a minimum guarantees a chance to be heard by a disinterested decision maker during which the party is largely in control of its own fate, which is at the heart of procedural fairness. Moreover, it does not take much of a leap of faith to believe that error minimization may result, even if it

cannot be formally demonstrated, which is surely a more important desideratum than expected utility maximization.

So, what does artificial intelligence have to say about all this? Perhaps quite a lot. In my article in the special issue of the *International Journal of Evidence and Proof*, I commented that the argument I had presented “suggests that the legal debates should move from Bayesianism to plausibility and explanations” (Allen 1997). Perhaps AI Research on plausibility and explanation could be usefully applied to that effect. Many of us are keenly interested to find out if this is so.

References

- Allen, R. (1980). Structuring Jury Decisionmaking in Criminal Cases: A Unified Constitutional Approach to Evidentiary Devices. *Harvard Law Review* 94(2): 321–368.
- Allen, R. (1981). Presumptions in Civil Actions Reconsidered. *Iowa Law Review* 66(4): 843–867.
- Allen, R. (1991). The Nature of Juridical Proof. *Cardozo Law Review* 13(2–3): 373–422.
- Allen, R. (1994). Factual Ambiguity and a Theory of Evidence. *Northwestern University Law Review* 77(2): 604–640.
- Allen, R. (1997). Rationality, Algorithms and Juridical Proof: A Preliminary Inquiry. *International Journal of Evidence and Proof* 1997 (Special Issue): 254–275.
- Allen, R., Grady, M., Polsby, D., and Yashko, M. (1990). A Positive Theory of the Attorney-Client Privilege and the Work Product Doctrine. *Journal of Legal Studies* XIX(2): 359–397.
- Allen, R. and Kuhns, R. (1989). *An Analytical Approach to Evidence: Text Problems, and Cases*. Little and Brown: Boston, MA.
- Allen, R., Kuhns, R., and Swift, E. (1997). *Evidence: Text, Cases and Problems*, 2nd edn. Aspen Publishers: New York, NY.
- Byrne, M. D. (1995). The Convergence of Explanatory Coherence and the Story Model: A Case Study in Juror Decision. In Moore, J. D. and Lehman, J. F. (eds.), *Proceedings of the Seventeenth Annual Conference of the Cognitive Science Society*, 539–543. Erlbaum: Mahwah, NJ.
- Kaplan, J. (1968). Decision Theory and the Factfinding Process. *Stanford Law Review* 20(6): 1065–1090.
- Kaye, D. (1980). Naked Statistical Evidence. *Yale Law Journal* 89(3): 601–611.
- Kaye, D. (1982). The Limits of the Preponderance of the Evidence Standard: Justifiable Naked Statistical Evidence and Multiple Causation. *American Bar Foundation Research Journal* 1982(2): 487–516.
- Pennington, N. and Hastie, R. (1991). A Cognitive Theory of Juror Decision Making: The Story Model. *Cardozo Law Review* 13(2–3): 519–557.
- Rashkin, D. and Yuille, J. (1989). Problems in Evaluating Interviews of Children in Sexual Abuse Cases. In Ceci, S. J., Ross, D. F., and Tolia, M. P. (eds), *Perspectives on Children’s Testimony*, 184–207. Springer-Verlag: New York, NY.
- Savage, L. (1972). *The Foundations of Statistics*. Dover Publications: New York, NY.
- Shafir, Eldar (1994). Uncertainty and the Difficulty of Thinking Through the Disjunction. *Cognition* 50: 403–430.
- Thagard, P. (1992). *Conceptual Revolutions*. Princeton University Press: Princeton, NJ.
- Tversky, A. and Shafir, E. (1992). The Disjunction Effect in Choice Under Uncertainty. *Psychological Science* 3(5): 305–309.