Artificial Intelligence and Law **10:** 163–218, 2002.

An AI model of case-based legal argument from a jurisprudential viewpoint

KEVIN D. ASHLEY*
Professor of Law and Intelligent Systems, Senior Scientist, Learning Research and Development Center, University of Pittsburgh, Pittsburgh, Pennsylvania, USA

Abstract. This article describes recent jurisprudential accounts of analogical legal reasoning and compares them in detail to the computational model of case-based legal argument in CATO. The jurisprudential models provide a theory of relevance based on low-level legal principles generated in a process of case-comparing reflective adjustment. The jurisprudential critique focuses on the problems of assigning weights to competing principles and dealing with erroneously decided precedents. CATO, a computerized instructional environment, employs Artificial Intelligence techniques to teach law students how to make basic legal arguments with cases. The computational model helps students test legal hypotheses against a database of legal cases, draws analogies to problem scenarios from the database, and composes arguments by analogy with a set of argument moves. The CATO model accounts for a number of the important features of the jurisprudential accounts, including implementing a kind of reflective adjustment. It also avoids some of the problems identified in the critique; for instance, it deals with weights in a non-numeric, context-sensitive manner. The article concludes by describing the contributions AI research can make to jurisprudential investigations of complex cognitive phenomena of legal reasoning. For instance, unlike the jurisprudential models, CATO provides a detailed account of how to generate multiple interpretations of a cited case, downplaying or emphasizing the legal significance of distinctions in terms of the purposes of the law as the argument context demands.

1. Introduction

Analogical legal reasoning presents a conundrum to American law school professors. Most professors employ cases in teaching. They assign readings from

* The research described here has been supported by grants from the National Science Foundation, West Publishing Company, Digital Equipment Corporation, Tektronix, the National Center for Automated Information Research, and the University of Pittsburgh ECAC Advanced Instructional Technology Program. As my Graduate Research Assistant and Ph.D. dissertation advisee, Vincent Aleven, designed and built the CATO program as part of his Ph.D. dissertation project in the University of Pittsburgh Graduate Program in Intelligent Systems (ISP) (Aleven 1997). His careful reading of a draft of this article and his thoughtful comments have been invaluable. I am deeply grateful for the opportunity to have taught, advised, and learned from Vincent. Another ISP GRA and Ph.D. dissertation advisee, Stefanie Brüninghaus, developed a web site for the CATO teaching materials. A third ISP GRA, Ravi Desai, is working on modifying CATO to generate more realistic legal dialogues. My Law School Research Assistant, Sean Bamford, helped me greatly in the preparation of citations and quotations for this article. I thank them all.

casebooks designed especially to encourage critical comparisons of cases, and they encourage students in class to make arguments by analogy with the cases published in the casebooks. The judges who wrote the legal opinions featured in the casebooks often appear to have drawn inferences from analogical comparisons of the current problem and past cases. The attorneys who argued the legal points before those judges probably made arguments by analogy to cases. Indeed, the professors of the attorneys, judges and law professors probably all employed a roughly similar case method of teaching.

Despite the pedagogical and professional prominence of analogical legal reasoning, however, jurisprudential enthusiasm for it ebbs and flows. While jurisprudential scholars recognize the popularity of analogical reasoning in the practice of law, and as law professors, they probably assign casebooks and entertain arguments by analogy in class, they do not agree about the epistemological underpinnings or normative value of analogical legal reasoning. They criticize analogical legal reasoning because it lacks certainty of inference, and because models of analogical legal reasoning have not adequately specified a theory of relevance or a theory of error. As explained in Section 2.4, one critic has even called analogical legal reasoning "perverse", a "monster" because no matter how hard scholars try to slay it, it rears its ugly head anew, and a "phantasm" because "it does not really exist, though the [widespread] belief that it does can lead to methodological error" (Alexander 1996, p. 86).

This article describes a computational model of case-based legal argument embodied in CATO, an intelligent tutoring environment intended to teach law students how to make basic legal arguments with cases. The model is sophisticated enough to address a number of the aspects and issues of analogical legal reasoning that the jurisprudential models deem important. It largely accounts for the theory of relevance of the main proponents of analogical legal reasoning while avoiding at least some of the problems urged by critics. It also addresses at least one phenomenon that the jurisprudential models do not seem to address in detail: the strategic reinterpretation of cases in arguments. At the same time, the model has been realized concretely enough to perform some real work in the world. The model runs on a computer, has a database of legal cases from which it draws analogies to problem scenarios, composes arguments by analogy from seven argument moves, and powers a tutoring environment which has been empirically evaluated in a controlled experiment with positive results.

In approximate order of exposition and ambition, this article has four goals.

The first is to describe the phenomenon of analogical legal reasoning to be modeled. In Section 2, I summarize two recent philosophical treatments of analogical legal reasoning by authors who believe in its value (Sunstein 1993; Brewer 1996). Their accounts may represent a recent high water mark in jurisprudential regard for analogical legal reasoning, and have occasioned the impassioned response from Larry Alexander, part of which is quoted above.

The second goal is to describe CATO's model of case-based legal argument and illustrate its use as an instructional environment designed to teach some basic argumentation skills (Section 3). CATO teaches argumentation by example using examples it generates itself. My goal is to illustrate some of these sample arguments and show how they are constructed so that I may then relate CATO's model to the jurisprudential models of analogical legal reasoning. Specifically, I describe (in Section 3.1) how CATO represents four sources of legal knowledge (i.e., factors, cases, relations of factors to legal issues, and argumentation techniques) and illustrate how those knowledge sources interact to test legal hypotheses (Section 3.2) and produce the argumentation examples in CATO's curriculum (Section 3.3).

A computational model helps account for a reasoning phenomenon by demonstrating how a complex interaction among knowledge sources may be orchestrated to achieve an intelligent result. In Section 3.4, I discuss CATO's algorithmic account of this interaction as CATO interprets cases in a context sensitive way, alternatively downplaying and emphasizing the legal significance of a distinction between a problem and a case cited by an opponent.

The third goal is to relate CATO's model to the jurisprudential models. I argue that CATO's model accounts for some of the important features of Brewer's and Sunstein's jurisprudential models of analogical legal reasoning and avoids some of the problems identified in Alexander's critique of the jurisprudential models (Section 4). Although the CATO model may not convince the Alexanders of the world that analogical legal reasoning is a legitimate form of legal inference, it may help to convince the Brewers and Sunsteins that it captures significant aspects of analogical legal reasoning. Comparing the components of CATO's computational model with those of the jurisprudential models also reveals CATO's jurisprudential assumptions and their ramifications for its design, implementation, and instruction. As discussed in Section 4, issues and concerns developed in the two jurisprudential models and their major critiques correspond to actual decisions in the design of CATO and its instructional uses.

The fourth goal is to suggest a role for Artificial Intelligence (AI) research into epistemological phenomena of interest to jurisprudential scholars (Section 5). As a sub field of computer science, AI endeavors to create computer programs whose behavior, if performed by a human being, would be regarded as "intelligent" (Minsky 1968). The virtue of applying AI to research in legal or practical ethical reasoning "is that the nature of the subject forces additional explicitness and clarification, because ultimately its products must be encoded and run on a computer. Thus hidden and unclear assumptions can often be exposed in such a context and applied back in a conscious manner in these and other contexts" (Schaffner 1990).

One important difference between CATO and the jurisprudential models is that CATO generates multiple interpretations of the significance of differences among similar cases. The argument's requirements lead it to make decisions about how to characterize the importance of relevant similarities and differences. It can perform this task for any pair of a problem scenario and case-to-cite drawn from its 150

cases. One can point to exactly those knowledge sources employed and detailed algorithms by which the program chooses a legally plausible argument.

CATO's ability to interpret cases' relevant similarities and differences in terms of more abstract legal issues and its other abilities (e.g., to help students learn to test hypotheses about the importance of factual features in a legal domain and to differentiate among mere differences and true distinctions among cases) relate to epistemological phenomena of interest to jurisprudential scholars. Specifically, they relate to the task of generating a rule, or versions of a rule, with which to epitomize a legal analogy, a cognitive task that plays an important role in Brewer's and Sunstein's jurisprudential models. More generally, they relate to the dialectical relationship between normative principles and concrete cases.

Comparing an AI model of case-based argument and jurisprudential models of analogical legal reasoning is especially appropriate in a special issue dedicated to the memory of Professor Donald Berman. Don was a most enthusiastic supporter of the CATO research, but also a probing critic. In a classic article, his colleague Carole Hafner and he criticized AI models of legal case-based reasoning as lacking a deep model, in particular lacking a teleological component for reasoning with the purposes underlying legal rules (Berman and Hafner 1993). Specifically, they called for case-based reasoning systems that can explain, "why a factor favors the plaintiff or defendant, and why that factor is considered legally relevant" (Berman and Hafner 1993, p. 56). As argued below, CATO satisfies that challenge.

Berman and Hafner were also profoundly aware of the significance of legal principles in interpreting legal rules and cases (Berman and Hafner 1986). In this respect, I am sure Don would find much to agree with in the jurisprudential accounts of Sunstein and Brewer, who, as I explain below, also sought to provide a deeper model of case relevance. Brewer's and Sunstein's accounts root relevance in a cognitive process of expounding a low level legal principle to epitomize the decision and inform subsequent assessments of relevance. Indeed, one of the motivations of this article has been to determine exactly the philosophical standard by which CATO's model would be judged, and to argue that it begins to satisfy that standard.

2. What is analogical legal reasoning?

Before addressing specific jurisprudential models of analogical legal reasoning, it is desirable to define some terminology. "Analogical legal reasoning" involves drawing legal inferences about target problems by comparing them to source cases. Often, the target problems are factual scenarios actually or potentially involving a legal dispute The source cases include past, decided legal cases (i.e., precedents) and hypothetical factual scenarios related to the target problem or precedents in legally significant ways.

"Case-based legal argument" involves a subset of analogical legal reasoning. One draws legal inferences about the target disputes by comparing them to pre-

cedent source cases and constructing legal arguments supporting the inferred conclusions based on the comparison. A model of case-based legal argument attempts to elaborate criteria differentiating reasonable from unreasonable arguments by analogy, and criteria for assessing which reasonable legal arguments by analogy should prevail.

A model of case-based legal argument should be distinguished from a model of legal precedent. The latter delimits the extent to which a legal precedent constrains, or should constrain, judicial decision-making in subsequent cases. Models of analogical legal reasoning and case-based legal argument are related to models of legal precedent in the sense that they must make some assumptions about the extent to which a legal precedent constrains judicial decision-making. These assumptions inform the criteria both of reasonableness and quality of legal arguments by analogy. Conversely, a model of legal precedent may assume a model of analogical reasoning, but it does not have to.

CATO is not a model of legal precedent. It does not purport to delimit the extent to which a legal precedent constrains, or should constrain, judicial decision-making in subsequent cases. Its criteria for assessing the reasonableness and quality of its case-based arguments, however, do embody some assumptions about the extent to which a legal precedent constrains judicial decision-making.

Although CATO models case-based legal argument, it does not model the whole range of case-based legal arguments. All of the analogies CATO draws are among cases within the framework of a single type of legal claim. While the bridgework connecting CATO's analogies is fairly direct, it can draw more abstract characterizations of the analogies, which may offer strategically useful reinterpretations of cases in terms of the law's purposes.

CATO is a *computational* model of case-based legal argument, as opposed to a jurisprudential one. The distinction between a jurisprudential and a computational model underscores the different approaches and domains of discourse of two scholarly communities who attempt to formalize models of legal analogical reasoning. A jurisprudential model is a philosophical model of the underlying phenomenon, one formalized in philosophical prose descriptions, sometimes augmented by logical or other formalisms. A computational model is also written down, but written in the formalism of a computer language such that it can be implemented and run on a computer. One may, of course, describe a computational model at multiple levels of abstraction. Conceptually, the code version that runs on a computer may not always be the most interesting or accessible version. By developing the code version, however, one identifies and addresses lacunae in ones understanding of a phenomenon.

Presumably, there is no reason why a computational model could not also be a jurisprudential model. Indeed, one of the arguments of this paper is that a computational model can be an effective way of capturing certain phenomena of jurisprudential interest.

2.1. BREWER'S MODEL OF ANALOGICAL LEGAL REASONING

As noted, Cass Sunstein and Scott Brewer recently have defended jurisprudential models of analogical legal reasoning from attacks that they lack certainty of inference and a well-articulated theory of relevance. Brewer points to the rule underlying an analogy as the source of its rational force. Both authors assert the need to discover, formulate and test (i.e., abduce) such a rule as an integral step. Both emphasize the role of comparing cases to formulate normative principles in the abduction step. In order to streamline the exposition, I will first summarize Brewer's chronologically later but more detailed account.

Brewer treats argument by analogy in law as one of a family of "exemplary reasoning" methods. His exemplary reasoning family also includes argument by counterexample, reasoning by example as described by Levi, and arguments that proceed "by effecting a 'reflective equilibrium' between general norms and particular applications of those norms" (Brewer 1996, p. 927).[1] For Brewer, "argument by analogy", "exemplary argument", "analogical reasoning", "reasoning by example" and "exemplary reasoning" are equivalent terms. They all share "a logical structure whose defining feature is the focus and reliance on examples in the process of inferring conclusions from premises" (Brewer 1996, pp. 927, 941, n. 52).

In reasoning by example, according to Brewer, a deductively applicable rule informs the example and supplies the analogical argument's "rational force". A legal arguer presents an analogical example (i.e., a source case) to support a conclusion that a current problem scenario (the target case) has a particular legal property. For Brewer, the source case may not serve as an analogical example, however, unless it is presented along with a rule to specify "in what its exemplariness consists". If the legal argument is to be compelling, such an "analogy-warranting rule" or AWR should have a deductive logical structure and satisfy the "entailment requirement". That is, the AWR must serve as a premise, which when applied to the target case (or the source case) deductively entails the conclusion that the target case (source case) has the desired characteristic (Brewer 1996, pp. 971, 975).

As the warrant supporting the analogical argument, the AWR supplies the analogical argument's rational force. The AWR picks out the particular, shared features of the source and target, which justify the conclusion that the target has the same property as the source. In addition, according to Brewer, a compelling analogy also requires an analogy-warranting rationale (AWRa). The rationale explains why, in the "eyes of the law", "the logical relation among the characteristics articulated by the analogy-warranting rule either does obtain or should obtain" (Brewer 1996, p. 965).

Brewer's example of the form of argument by analogy focuses on its use in justification. As shown in Figure 1, he schematizes a judge's argument that a steamboat owner is strictly liable to a passenger for a loss occasioned by the theft of valuables from the passenger's rented steamboat cabin (Brewer 1996, p. 1005). At the heart of the analogy, the analogy-warranting rule states conditions under which an owner

Target (y) = the steamboat owner. [Is a steamboat owner strictly liable to a passenger for a loss occasioned by the theft of valuables from the passenger's rented steamboat cabin?]
Source (x) = the innkeeper. [Cases hold innkeeper was strictly liable for the theft of a boarder's valuables from the boarder's room at the inn.]
Shared characteristic:
F: [Owner] has a client who procures a room for specified reasons R (privacy, etc.)
G: [Owner] has tempting opportunity for fraud and plunder of client.
Inferred characteristic:
H: [Owner] is strictly liable.
Argument:
(1) y has F and G (target premise);
(2) x has F and G (source premise);
(3) x also has H (source premise);
(4) AWR: if anything that has F and G also has H, then everything that has F and G also has H;
(5) Therefore, y has H.

Figure 1. Brewer's example of analogical legal reasoning.

is strictly liable. The AWR subsumes the source case of an innkeeper and the target case of the steamboat owner.

Although Brewer does not identify the analogy-warranting rationale that accompanies the rule in this example, he quotes a passage from the opinion, which appears to provide a suitable AWRa: "The principle upon which innkeepers are charged ... as insurers ... [is that they] should be subjected to a high degree of responsibility in cases where an extraordinary confidence is necessarily reposed in them, and where great temptation to fraud and danger exists by reason of the peculiar relations of the parties" (Brewer 1996, p. 1004).

Brewer also offers an account of "argument by disanalogy" in which "one argues that because the two or more items – which seem prima facie to be relevantly similar – do *not* share some characteristics, we may not infer that they share some additional characteristic that one of them (the source) is known to have". Brewer illustrates two types of disanalogy. The first involves rewriting the AWR to "impose additional conditions on the rules stated (or implied) in prior cases" (Brewer 1996, pp. 1006, 1011). For instance, where appropriate, a distinguisher might rewrite the proposed analogy-warranting rule in such a way that it excludes anything that has F and G but not E. In other words, starting with the following AWR,

> AWR: if anything that has F and G also has H, then everything that has F and G also has H,

the resulting disanalogy-warranting rule would be:

> DWR: if anything that has F and G also has H, then everything that has F and G also has H unless it also has not E.

The DWR might also take another form:

> DWR: if anything that has E, F and G also has H, then everything that has E, F and G also has H.

For instance, one might suppose a distinguisher, arguing on behalf of the steamship owner, happened to observe that in the cases where innkeepers were held strictly liable, they had not posted warning notices to customers to protect their valuables. Knowing that the steamship owner had posted such notices, the distinguisher might offer the above DWR where E is:

E: Owner failed to post a warning notice.[2]

According to Brewer, the distinguisher should also provide a disanalogy-warranting rationale (DWRa), perhaps to the effect that innkeepers who fail to post warning notices assume the risk that uninformed, unwary customers will leave valuables in the room for the convenience of thieves.

Brewer's second but related kind of disanalogy involves distinguishing a target problem from a competing line of source cases, which had the opposite result. For instance, the steamship owner might attempt to draw an analogy to a line of cases which held "that the owner of a railroad was not strictly liable to railroad passengers who had personal goods stolen from the open-berth sleeping cars on trains" (Brewer 1996, p. 1013). Given the above analogy-warranting rule, a distinguisher, arguing against the steamship owner, could maintain that the situation of the railroad is not analogous to that of the steamship. Arguably G does not apply. Unlike a steamship berth, an open booth in a train does not afford the privacy of a room and thus does not present the railroad owner with the same tempting opportunity for fraud and plunder of the client. It follows that the line of railroad cases "does not satisfy the sufficient conditions for the inferred characteristic" H, strict liability.

Finally, Brewer offers an important reason, based on the concept of defeasibility, why a reasoner should remember the facts of source examples even after an appropriate AWR has been adduced to explain the analogy. "A defeasible argument is one in which the addition of premises can weaken the force of the conclusion". A judge knows "that later judges may well come along and rewrite the AWR". Therefore, "in a context of doubt, the legal reasoner uses the resources of analogy both to build and to maintain confidence in her judgment about how that doubt is to be resolved" (Brewer 1996, pp. 1017, 1020). "[T]he reasoner keeps her eye on the shared characteristics of source and target and thus does not simply dispense with the example, because she is confident that source and target are alike in the respects specified by the AWR, that those respects are relevant to being "defeased" or not, that the source case managed to defeat defeasibility, and that therefore one ought to adjudge defeasibility as being likewise defeated in the target case as well" (Brewer 1996, p. 1020).

Brewer appears to be saying: a source case, which might have given rise to an exception to a rule, but did not, lends support for the conclusion that no exception is warranted in the target case either. Consider a source case with characteristics E, F and G where the court held that H applied. The target case also has E, F and G. Assume that the source case's enunciated AWR is: "if anything that has F and G also has H, then everything that has F and G also has H". Assume also that had

the result been the opposite, the source case might have given rise to an exception: "if anything that has F and G also has H, then everything that has F and G also has H *unless it also has E*". The fact that the source case did *not* have the opposite result suggests that the opposite result is not warranted in the target either. Thus, the source case provides support for an assertion that E does not overpower or outweigh F and G. That information would be lost, however, if the source case's facts were not retained along with its enunciated AWR. (See Section 4.1.2).

2.2. SUNSTEIN'S MODEL OF ANALOGICAL LEGAL REASONING

Sunstein's account of analogical legal reasoning is more general than Brewer's, and focuses more on the intellectual work analogical reasoning can perform as an instrument for analytical investigation of a legal issue. Sunstein eschews defining any formula for a compelling legal argument by analogy, opting instead for a more general description. "Without relying on general theories, and without achieving reflective equilibrium, lawyers develop low-level principles to account for particular judgments, and apply those low-level principles to new cases in which there is as yet no judgment at all" (Sunstein, 1993, pp. 758f.). His account focuses less on analogical reasoning's justificatory role than on its role in facilitating legal inquiry to help formulate and test reasonable legal positions. His extended example, excerpted below, illustrates how posing analogies frames and assesses a legal analysis. The analogies suggest reasonable propositions, test them, suggest additional inquiries, and sometimes lead to plausible, tentative conclusions (Sunstein 1993, pp. 759–767). Specifically, it "outline[s] a set of responses to the cross-burning issue [based on R.A.V. v. City of St. Paul, 112 S. Ct. 2538 (1992)], attempting to show how much progress might be made by working from analogies and low-level principles".

The starting proposition in Sunstein's extended example is that cross-burning is not covered by the First Amendment.

Proposition (I): "Cross-burning is action, not speech, and is therefore outside of the First Amendment altogether. .... To claim constitutional protection, a person must be saying or writing words" (Sunstein 1993, pp. 759f).

← *To disconfirm (I)*, *analogize* to flag-burning cases. United States v. Eichman, 496 U.S. 310, 315 (1990); Texas v. Johnson, 491 U.S. 397, 404–06 (1989) (Rejecting claim that flag-burning as mode of expression is not fully protected under First Amendment.)

This leads to a modified proposition:

Proposition (II): "Content-neutral restrictions on acts that qualify as speech are generally permissible", such as ordinary criminal trespass law (Sunstein 1993, p. 760).

→ *To confirm (II) analogize* to trespass cases. Hudgens v. NLRB, 424 U.S. 507, 520-21 (1976) (Finding no First Amendment right to advertise on shopping mall property strike against store.); Lloyd Corp. v. Tanner, 407 U.S. 551, 567–70 (1972) (Finding no First Amendment right to distribute on shopping mall property handbills concerning Vietnam).

The next step is to pose a hypothetical rule and test its viability with an analogy.

Proposition (III): "Acts that qualify as speech can be regulated if they produce anger or resentment", as in *hypothetical law (A)*: It is a crime to "place on public or private property a symbol, including but not limited to a burning cross or a Nazi swastika, which one knows or has reason to know arouses anger or resentment in others on the basis of race, color, or creed" (Sunstein 1993, p. 760).

← *To disconfirm (III) analogize* to flag-burning cases and *See* Terminiello v. Chicago, 337 U.S. 1, 4–5 (1949) (Conviction of petitioner under ordinance construed as permitting conviction if his speech stirred people to anger, invited public dispute, or brought about a condition of unrest held in violation of First Amendment.) Leads to a modified proposition ....

Beside illustrating the investigative uses of posing analogies, Sunstein's extended example illustrates his four attributes of analogical reasoning: "principled consistency; a focus on particulars; incompletely theorized judgments; and principles operating at a low or intermediate level of abstraction". He says, "judgments about specific cases must be made consistent with one another .... [I]n producing the necessary consistency, some principle, harmonizing seemingly disparate outcomes, will be invoked to explain the cases". Sunstein characterizes analogical reasoning as "bottom-up" in its focus on particulars. "We cannot say whether decided case X has anything to do with undecided case Y unless we are able to abstract, a bit, from the facts and holding of case X. The key point is that analogical reasoning involves a process in which principles are developed with constant reference to particular cases". Analogical reasoning, Sunstein maintains, "operates without a comprehensive theory that accounts for the particular outcomes it yields. Finally, the principles operate "at a low or intermediate level of abstraction. ... Analogical reasoning usually operates without express reliance on any general principles about the right or the good" (Sunstein 1993, pp. 746ff).

2.3. PARALLELS IN SUNSTEIN'S AND BREWER'S ACCOUNTS OF RELEVANCE

Both authors recognize the need for a set of criteria by which analogical reasoning can assess relevant similarities and differences. Sunstein regards this as the most powerful criticism of analogical reasoning: that it assumes a theory of relevant similarities and differences which it does not provide.

> The method of analogy is based on the question: Is case A relevantly similar to case B, or not? ... To answer such questions, one needs a theory of relevant similarities and differences. By itself, analogical reasoning supplies no such theory. It is thus dependent on an apparatus that it is unable to produce. (Sunstein 1993, p. 774; see also Brewer 1996, pp. 932f)

For each author, the relevance criteria are supplied by the low-level principle (Sunstein), which informs the analogy, or by the analogy-warranting rule (Brewer), which subsumes the source and target cases and which lends justificatory weight to the analogies. In Sunstein's view the process of reasoning by analogy generates principles, which, in turn, define relevance (Sunstein 1993, p. 778). Brewer relates the criteria of relevant similarity and difference to the analogy-warranting and disanalogy-warranting rules and rationales, which supply analogies' rational force. Basically, "the AWR supplies the criteria of 'relevant similarity' " (Brewer 1996, p. 1020). The AWR's antecedent picks out the particular shared features of the source and target which justify the conclusion that the target has the same property as the source. The DWR picks out the additional, unshared features in the seemingly similar source and target, which justify the conclusion that the source's property does not apply to the target (Brewer 1996, pp. 965, 1010f).

Although Brewer's theory of relevance focuses on analogy-warranting rules and Sunstein's on principles, the difference probably is not significant. Brewer sees a connection between AWRs and "norms or principles". His account of analogical reasoning draws on a similar process of analogical reasoning in ethics, in which "relatively precise norms or principles" play the role of AWR (Brewer 1996, p. 979). Furthermore, Sunstein makes clear that the principles produced by analogical reasoning "operate at a low or intermediate level of abstraction" (Sunstein 1993, p. 747).

Both authors emphasize the importance of distinguishing and disanalogy and its relationship to AWRs or principles. Brewer explicitly relates relevant differences to AWRs, characterizing two types of arguments by disanalogy (Brewer 1996, pp. 1011ff). While Sunstein is less explicit, he emphasizes that "(t)he major challenge facing analogical reasoners is to decide when differences are relevant" and that "analogical reasoning goes wrong when there is an inadequate inquiry into the matter of relevant differences and governing principles" (Sunstein 1993, p. 745).

Having espoused relevance theories based on principles, each author anticipates a "reductionist" criticism: if legal analogy requires underlying principles or rules to tell what is relevantly similar or different, why not dispense with analogies altogether? Why not simply reason with principles, as one does in reflective equilibrium? (Sunstein 1993, p. 774). Both defend against the reductionist critique by underscoring the cognitive role of comparing cases in helping reasoners to abduce and to evaluate the rules or principles. Both maintain that analogical reasoning requires and facilitates discovery and confirmation of the principles that determine relevance, a process in which comparing the target case and the source cases plays a vital role. As Sunstein says,

> Of course one needs criteria to engage in [analogical] reasoning. But those criteria will emerge from the process of comparing various cases; often they are not given or even describable in advance, except at an unhelpful level of generality . . . (Sunstein 1993, pp. 778f).

Similarly, according to Brewer's model, the AWR is not a given. It must be abduced in a process of inference that requires one to compare source and selected target cases. The AWR emerges in the first two of his three-step process of reasoning by analogy:

> 1. " 'discovery' (or abduction) . . . [of] a proposed AWR using . . . heuristically well chosen source cases (such as actual precedents or hypotheticals.)"
> 2. "confirmation or disconfirmation of the proposed AWRs that have been abduced ('discovered')" in a process of "reflective adjustment".
> 3. application of the "confirmed AWR to the case under consideration" (Brewer 1996, pp. 978, 1022f).

Both authors posit a process of confirmation in which one determines whether the proposed analogy-warranting rules or principles apply to the source case with a result that is both acceptable and coherent in light of comparison to other cases and rules (Brewer 1996, p. 1021).

> In legal argument, the demonstration that the sorting effected by the AWR is acceptable is a staple of briefs and oral arguments in which the advocate shows the court that the rule he advocates – which surely will entail the result the lawyer wishes for his client – manages to draw an acceptable line or is needed to stop a slide along a slippery slope (Brewer 1996, p. 1021, n. 281).

Brewer calls the process of achieving confirmation "reflective adjustment" of an AWR and its proposed application as well as of the accompanying AWRa.

> [A]n abduced AWR might be rejected because . . . it does not, as applied to some particular cases, cohere sufficiently with explanatory and justificatory rationales that the reasoner is unwilling to amend. Or the AWR might be so compelling that the reasoner chooses to hold onto the AWR and effect a modification of the rationales (Brewer 1996, p. 1023).

For Sunstein also, confirmation is integrally related to analogical reasoning as a judge attempts to achieve horizontal coherence among principles and decisions. Sunstein allows that analogical reasoning may then ascend into reflective equilibrium as a judge attempts to achieve more vertical coherence. For a variety of reasons, however, analogical reasoning need not ascend, thus achieving some pragmatic benefits. For example, "reasoning by analogy may be the best approach available for people of limited time and capacities" (Sunstein 1993, p. 782).

The authors disagree as to the dividing line between analogical reasoning and reflective equilibrium and how much "reflective equilibrium" the confirmation step requires (Brewer suggests that Sunstein does not regard the confirmation

step as internal to the analogy (Brewer 1996, p. 1028). They also disagree as to the direction of analogical reasoning. Sunstein calls it bottom-up (Sunstein 1993, pp. 746f), but Brewer says it is can be both bottom-up and top-down (Brewer 1996, p. 982, n. 180). The authors also focus on different models of reflective equilibrium. Brewer's sources are Nelson Goodman and John Rawls, whose models have less to do with reconciling current decisions with judgments about past particular situations (Brewer 1996, p. 938). By contrast, Sunstein's main source (beside Rawls) is Ronald Dworkin, who he says, "describes the search for reflective equilibrium in a way that relies heavily on particular judgments about particular situations" (Sunstein 1993, p. 751, n. 38).

2.4. A CRITIQUE OF MODELS OF ANALOGICAL REASONING IN LAW

A proponent of a rule model of precedent has attacked Brewer's and Sunstein's models of analogical legal reasoning, along with Dworkin's model of reflective equilibrium, on the grounds that they are "not a satisfactory methodology for [explaining how courts follow precedents], either descriptively or normatively" (Alexander 1996, p. 58).

In a rule model of precedent, a precedent controls only those cases logically subsumed by the precedent's rule (Alexander 1989, pp. 17ff). According to Alexander, to be a rule model: (1) "the rule must have a canonical formulation, even if that canonical formulation does not appear in the original opinion, such as, 'Whenever facts A, B, and C, and not fact D, decide for P". (2) "Only the rule and not the [precedent court's] reasoning . . . binds the constrained court". (3) "The formulation of the rule must be fixed at the time of the precedent decision; . . . it must not be dependent on what any court other than the precedent court did". As such, "the constrained court faces a binary choice: it can either follow the precedent rule in its canonical form or overrule it" (Alexander 1989, p. 19).

Given the importance of analogy-warranting rules in Brewer's model of analogical legal reasoning, one might think that Brewer's is also a rule model of precedent. Alexander rejects that characterization, however, because of Brewer's insistence on the need to abduce and confirm an AWR (Alexander 1996, p. 64).

Alexander's main criticism of models of analogical legal reasoning focuses on the confirmational part of the process, which, as described above, involves or, at least, connects to a process of reflective equilibrium. He maintains that Brewers' and Sunstein's theories stumble fatally on the problem of "bad beginnings". They fail to deal adequately with the problem of erroneously-decided precedents or source cases. It is a fact, he maintains, "[j]udges may decide cases in ways that are morally unjustifiable" (Alexander 1996, p. 80). Because of these mistaken decisions, "[l]egal principles for Dworkin [and by extension, for Brewer and Sunstein] can be characterized counterfactually as those principles that would be correct moral principles in a world in which the morally incorrect decisions were morally correct" (Alexander 1996, p. 84). Alexander insists that approach

cannot work. "Correct moral principles will never dictate their own abandonment .... [They] will reject all other purported justificatory principles as counterfeit" (Alexander 1996, p. 85).

Alexander's challenge raises the question of whether a model of legal analogical reasoning which includes a theory of relevance based on principles, also needs to elaborate a theory of error to deal with mistaken judicial decisions. As far as I can tell, Brewer does not deal with this issue. Sunstein recognizes the problem of judicial errors, but does not provide a theory for identifying such errors. Instead, he regards analogical reasoning as an institutionally realistic and justifiable means for dealing with the inevitable errors (Sunstein 1993, p. 778).

More generally, Alexander's critique can be seen as challenging any model of analogical legal reasoning to explain the criteria of "coherence" by which a decision "fits" with past cases and general principles, as required by the abduction/confirmation step and its search for horizontal and some measure of vertical coherence. Although neither Brewer nor Sunstein emphasize the concept of "fit" as much as Dworkin, the notion of coherence does seem to play an important role in their accounts of analogical legal reasoning, but is not adequately explained (Sunstein 1993, p. 784; Brewer 1996, pp. 1022ff).

Critics like Alexander assail Brewer's and Sunstein's claims that analogical legal reasoning has any normative justification at all. Nevertheless, Alexander believes that moral reasoning and the interpretation of rules "have room for analogical reasoning of the type Brewer endorses[.]" (Alexander 1996, p. 72). Apparently, a decision maker may take account of past court decisions as facts; but not as normative "anchors" (Alexander 1996, p. 76). If this is true, then Alexander might still agree with the desirability of teaching law students skills of comparing cases on their facts, even though he does not regard analogical legal reasoning as a legitimate form of legal inference.

3. CATO: An AI model of case-based legal argument

In designing CATO, we have assumed that teaching analogical legal reasoning skills is an important and desirable goal. The design question we faced was how to help students become exemplary reasoners in both senses.

CATO is a computer program that implements a model of case-based legal argument in the context of an intelligent tutoring environment. The goal in designing CATO has not been to teach a particular model of legal argument by analogy but to use a particular model to assist students in learning some important skills. CATO models a relatively small subset of the features of analogical legal reasoning explained in the jurisprudential models, but for those features it models, its model is quite detailed.

In the configuration in which students encounter it, CATO comprises an on-line Casebook Chapter (Ashley and Aleven 1996, Parts I, II) introducing an area of law, an on-line Workbook (Aleven 1996), and a set of six computational tools in

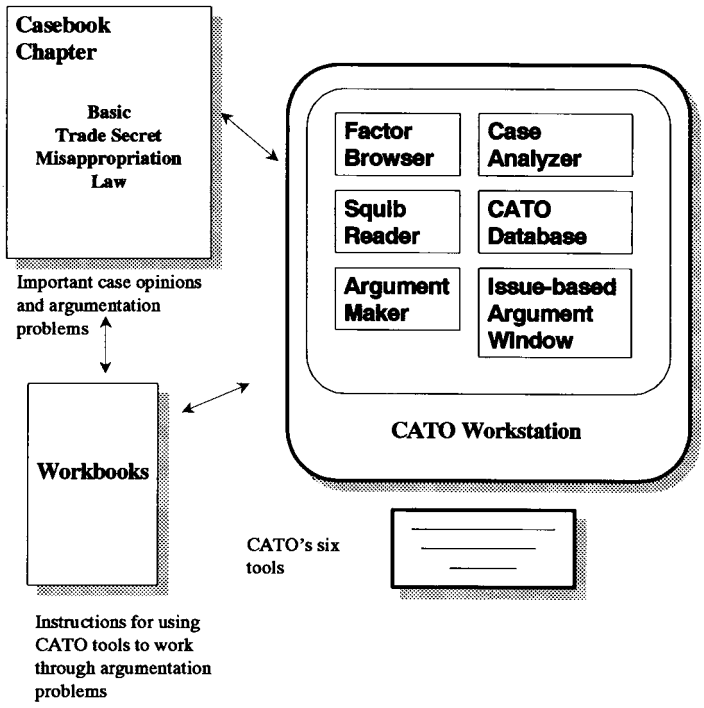

Figure 2. CATO environment and tools.

windows through which students interact with CATO (see Figure 2). By responding to students' interactions with the tools, CATO's intelligent instructional software assists them in learning the skills.

Although the Casebook Chapter deals with the law of trade secret misappropriation, the primary goal is to teach students basic skills of reasoning and arguing with cases. As in traditional casebooks, the Chapter contains the excerpted opinions of a handful of important cases discussed or noted in Goldstein (1990), each of which is followed by a small set of discussion and argumentation questions.

The Workbook lays out a curriculum for teaching the skills through a combination of chapter readings and end-of-case exercises using CATO's tools. It explains how students can use CATO's six computational tools to research and respond to the questions at the end of each case.

After a number of lessons introducing CATO's windows and teaching students to identify factual strengths and weaknesses, the Workbook treats two major topics: (1) Testing hypotheses about a legal domain against cases (discussed below in Section 3.2) and (2) Using cases in arguments (discussed below in Sections 3.3 and 3.4.) The Workbook devotes the most attention to the second topic. It addresses the general topics of analogizing and distinguishing, selecting the best cases to cite in an argument, organizing an argument around issues, and responding to an argument.

In working through this curriculum, CATO employs many examples of making case-based arguments, examples which the program generates itself at the request of the authors or students. The examples selected by the authors appear in the Casebook and Workbook. The Workbook sometimes guides students in using CATO's tools to produce appropriate additional examples.

In constructing examples, CATO employs its computational model of case-based legal argument. The model enables CATO to perform seven basic argument moves with the cases in its database:

1. **Analogizing** a target problem to a relevant source case with a favorable outcome (i.e., a favorable source case).
2. **Distinguishing** a target problem from a relevant source case with an unfavorable outcome (i.e., an unfavorable source case).
3. **Downplaying** the significance of a distinction between a target problem and a favorable source case.
4. **Emphasizing** the significance of a distinction between a target problem and an unfavorable source case.
5. **Conflict-Resolution** by citing a favorable source case to show that the target problem's factual strengths overcome its weaknesses.
6. Citing a favorable source case to argue that a target problem's factual weaknesses are **not fatal**.
7. Citing a **counterexample** to an unfavorable source case cited by an opponent.

By combining these moves, CATO builds and organizes more complex multi-case, multi-issue arguments (Aleven 1997, p. 19).[3] It composed all of the examples in the Workbook and this paper from these moves.

Before examining how CATO generates argumentation examples and employs them in its curriculum, I will first discuss CATO's techniques for representing source cases and target problems.

3.1. KNOWLEDGE REPRESENTATION IN CATO

CATO employs interrelated techniques for representing four sources of legal knowledge:[4] a set of factors, a database of indexed cases, a Factor Hierarchy, and a set of argumentation techniques (i.e., CATO's argumentation templates, recipes, and algorithms).[5]

Each case in CATO's database is represented by (1) a brief textual description of the facts and decision (i.e., the case squib), and (2) a set of applicable "factors" which indexes the case and represents its important factual features. An example of a squib, summarizing a real trade secret dispute between a bar owner and the Jack Daniel Distillery, is shown in Figure 3. As described below, the case of *Mason v. Jack Daniel Distillery* became a centerpiece of the curriculum. Students first encounter it and other squibs in the Squib Reader window. The format of each case squib is similar to that taught to first year students in "briefing" cases: title of the case, citation information, date, parties, statement of the facts, issues and holding.

Unlike a law student's brief of an opinion, a squib does not contain the Court's rationale.

Factors represent stereotypical collections of facts, which tend to strengthen or weaken a plaintiff's legal claim.[6] CATO's 27 factors all deal with the legal claim of trade secret misappropriation. Collectively, they represent the stereotypical patterns of facts, which tend to strengthen or weaken a plaintiff's claim for misappropriation of trade secrets. Each factor indicates the side, either plaintiff (p) or defendant (d), which it typically supports. CATO's factors cover such conditions, generally favoring the plaintiff's side, as plaintiff's having adopted security measures (F6).

The *Mason* case in Figure 3 can be represented as a set of factors (Aleven 1997, p. 20). Factor F1, Disclosure-In-Negotiations captures the plaintiff's weakness that the plaintiff bar owner disclosed the allegedly secret information himself to the distillery's agent in an attempt to secure a deal. In addition, there was expert testimony, that the recipe could easily be duplicated (Factor F16, Info-Reverse-Engineerable). On the other hand, plaintiff has some strengths. The bar owner did take some steps to protect the security of the information (Factor F6, Security-Measures), for instance, he cautioned his bartenders. The distillery's agent apparently was aware that the bar owner regarded the information as confidential (Factor F21, Knew-Info-Confidential), and the bar owner's product was both popular and unique among his competitor's products (Factor F15, Unique-Product).

Although students can read case squibs, CATO cannot. For CATO to process a target problem and source cases, the problem and cases must be represented as a set of the applicable factors. A user or case enterer determines which factors apply after reading the text of the opinion or squib. Users entering source cases or target problems can inspect CATO's factors in the Factor Browser tool and select those that apply. CATO has a database of 147 trade secret misappropriation cases drawn from reported legal cases in a variety of jurisdictions. All are represented and indexed manually in terms of the applicable factors.

In general, the presence of a factor in CATO's representation of a case indicates that the case's opinion contains statements from which it may be inferred directly that the stereotypical collection of facts associated with the factor were present in the case's facts. The absence of a factor from the case's representation indicates that there are no such statements in the case's opinion, or that there are statements from which it may be inferred directly that the collection of facts were not present.

CATO's Factor Hierarchy represents the underlying meaning of factors in terms of the purposes of trade secret law, for example, to protect confidential relationships, discourage unfair competitive practices, or reward inventiveness. By representing the connections between factors and those legal issues to which they are relevant (Aleven 1997, pp. 44–49), it enables the program to explain why a factor's stereotypical collection of facts makes a difference to the legal claim. More specifically, each factor is related to one or more of five major legal issues, whether: plaintiff's information is a trade secret (F101), defendant acquired plaintiff's information through improper means (F110), defendant used plaintiff's information

Title: Mason v. Jack Daniel Distillery
Cite: 518 So.2d 130 (Ala.Civ.App. 1987)
Date: Aug. 5, 1987.
Parties: Plaintiff: Mason; Defendant: Distillery.
Claim: Misappropriation of trade secret
Procedural setting: Trial court denied Distillery motion for directed verdict and jury awarded nominal damages to Mason. Mason appeals court decision barring punitive damages. Distillery cross-appeals denial of directed verdict.
Decision: For plaintiff; remanded.
Facts: In 1980, a restaurant owner named Mason developed a combination of Jack Daniel's whiskey, Triple Sec, sweet and sour mix, and 7-Up to ease a sore throat. He promoted the drink, dubbed "Lynchburg Lemonade" for his restaurant, "Tony Mason's, Huntsville", served it in Mason jars and sold t-shirts. Mason told the recipe only to his bartenders and instructed them not to reveal the recipe to others. The drink was only mixed out of customers' view. Despite its extreme popularity (the drink comprised about one third of the sales of alcoholic drinks), no other establishment had duplicated the drink, but experts claimed it could easily be duplicated. In 1982, Randle, a sales representative of the distillery, visited Mason's restaurant and drank Lynchburg Lemonade. Mason disclosed part of the recipe to Randle in exchange, Mason claimed, for a promise that Mason and his band would be used in a sales promotion. Randle recalled having been under the impression that Mason's recipe was a "secret formula". Randle informed his superior of the recipe and the drink's popularity. A year later, the Distillery began using the recipe to promote the drink in a national sales campaign. Mason did not participate in the promotion or receive other compensation.
Issues: (1) Was there sufficient evidence to allow a jury to determine that the recipe for Lynchburg Lemonade was a trade secret? (2) Was there malice, willfulness, or wanton and reckless disregard of the rights of Mason allowing recovery of punitive damages?
Holding: (1) Yes. For plaintiff. (2) Possibly; remanded for determination by jury.

Figure 3. Case squib of Mason vs. Jack Daniel distillery.

and usurped a competitive advantage (F112), a confidential relationship existed between plaintiff and defendant (F114), and defendant has ownership rights in the information (F124) (Aleven 1997, p. 239, Appendix 1).

As shown in the excerpts in Figures 4 and 5, the Factor Hierarchy is a collection of graphs. The root nodes (at the top of each graph) represent the main legal issues in a claim of trade secret misappropriation (Aleven 1997, p. 239, Appendix 1). The leaf nodes (at the bottom of the graphs) are the factors, each representing a stereotypical collection of factual strengths or weaknesses of legal disputes involving a trade secret claim. Linking the leaf nodes to the root nodes are layers of intermediate nodes.

As one moves upward from a leaf node, each intermediate node represents an increasingly abstract characterization of the legal significance of the factors below that node, leading ultimately to the legal issue nodes to which those factors are relevant (Aleven 1997, p. 44–49). For purposes of discussion, the issues will also be referred to as abstract factors. The Factor Hierarchy contains 26 leaf node factors and 16 abstract factors, 5 of which are top-level legal issues. (For ease of reference, the leaf nodes (i.e., factors) have numbers below 30; abstract factors have numbers above 100.)

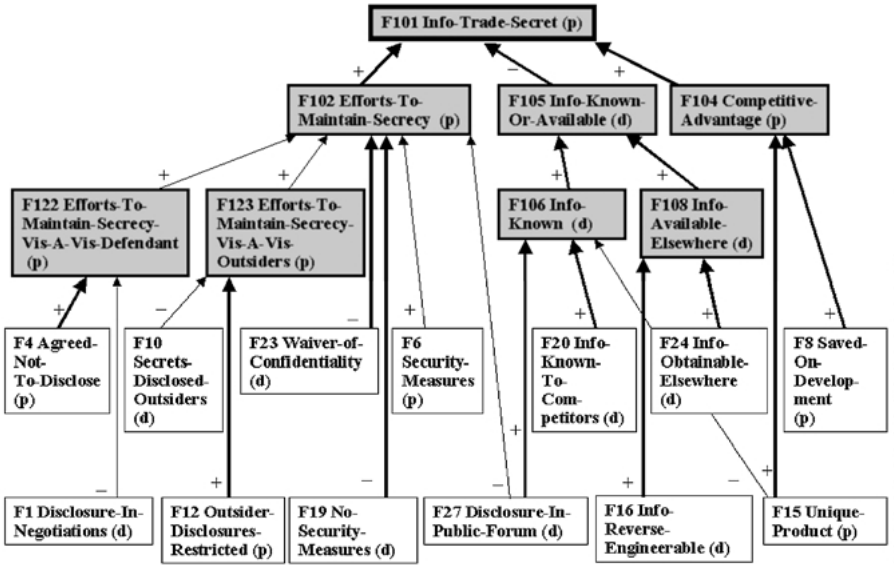

Figure 4. CATO's Factor Hierarchy for Trade Secret Misappropriation Issue: Whether plaintiff's information was a trade secret.

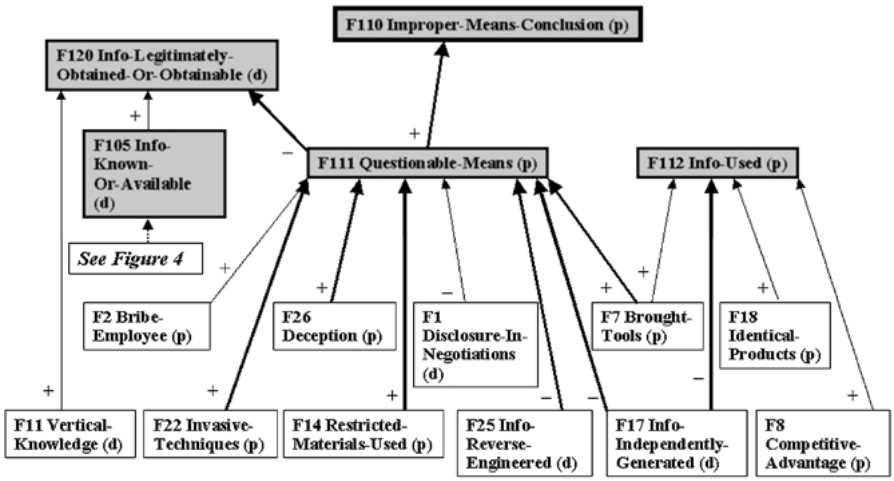

Figure 5. CATO's Factor Hierarchy for Trade Secret Misappropriation Issue: Whether defendant obtained plaintiff's information using improper means.

Each abstract factor represents two possible, opposing conclusions about its legal effect in a problem: one conclusion is that it favors the plaintiff and the other that it favors the defendant (Aleven 1997, p. 44–49). The name of the abstract factor refers to one of the two associated conclusions and the side it favors, plaintiff (p) or defendant (d). For example, factor F102, Efforts-To-Maintain-Secrecy, in Figure 4, represents two possible conclusions, that "Plaintiff took efforts to maintain the

secrecy of its information" or that "Plaintiff showed a lack of interest in maintaining the secrecy of its information". Each factor or abstract factor is linked to (at least one) parent abstract factor, and supports one or other of the conclusions associated with the parent. A child node's pro-plaintiff factor or abstract factor supports the pro-plaintiff conclusion in the parent node (and, indeed, in all of its ancestor nodes.) A pro-defendant factor supports the pro-defendant conclusion in the parent node (and all of its ancestor nodes). In Figure 4, for example, all of the factors and abstract factors linked to F102 provide evidence for or against the conclusions associated with F102, which itself provides evidence concerning the more abstract issue of whether the information is a trade secret, F101.

Significantly, a particular factor may relate to more than one abstract factor (including more than one issue). In the Factor Hierarchy, this occurs where a factor node has more than one parent (i.e., the Factor Hierarchy is a graph.) As discussed below, CATO makes use of this property in generating arguments emphasizing or downplaying the significance of a distinction.

The degree of support a factor (or abstract factor) lends to a conclusion may be strong or weak, represented as thick or thin links in Figures 4 and 5 (Aleven 1997, pp. 44–49, 64–66). This degree of support is *not* a numerical weight. Instead, it helps CATO identify those issues for which the evidence is in conflict and blocks certain inferences. For instance, in the Factor Hierarchy, Figure 4, there is a strong negative link from F23, Waiver-Of-Confidentiality (d), to F102, Efforts-To-Maintain-Secrecy (p). For certain purposes, this link would block CATO from making the inference that plaintiff undertook efforts to maintain the secrecy of the information. Although typically a target problem presents evidence for and against a high-level conclusion, as discussed below, the Factor Hierarchy is not the primary vehicle for resolving such conflicts. For that purpose CATO makes arguments citing conflict-resolving cases. *See infra*, Section 4.1.2.[7]

CATO employs the Factor Hierarchy to characterize the importance of factors present in a target problem or to characterize the significance of differences between a target problem and source case. Specifically, the Factor Hierarchy has four roles. It helps to: (1) explain the significance of factors, (2) identify issues in a case or problem, (3) organize an argument by issues, and (4) reason in a context-sensitive way about the significance of similarities and differences among cases. When students click on a factor in the Factor Browser tool, they see a description of the factor, which explains its significance in terms of the law's purposes based on the factor's connections to higher-level abstract factors and to the top-level legal issues in the Factor Hierarchy. For example, F4 Agreed-Not-To-Disclose (p) relates to abstract factors F115, F121 and F102. It shows defendant was on notice that using or disclosing the information would be a breach of confidentiality, there was an express agreement to keep the information confidential, and plaintiff took efforts to maintain the secrecy of its information.

The Factor Hierarchy's second role is to help identify issues in a case or problem (Aleven 1997, pp. 44–50). For instance, in the *Mason* case, there is evidence

that the plaintiff (a bar owner) "told the recipe only to his bartenders and instructed them not to reveal the recipe to others". A student entering the problem, let us assume, knows that based on such evidence, one may infer that Factor F6, Security-Measures (p), applies in the problem, and he/she so indicates, along with whatever other factors may apply. In processing the *Mason problem*, CATO can trace in the Factor Hierarchy, shown in Figure 4, the connections from factor F6 (at the bottom) to abstract factor F102, Efforts-To-Maintain-Secrecy (p), to issue F101, Info-Trade-Secret (p). By virtue of tracing this path, the program can make the following (provisional) inferences: Plaintiff should win its trade secret misappropriation claim because its efforts to maintain the secrecy of the information support the conclusion that the information is a trade secret. The Factor Hierarchy's remaining two roles are discussed below.

The factors, issues, and other abstract factors have been gleaned from a variety of sources including the Restatement (First) of Torts, Section 757, Liability for Disclosure or Use of Another's Trade Secret, which many jurisdictions have adopted as an authoritative statement of the law of trade secrets,[8] from the opinions of trade secret cases, and from secondary sources such as treatises and law review articles (e.g., Milgrim 1995). The secondary sources tend to group cases in footnotes that illustrate the effect on the decisions of particular factual strengths and weaknesses (or illustrate counterexamples, where a court decides in spite of the strengths or weaknesses.)

Organizing the factors into the Factor Hierarchy was a manual task conducted by Vincent Aleven in a trial and error process guided by legal issues identified in §757's statement of General Principle,[9] and various important cases.

Although the Factor Hierarchy is one reasonable interpretation of the important issues, concepts, and fact-patterns in trade secret law, it is certainly not the only such interpretation. A designer beginning with the Uniform Trade Secrets Act,[10] Restatement 3d on Unfair Competition Sections 39-45 or a state trade secret statute (or a combination of one or more of these sources) might well develop a somewhat different Factor Hierarchy.

Other types of legal claims, beside trade secret misappropriation, would involve their own sets of factors, although a given factor could be relevant to more than one type of legal claim.

Deciding which factors apply to a particular source case is a manual task for a human indexer to perform, and, indeed, a subjective task. Our goal has been to design CATO in such a way as to minimize the consequences of this inevitable subjectivity. For one thing, there is by now a fairly stable list of factors to choose from. For another, the Factor Browser tool specifies conditions under which each factor applies. As a result, under the supervision of a law professor it is possible for a set of graduate or law student research assistants to index cases in a fairly consistent way.

We have also implemented techniques to minimize the consequences of subjectivity when students enter problems. Students using CATO need not agree with

Step 1: Make a prediction (i.e., state a hypothesis)
Step 2: Frame a CATO query to test the hypothesis
Step 3: Interpret the results of the query
Step 4: If necessary, modify the hypothesis in light of the retrieved cases. (Aleven 1996, *Workbook I* and *Reference Pages*, Section 7, 1997)

Figure 6. Workbook process for testing legal hypothesis.

CATO's representation of a given case. As described above, the representation of a case in the database has been prepared manually by a case enterer and represents that person's best effort to identify the plaintiff's factual strengths and weaknesses in terms of CATO's factors. Since interpreting the facts of a case is to some extent a subjective opinion, students may select factors not included in CATO's list of factors for a given case. In effect, students may disagree with the case enterer. If a student rejects the Case Analyzer's feedback, the student is encouraged to stand by his/her representation of a case. In subsequent processing of the case as a problem situation, such as generating argumentation examples, CATO will employ the student' representation of that case, not its own (Aleven 1997, pp. 104–107).

Although CATO's set of factors for trade secret law is fairly complete, it may also occur that students identify a factual strength or weakness for which they cannot find a corresponding factor. In this event, students are encouraged to use the closest corresponding factor they can find. The program does not support students' adding to CATO's list of factors (Aleven 1997, pp. 104–107).

3.2. USING CATO TO TEST HYPOTHESES ABOUT A LEGAL DOMAIN

After introductory exercises introducing students to using factors to represent factual strengths and weaknesses in a case, the CATO curriculum divides naturally into two parts: using factors to express and test hypotheses about trade secret law and argumentation.

CATO and its Workbook instruct students in one of the main tasks for which factors are useful: expressing hypotheses or theories about a domain of law in terms of stereotypical factual strengths and weaknesses and testing them against a database of cases (Aleven 1997, pp. 108–118). These exercises instruct students about the utility of empirical legal research and about the ways cases may confirm or disconfirm hypotheses. CATO's simplification of the research task, it is hoped, can make the process of testing a hypothesis more apparent to students (Aleven 1997, pp. 118–119). The exercises also introduce students to using CATO's Database and query language.

In instructing students to test legal hypotheses, the Workbook leads students through a four-step process shown in Figure 6:

Some hypotheses to test are introduced in two discussion questions in the Casebook following the *Forest Laboratories* case[11] whose facts prompt the particular

hypotheses (Ashley and Aleven 1996, Parts I and II, Section 3). The first exercise is:

> 3.5 Suppose a defendant to whom confidential information was disclosed knew that the information was confidential, but there was no written non-disclosure agreement? Is defendant under an obligation not to use or disclose the information?

Students first predict the answer they expect to find in trade secret law. For instance, in 3.5, students might predict, "Defendants to whom confidential information is disclosed but who have not entered into a written non-disclosure agreement, are under *no* obligation not to use or disclose the information even though they know that the plaintiff regards the information as confidential". Alternatively, they may predict that the defendants are under an obligation.

Next, students test their prediction using CATO's Database Window to frame a CATO query. In order to "translate" the prediction into a specification of factors, they must find (with the Factor Browser window) the CATO factors, which best correspond to the elements in the prediction. For instance, the following query represents one interpretation of the circumstances in problem 3.5:

> List all cases with factors F1, Disclosure-In-Negotiations (d), F21, Knew-Info-Confidential (p), but without factor F4, Agreed-Not-To-Disclose (p).

For this kind of query, students would type the following in the Query Entry box: f1 f21 % f4 (where % means "not". An implicit "and" requires that all three constraints be satisfied.) This query returns eight cases won by the plaintiff and one case won by the defendant (Aleven 1997, p. 111). Students next interpret the prediction in light of the query's results. In particular, step 3 involves students in reading some or all of the cases returned (in the Squib Reader window) to determine whether the cases confirm or contradict the student's hypothesis. A student who predicted no obligation not to use or disclose the information, for example, might be surprised to see so few cases won by defendant and so many won by plaintiff, results which appear to be inconsistent with the hypothesis. Conversely, a student who had predicted that defendants *did* have an obligation not to use or disclose the information would be interested in examining and explaining away the one case won by the defendant.

Determining whether such a hypothesis is correct, of course, is not simply a matter of comparing numbers of returned cases either way. A student's prediction might still be valid despite lots of apparently contradictory cases. The pedagogical point is that the query results provide a context in which students investigate the cases that appear to be inconsistent with their predictions to determine whether the inconsistency is real.

With respect to the sample query above, a student, who predicted that there is no obligation not to use or disclose the information, needs to read the pro-plaintiff cases to determine whether all, most, or any of these cases can be "explained away". In other words, students need to determine whether they can think of some

rationale which both explains why those cases were decided as they were and is consistent with the student's hypothesis. Students are encouraged to examine the court's holding and decision (a short summary of which is reported in the Squib Reader) for a possible rationale. They may examine with the Database window the cases' factor representations to see, for example, how many of the cases have additional pro-plaintiff factors (i.e., factors in addition to F21, Knew-Info-Confidential) that might form a basis for explaining away the case and thus saving the hypothesis.

If the student can find an alternative explanation for the decisions in those cases based on their facts, he/she may rescue the hypothesis (Aleven 1996, *Workbook 1* and *Reference Pages*, Section 7). If not, these cases will force the student to modify or even abandon his/her hypothesis. (The workbook reminds students that CATO's database contains a small fraction of all trade secret cases, another reason why a conclusion either way may still be only provisional.)

In the fourth step, if the student cannot explain away the retrieved cases that are inconsistent with the hypothesis, the workbook encourages him/her to modify the hypothesis in light of the results. That might mean revising the prediction all together (e.g., instead of predicting no obligation, predicting that there is an obligation not to use or disclose the information.) Alternatively, a student's analysis of the inconsistent cases may indicate a qualification (e.g., an additional factor), which could strengthen the hypothesis. Having modified the hypothesis, students can test it by iterating the four-step process in Figure 6 (Aleven 1997, pp. 114–118).

3.3. CATO'S ARGUMENTATION CURRICULUM

The second part of the curriculum comprises a cumulative series of exercises involving argumentation with cases, which culminates in students' writing a short brief for the plaintiff and a response on behalf of the defendant in a problem situation. The lessons introduce students to the seven basic argument moves CATO implements and then guide them in composing more complex multi-case arguments organized around relevant legal issues. The next sections present examples of arguments CATO employs to illustrate the following topics: (1) analogizing, resolving conflicting factual strengths and weaknesses, and distinguishing, (2) making arguments about the importance of distinctions, and (3) organizing an argument around issues. The last topic deals with composing the basic moves into multi-case arguments organized by legal issues.

Most of these arguments appear in CATO's Argument Maker window. As pictured in Figure 7, this window has five buttons across the top corresponding to five basic argument moves: Analogize, Distinguish, "Not Fatal", Emphasize Distinction and Downplay Distinction (Aleven 1997, p. 124). Having selected a target problem and a source case from CATO's database, a student may see how CATO would use that case (if possible) in the corresponding argument move concerning the problem.

In presenting an argument, the Argument Maker employs three vertically stacked window panes, shown in Figure 7 on the right. In the top pane, it depicts

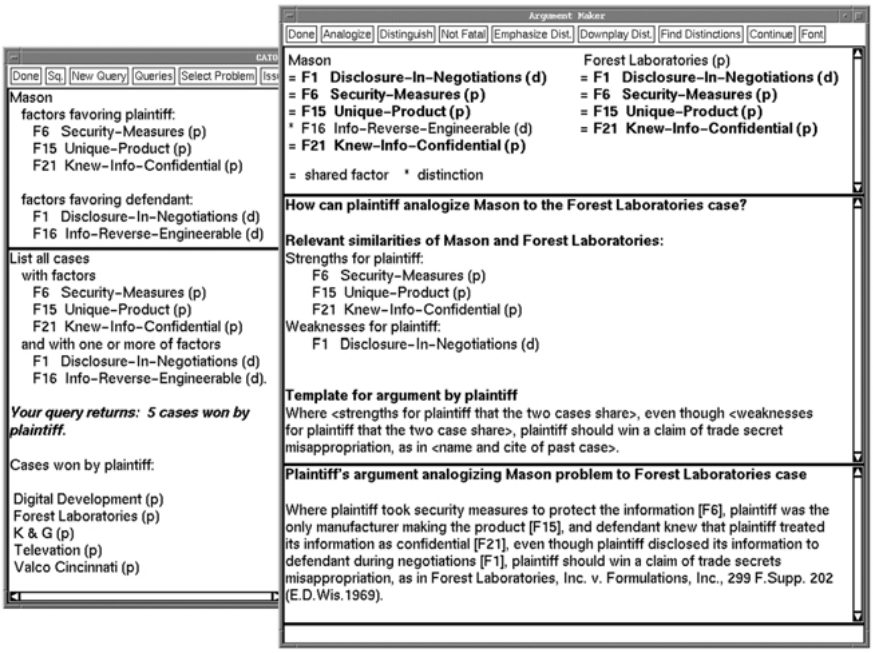

Figure 7. CATO's Argument Maker Window.

the comparison of the target problem and source case in terms of shared and unshared factors. The middle pane usually provides a general template or recipe for generating the argument move. The bottom pane demonstrates the actual argument created by filling in the template or following the recipe using the top pane's factor comparison information concerning the target problem and source case.

The simplest explanation of how CATO works is this: it "knows" when and how to fill out these argumentation templates and recipes. Its algorithms enable it to determine whether an argumentation template or recipe is appropriate and what information to plug into the template's slots or to select in response to the recipe's directions. It can select such information by virtue of its integrated knowledge representation. Its argumentation techniques (i.e., argumentation templates, recipes and algorithms) interact with its database of cases indexed by factors and the Factor Hierarchy. As I proceed through the examples, I will indicate how these knowledge sources are used.

3.3.1. *Analogizing, resolving factual strengths and weaknesses, and distinguishing*

Students first encounter the argument moves for analogizing and distinguishing in connection with the *Mason* target problem, shown in Figure 3. Although the *Mason* case is one of the CATO database's real legal cases, when using it as a target problem, the case is removed from the database so that neither the program

nor the students knows the outcome. In the Squib Reader window, the students see only the part of the *Mason* squib detailing the fact situation.

In the CATO model, analogizing a target problem and a source case involves making an argument that the target problem should be decided as the source case was, based on the relevant similarities between the two. Relevant similarities are defined as the set of factors, which the target problem and source case have in common (Aleven 1997, pp. 19–25, 58–61; Ashley 1990, pp. 29–34). These similarities give rise to reasons for deciding the two cases in the same way. (As discussed below, the reasons are derived from the abstract factors and issues to which the shared factors are linked in the Factor Hierarchy.) Distinguishing a case involves making an argument that the target problem should not be decided as the source case was decided, based on their relevant differences. Relevant differences are defined as a subset of the set of factors, which each case does not share with the other. In particular, the subset contains those unshared factors in the target problem, which favor the distinguisher's side, and those unshared factors in the source case favoring its winner. These differences give rise to reasons for deciding the two cases differently (again, derived from the Factor Hierarchy). The former strengthen the distinguisher's position in the target problem in ways not present in the source case. The latter strengthen the position of the source case winner in ways not present in the target (Aleven 1997, pp. 19–25, 58–61; Ashley 1990, pp. 29–34).

The following example illustrates how CATO compares a target problem and source case. Figure 8 shows the Argument Maker window's comparison of the *Mason* target and the *Forest Laboratories* source case. The comparison lists each case's factors and annotates them. Relevant similarities (i.e., shared factors) are marked with an "=". Relevant differences are marked with an "*". Other unshared factors (i.e., those that do not amount to distinctions) would be left unmarked.

The plaintiff in the *Mason* target having cited the *Forest Laboratories* case, the defendant can distinguish it by pointing out that factor F16 is not shared. Since this factor favors the defendant in *Mason* but not in *Forest Laboratories*, it is a distinction (i.e., a relevant difference).

In its middle pane, the Argument Maker window presents templates for using the case-comparison information to make arguments analogizing and distinguishing the target problem and source case. Figure 8 shows the templates for a plaintiff to analogise the target and source case, resolve a conflict among strengths and weaknesses, and for the defendant to distinguish the source case from the target.[12]

The template for analogizing/conflict-resolution induces students, in effect, to propose a rule for explaining the analogy between the target problem and the source case. I argue below that CATO's analogizing and distinguishing move templates serve the functions of Brewer's AWRs and DWRs. The rule is of the form:

> IF plaintiff's shared strengths in [a] apply THEN decide claim for plaintiff EVEN THOUGH plaintiff's shared weaknesses in [b] also apply.

The "even though" clause of [b] says, in effect, "Plaintiff's strengths outweigh the weaknesses. Here is a source case (i.e., a precedent) which involved the very

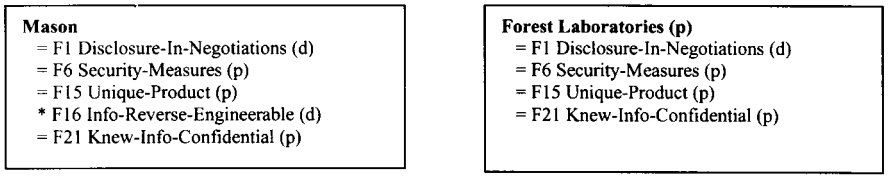

Template for Analogizing/Conflict-Resolution Move: [a] Where ⟨Insert strengths for plaintiff/defendant shared by target problem and source case⟩, *[b] even though ⟨Insert weaknesses for plaintiff/defendant shared by target problem and source case⟩*, ⟨Insert plaintiff/defendant⟩ should win a claim of trade secret misappropriation, as in ⟨Insert name and cite of source case⟩.

Distinguishing Move Template:
⟨Insert name of the source case⟩ is distinguishable.
It is stronger for ⟨Insert plaintiff/defendant⟩ than is the current problem.
[c] In the ⟨Insert name of the source case⟩, ⟨Insert extra strengths for plaintiff/defendant in the source case⟩.
This was not so in the current problem.
[d] Also, in the current problem, ⟨Insert extra weaknesses for plaintiff/defendant in the target problem⟩.
This was not so in ⟨Insert name of the source case⟩.

Analogizing/Conflict-Resolution Move: Where plaintiff adopted security measures (F6), plaintiff was the only manufacturer making the product (F15), and defendant knew that plaintiff's information was confidential (F21), even though plaintiff disclosed its product information in negotiations with defendant (F1), plaintiff should win a claim of trade secrets misappropriation, as in *Forest Laboratories, Inc. v. Formulations, Inc.*, 299 F. Supp. 202 (E.D. Wis. 1969).

Distinguishing Move: *Forest Laboratories, Inc. v. Formulations, Inc.* is distinguishable, because in Mason, plaintiff's product information could be learned by reverse-engineering (F16). This was not so in *Forest Laboratories*.

Figure 8. Factor comparison for *Mason* and *Forest Laboratories*, templates for analogizing/conflict resolution move and distinguishing move, and actual moves.

same conflicting strengths and weaknesses and which resolved them in favor of the winning side".

The template for distinguishing in Figure 8 leads students, in effect, to propose an alternative rule for explaining the distinguished source case, a rule that would no longer apply to the target problem. From the distinguisher's viewpoint, the rule of [a] and [b] offered in the analogizing/conflict-resolution move needs two kinds of modifications. First, at [c], the distinguishing template points out unshared strengths, if any, in the source case. The template implies that the revised rule for the source case should include these extra strengths of plaintiff as additional terms. Such a modification, however, renders the revised rule too specific to apply to the target (i.e., the rule explains a disanalogy between source and target). Second, at [d], the distinguishing template identifies unshared weaknesses that make the

target problem an exception to the rule offered in the analogizing move. In effect, the revised rule needs yet another additional term specifying the absence of those unshared weaknesses.

These arguments serve to communicate to students a basic model for using information about factual strengths and weaknesses in legal analogical arguments. In effect, given a target problem and source case, the students are taught to propose a rule for explaining the analogy between the target problem situation and source case in terms of the intersection of the target's and source case's sets of relevant factual strengths, [a] in Figure 8.

Resolving conflicting factual strengths and weaknesses is, of course, important in legal argument. As in most disputed cases, the plaintiff in the *Mason* problem has some factual strengths but also some weaknesses. The CATO model teaches three techniques for making arguments how to resolve such conflicts:

1. Citing source cases to show how courts have resolved the conflicts.
2. Citing source cases to show that the weaknesses are not fatal.
3. Arguing that the distinguishing extra strengths or extra weaknesses are not important.

The first and main technique is to cite cases that have resolved such conflicts in the past. CATO teaches students to seek cases that involve the same claim, share the same conflicting factual strengths and weaknesses, and where the corresponding side won. With cases like these, students can fill out the analogizing template's clauses [a] and [b].

Students try some queries to retrieve cases from CATO's database that share as large a set of *Mason*'s strengths and weaknesses as possible. One obvious query to try first is: f6 f15 f21 f1 f16. It seeks cases that share all of *Mason*'s strengths and weaknesses. As it turns out, however, the query is too restrictive. No case in the database has all of these factors. A less restrictive query is: f6 f15 f21 (or f1 f16). It seeks cases that share all of *Mason*'s pro-plaintiff strengths (Factors F1, F15, and F21), and one or more of its pro-defendant weaknesses (Factors F1 or F16). Among the cases this query returns are the *Forest Laboratories* and the *Televation* case,[13] discussed below (Aleven 1996, *Workbook* 2, Section 4).

The rule for explaining the analogy between the *Mason* target and the *Forest Laboratories* source, a case, which resolved some of the same conflicting strengths and weaknesses, states explicitly that the plaintiff prevails despite those weaknesses. When a target case resolves conflicting strengths and weaknesses, clause [b] in the template for the analogizing/conflict-resolution move, Figure 8, is not empty. The argument CATO makes for plaintiff with the *Forest Laboratories* case shows how one court has resolved the conflict between the factual strengths represented by Factors F6, F15 and F21 on the one hand, and the weakness F1 on the other.

Combining cases to overcome factual weaknesses is another argumentation strategy called "covering the bases". As CATO's argument distinguishing the *Mason* target from *Forest Laboratories* indicates, there is still one factual weakness in

Mason not accounted for: F16 Info-Reverse-Engineerable. The same query which returned *Forest Laboratories*, however, also returns a case that can form the basis of a similar argument resolving F16 in favor of the plaintiff: the *Televation* case. Together, *Forest Laboratories* and *Televation* support an argument that plaintiff should win despite both these weaknesses.

The second technique for arguing how to resolve conflicting factual strengths and weaknesses is to argue that the weaknesses are not fatal. The "not fatal" argument move anticipates an opponent's hypothesis that the weaknesses are fatal and provides a counterexample (Aleven 1997, p. 72). Conversely, if it proves impossible to find any cases to use in a "not fatal" move, one may have more reason to conclude that the weakness is significant. Since they focus on using counterexamples to test assertions about the significance of factors, "not fatal" moves also play an important role in the hypothesis-testing process earlier in the curriculum.

The activities of analogizing, distinguishing, and conflict-resolution are closely related to the hypothesis-testing process (Figure 6). The analogical argument and its explanatory rule are like the hypotheses students tested. Distinguishing the source case by further specifying the rule or stating an exception for the target are techniques students would use in modifying a hypothesis and explaining away counterexamples to a hypothesis.[14]

3.3.2. *Arguments about importance of distinctions*

A third technique for arguing how to resolve conflicting factual strengths and weaknesses is to argue that the extra factual strengths, which the distinguisher has focused upon (in [c] of the distinguishing template), and the extra factual weaknesses (in [d]) are not very important in the target problem. Conversely, a distinguisher can argue that these distinctions are very significant.

The Argument Maker window provides a recipe for downplaying a distinction *d*, shown in Figure 9. The idea is to argue that, at a more abstract level of description, the target problem and source case really are alike. One way to do this is by pointing out undercutting factors in the target problem (i.e., factors in the target which tend to contradict the reason why *d* matters.) This is an argument that the distinguishing factor is not really a concern in the target problem. Another is to show similar factors in the source case (i.e., factors which matter for the same reason as *d*) from which it may be inferred that the presence of *d* does not make the cases significantly different. (Aleven 1997, pp. 67ff).

By following the recipe, one can argue that the distinction does not matter because other facts in the target problem counteract the concern and (if possible) that, in any event, similar facts were present in the source case (Aleven 1996, *Workbook 2*, Section 3).

The idea underlying the recipe for emphasizing a distinction is to give a reason why the distinction matters in the target problem, preferably a reason that does not apply in the source case (because of the presence of contrasting factors in the source case or, at least, the absence of similar factors). If possible, one should also

Recipe for Downplaying Distinction:

1. Recite reasons why distinction *d* might matter in the cfs [source case].
2. Recite factors in cfs [source case] which contrast with or undercut *d* (i.e., matter for exactly the opposite reason.)
3. Recite factors in the source case [cfs] which matter for the same reason (i.e., are similar to *d*.)

Recipe for Emphasizing Distinction:

1. Recite reasons why distinction *d* matters in the cfs [source case] which are not contradicted in the cfs [source case] and not present in the source case [cfs].
2. Recite factors in cfs [source case] which matter for the same reason (i.e., are similar to *d*.)
3. Recite factors in the source case [cfs] which contrast with or undercut *d* (i.e., matter for exactly the opposite reason) and which are not present in the cfs [source case].

Figure 9. Recipes for moves downplaying/emphasizing distinctions.

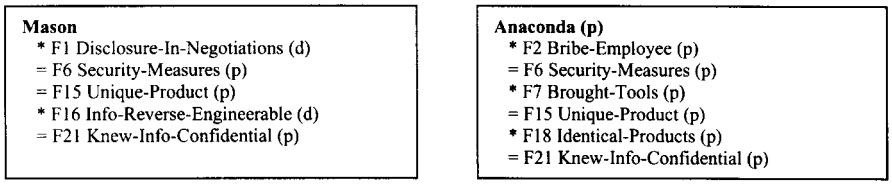

Figure 10. Factor comparison of *Mason* target problem and *Anaconda* case.

point to other facts in the target case that matter for the same reason (Aleven 1996, *Workbook 2*, Section 3).

An example of downplaying and emphasizing differences may be seen in the comparison of the *Mason* target problem and the *Anaconda* case,[15] shown in Figure 10. CATO's arguments downplaying and emphasizing two of the distinctions are shown in Figure 11 (left).

In downplaying the distinction F1, the fact that Mason disclosed information in negotiations with the Distillery, CATO finds another reason (i.e., factor) on which to base an assertion that Mason took steps to maintain the secrecy of its information. In effect, plaintiff reiterates its analogy-explaining rule based on *Anaconda*, that where plaintiff takes security measures, it should win, and suggests that in light of those security measures, the disclosure to defendant does not matter.

The gist of the emphasizing argument, on the other hand, is to suggest that *Mason* is a very different kind of case from *Anaconda*, one where the allegedly secret information readily could be had by legitimate means. The defendant explains away *Anaconda* as a case where defendant used improper means. Here the program points out that F2 Bribe-Employee (p) and F7 Brought-Tools (p) in *Anaconda* imply that the defendant used improper means to gain access to the plaintiff's information.

As discussed below, CATO uses the Factor Hierarchy to categorize the cases more abstractly in terms of the higher-level abstract factors or legal issues. The arguments in Figure 11, show the numbers (in the Factor Hierarchy) of the factors

1. Arguments re Distinction F1 Disclosure-In-Negotiations (d): Plaintiff disclosed its information to defendant during negotiations (in *Mason*, not in *Anaconda*).

Downplaying: Plaintiff's argument downplaying the distinction F1:
In *Mason*,[16] plaintiff disclosed its product information in negotiations with defendant [F1]. This was not so in *Anaconda*.[17] This however is not a major distinction. In spite of the fact that in *Mason*, plaintiff disclosed its information to defendant during negotiations [F1], there was evidence that plaintiff took efforts to maintain the secrecy of its information [F102]: In *Mason*, plaintiff adopted security measures [F6].

Emphasizing: Defendant's argument emphasizing the distinction F1:
In *Mason*, plaintiff disclosed its product information in negotiations with defendant [F1]. This was not so in *Anaconda*. This is a marked distinction. It shows that in *Mason*, defendant obtained or could have obtained its information by legitimate means [F120]. Other facts in *Mason* further support this: Plaintiff's product information could be learned by reverse-engineering [F16]. In *Anaconda*, by contrast, defendant may have acquired plaintiff's information through improper means [F120]: Defendant offered plaintiff's employee a raise or bonus to switch employment, in order to gain access to plaintiff's information [F2] and plaintiff's former employee brought product development information to defendant [F7].

2. Arguments re Distinction F16 Info-Reverse-Engineerable (d): Plaintiff's information could be discovered by reverse engineering plaintiff's product (in *Mason*, not in *Anaconda*).

Downplaying: Plaintiff's argument downplaying the distinction F16: None.

Emphasizing: Defendant's argument emphasizing the distinction F16:
In *Mason*, plaintiff's product information could be learned by reverse-engineering [F16]. This was not so in *Anaconda*. This is a clear and striking distinction.
It shows that in *Mason*, defendant obtained or could have obtained its information by legitimate means [F120]. Other facts in *Mason* further support this: Plaintiff disclosed its information to defendant during negotiations [F1]. On the other hand, in *Anaconda*, defendant may have acquired plaintiff's information through improper means [F120]: Defendant offered plaintiff's employee a raise or bonus to switch employment, in order to gain access to plaintiff's information [F2] and plaintiff's former employee removed documents containing plaintiff's information [F7]....

Between the *Mason* Target Problem and *Anaconda* Case

1. Arguments re Distinction F1 Disclosure-In-Negotiations (d): Plaintiff disclosed its information to defendant during negotiations (in *Mason*, not in *Anaconda*).

Downplaying: Plaintiff's argument downplaying the distinction F1:
In *Mason*, plaintiff disclosed its product information in negotiations with defendant [F1]. This was not so in *Anaconda*. However, this difference is insignificant. First, even though in *Mason*, plaintiff conveyed its information to defendant in the course of negotiations [F1], defendant may have acquired plaintiff's information through questionable means [F111]: In *Mason*, defendant procured plaintiff's information by deceit [F26]. Second, in spite of the fact that in *Mason*, plaintiff conveyed its information to defendant in the course of negotiations [F1], there was evidence that plaintiff took efforts to maintain the secrecy of its information [F102]: In *Mason*, plaintiff adopted security measures [F6].

Emphasizing: Defendant's argument emphasizing the distinction F1: None.

2. Arguments re Distinction F16 Info-Reverse-Engineerable (d): Plaintiff's information could be discovered by reverse engineering plaintiff's product (in *Mason*, not in *Anaconda*).

Downplaying: Plaintiff's argument downplaying the distinction F16:
In *Mason*, plaintiff's product information could be learned by reverse-engineering [F16]. This was not so in *Anaconda*. However, this does not amount to an important distinction. In spite of the fact that in *Mason*, plaintiff's information could be discovered by reverse engineering plaintiff's product [F16], defendant may have acquired plaintiff's information through improper means [F120]: In *Mason*, defendant deceived plaintiff to gain access to the information [F26].

Emphasizing: Defendant's argument emphasizing the distinction F16:
In *Mason*, plaintiff's product information could be learned by reverse-engineering [F16]. This was not so in *Anaconda*. This distinction is highly significant. It shows that in *Mason*, plaintiff's information was available from sources outside plaintiff's business [F108]. This was not so in *Anaconda*....

Between the Modified *Mason* Target Problem and *Anaconda* Case

Figure 11. CATO's arguments downplaying/emphasizing distinctions.

and those of the legal issues or other abstract factors employed to characterize the target and source case. For instance, in the emphasizing argument, it applies abstract factor F120 in a pro-defendant sense to *Mason*. F120 can be seen in the Factor Hierarchy in Figure 5. At the same time, CATO characterizes *Anaconda* as a case in which defendant used improper means, applying F120 in a pro-plaintiff sense. In Section 3.4 infra, I illustrate the context-sensitivity of CATO's algorithms for downplaying and emphasizing distinctions by contrasting these arguments with those CATO generates in comparing the *Anaconda* case and a slightly modified version of the *Mason* target.

The argument moves downplaying and emphasizing distinctions are related to the hypothesis testing process of Figure 6. A hypothesis can be generalized by expressing it in terms of the more abstract characterizations. A more abstract characterization may be used to explain away a potential counterexample to a hypothesis. It may also suggest an underlying rationale for explaining the hypothesis.

3.3.3. *Organizing an argument around issues*

Finally, the CATO instruction helps students construct a multi-case argument on behalf of a plaintiff in a trade secret misappropriation claim, organized by legal issues, and composed of multiple applications of the above argument moves.

Although CATO has a recipe for organizing such arguments, it is fairly elaborate. Instead of presenting it to students directly, the Workbook presents a more discursive description and illustrates it with, first, an outline of an argument for the plaintiff in the *Mason* problem, and later a complete argument incorporating the cases a student has selected for his/her argument. Then, CATO generates an outline for an argument responding on behalf of the defendant followed by a full argument. The Workbook encourages students to use the Issues tool to compare their arguments with CATO's.

Each outline and argument relate the plaintiff's [defendant's] factual strengths and weaknesses to relevant legal issues and provide notes to the student specifying the types of cases needed to support the argument. CATO's issue-based argument on behalf of plaintiff in the *Mason* problem is shown in Figure 12 (Aleven 1997, p. 146). It employs eight cases students will already have encountered in the Workbooks. Organized around three issues, the argument claims that plaintiff's information is a trade secret, a confidential relationship existed between plaintiff and defendant, and that the defendant acquired plaintiff's information through improper means. The "F" numbers refer to the numbers of the applicable factors and abstract factors in the Factor Hierarchy.

Perusing the sample argument shows examples of three of the argument moves: Analogizing, Conflict-Resolution, and Not Fatal. The corresponding argument for defendant would include these three moves applied to pro-defendant cases, as well as Distinguishing and Counterexample moves to respond to plaintiff's cases.

Argument for Plaintiff in the Mason v. Jack Daniel Distillery Problem
Plaintiff should win a claim of trade secrets misappropriation. Plaintiff's information is a trade secret [F101], a confidential relationship existed between plaintiff and defendant [F114], and defendant acquired plaintiff's information through improper means [F110].

Plaintiff's information is a trade secret [F101]
Plaintiff's information is a trade secret [F101]. Restatement 1st of Torts s 757, and Comment b, factors 1–6 (1939). In the problem at hand, plaintiff took security measures to protect the information [F6] and plaintiff was the only manufacturer making the product [F15]. This shows that plaintiff took efforts to maintain the secrecy of its information [F102], the information apparently was not known or available outside plaintiff's business [F105], and plaintiff's information was valuable for plaintiff's business [F104]. In Television Telecommunication Systems, Inc. v. Saindon, 522 N.E.2d 1359 (Ill.App. 2 Dist. 1988), Anaconda Co. v. Metric Tool & Die Co., 485 F.Supp. 410, USM Corp. v. Marson Fastener Corp., 379 Mass. 90, 393 N.E.2d 895 (1979), Valco Cincinnati, Inc. v. N & D Machining Service, Inc., 24 Ohio St.3d 41, 492 N.E.2d 814 (1986), and Forest Laboratories, Inc. v. Formulations, Inc., 299 F.Supp. 202 (E.D.Wis.1969), which held for plaintiff, there was similar evidence that plaintiff's information is a trade secret [F101]. In *Television*, *Anaconda*, *USM*, *Valco Cincinnati*, and *Forest Laboratories*, which held for plaintiff, plaintiff adopted security measures [F6] and plaintiff's product was different from products made by competitors [F15], as in the current case.
The fact that plaintiff disclosed its product information in negotiations with defendant [F1] does not rule out a conclusion that plaintiff's information is a trade secret [F101]. Kamin v. Kuhnau, 232 Or. 139, 374 P.2d 912 (1962), Smith v. Dravo Corp., 203 F.2d 369 (7th Circuit). This is especially so where, as in the current problem, plaintiff took measures to keep its information secret [F6]. *Valco Cincinnati*, *Forest Laboratories*.
The fact that plaintiff's product information could be learned by reverse-engineering [F16] does not preclude a conclusion that plaintiff's information is a trade secret [F101]. *Television*, *Dravo*, *Kamin*.

A confidential relationship existed between plaintiff and defendant [F114]
A confidential relationship existed between plaintiff and defendant [F114]. Restatement 1st of Torts s 757 (1939). In the current problem, defendant knew that plaintiff intended its information to be kept confidential [F21]. This shows that defendant was on notice that using or disclosing the information would be a breach of confidentiality [F115]. In *Television*, *Anaconda*, *Valco Cincinnati*, and *Forest Laboratories*, there was similar evidence that a confidential relationship existed between plaintiff and defendant [F114], and plaintiff won. In *Television*, *Anaconda*, *Valco Cincinnati*, and *Forest Laboratories*, which held for plaintiff, defendant knew that Plaintiff's information was confidential [F21], as in the current case.
The fact that plaintiff conveyed its information to defendant in the course of negotiations [F1] does not necessarily rule out a conclusion that a confidential relationship existed between plaintiff and defendant [F114]. *Forest Laboratories*, *Valco Cincinnati*, *Dravo*, *Kamin*.

Defendant acquired plaintiff's information through improper means [F110]
Defendant acquired plaintiff's information through improper means [F110]. Restatement 1st of Torts s 757 (1939). The fact that plaintiff disclosed its information to defendant during negotiations [F1] does not rule out a conclusion that defendant acquired plaintiff's information through improper means [F110]. *Forest Laboratories*, *Valco Cincinnati*, *Dravo*, *Kamin*.

Figure 12. CATO's issue-based argument for plaintiff in the *Mason* problem.

While the examples above demonstrate CATO's argument with eight cases students encounter in the workbook, CATO's ability to generate arguments is not limited to those cases. CATO can automatically generate an argument for any set of cases, which students select from its database. This makes it easy for students to compare their arguments with a given set of cases to the arguments CATO would make with those same cases.

3.3.4. *Berman's challenge: An algorithmic model of context-sensitive arguments with reasons*

CATO can interpret a case differently depending on how it is to be used in an argument. It takes the argument context into account in selecting what legal reasons to offer as justifications for its assertions that, given the precedents, a case's factors support a particular result. Working with the Factor Hierarchy its algorithms enable it to organize arguments by issues and to argue in a context-sensitive way about the significance of similarities and differences. (Aleven 1997, pp. 41–148, 248–252; Aleven and Ashley 1996, 1997; Ashley and Aleven 1997).

This goes a long way toward meeting the challenge of Berman and Hafner for case-based legal argument systems to provide reasons why factors matter. In addition, as discussed in Section 5 *infra*, modeling the interpretation of precedents is an area where an AI model of case- based argument may contribute to jurisprudential discussions of analogical legal reasoning.

CATO's three-step recipe for organizing an argument for plaintiff or defendant by issues comprises: (1) Identifying issues, (2) Organizing source cases by issues, and (3) Generating English text for the arguments organized by issues. Each step is broken into sub-steps implemented algorithmically. In generating the argument in Figure 12 (Aleven 1997, p. 57) for each of the possible issues, CATO has:

> [1a] Identified the factors in the Mason problem, which are related to the issue according to the Factor Hierarchy.
> [1b] Selected abstract factors in the Factor Hierarchy with which to characterize the significance of those factors and explain how they strengthen the plaintiff's position on the legal issue. (These abstract factors will be referred to as "intermediate legal concerns".)
> [2a] Selected from among the eight specified input cases those relevant to the issue (i.e., those whose factors shared with the target are related to that issue.) It deems these as appropriate for Analogizing Moves.
> [2b] Checked whether some of the strengthening factors may compensate for some of the weakening factors with respect to the issue (i.e., if they share an intermediate legal concern.)
> [2c] Selected the relevant cases, which can be used to justify a conclusion that plaintiff should prevail on the issue. It will use these cases in the following argument moves as appropriate: Analogizing, Conflict-Resolution, or "Not Fatal".

[3a] In composing the materials selected above into a textual argument that plaintiff should prevail on that issue, CATO stresses the target problem's factual strengths (i.e., favorable factors) related to that issue using reasons associated with the intermediate legal concerns.

[3b] It underscores the strengths by citing the favorable input source cases that share those strengths.

[3c] It covers the weaknesses related to that issue by reciting opposing factors and covering examples. Opposing factors are those strengths in the target problem, which are related to the issue in the Factor Hierarchy and may be seen as counteracting the weakness. Covering examples are input source cases that reached a favorable result despite having those opposing factors (e.g., "cover the bases" cases and "not fatal" cases.)

In following its recipe for generating multi-case, multi-issue arguments, CATO performs a simplified kind of reflective adjustment, as Brewer uses that term. It sorts the cases by the explanatory rules implicit in the argument by analogy and draws a line around those cases in which the strengths overcame the weaknesses. CATO does not make a slippery slope argument, but if an opponent cites counter-examples to attack the proposed line, CATO responds to them by distinguishing and downplaying. At almost every point, the Factor Hierarchy provides information relating factors to issues enabling the program to organize the argument.

CATO's algorithms for emphasizing or downplaying distinctions enable it to make fairly subtle choices how best to characterize relevant similarities and differences among cases for purposes of argument. Its algorithms implement the recipes for emphasizing or downplaying, Figure 9, in sufficient detail to enable the program to generate the corresponding arguments.

In emphasizing a distinction, the algorithm leads CATO to select an abstract characterization to use as a focal point for interpreting the target problem and source case as fundamentally different. It finds the right focal abstraction in the Factor Hierarchy. Given a distinction *d* to emphasize, the algorithm finds factor *d* in the Hierarchy, selects which upward path to take through the Hierarchy, and decides how high up that path to go in search of the right abstract factor to characterize the two cases as fundamentally different.

Having found the right focal point, the program characterizes the Distinction's significance in the source case accordingly, points out other factors in the source case which support the interpretation (i.e., "corroborating factors"), and shows factors in the target problem which support a contrary characterization (i.e., "contrasting factors").

In downplaying a distinction, on the other hand, the goal is to minimize the distinction's significance and find an abstract characterization to use as a focal point for interpreting the target problem and source case as fundamentally similar.

In each algorithm, the central step is to select the focal point for emphasizing (or downplaying). Using the Factor Hierarchy, CATO identifies the distinguishing factor *d*, the set of contrasting factors, and the set of corroborating factors. It then

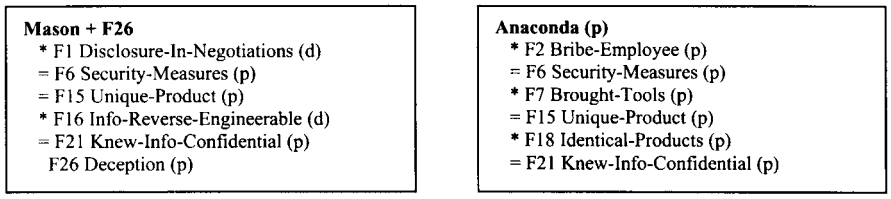

Figure 13. Factor comparison of *Anaconda* case and modified *Mason* problem.

selects the *most specific common ancestor* of these factors in the Factor Hierarchy. That is the lowest abstract factor to which *d* and all of the corroborating and contrasting factors are linked. There usually are many possible focal points. The algorithm finds them all and then filters unsuitable ones, consolidates others, and orders them by their estimated strength. Various constraints inform the generation and testing of focal points. They include: (1) characterizing the distinction broadly enough to employ all relevant contrasting and corroborating factors, (2) avoiding characterizing the distinction so broadly as to destroy the contrast or to allow obvious objections, and (3) making the most parsimonious argument possible. In other words, the focal point is located as high up in the Factor Hierarchy as needed, but not so high that the cases can no longer be contrasted abstractly (Aleven 1997, p. 78).

A small change to the *Mason* target illustrates the sensitivity of the algorithms for emphasizing and downplaying distinctions. For instance, suppose one adds Factor F26, Deception, to the facts. In other words, suppose that the Distillery's agent deceived the bar owner into disclosing a portion of his secret recipe. Figure 13 shows the comparison between *Anaconda* and the modified *Mason* target problem.

CATO's new arguments downplaying and emphasizing two distinctions (F1, F16) between *Anaconda* and the modified *Mason* target are shown in Figure 11 (right). Compare these to its former arguments comparing *Anaconda* and the original *Mason* target, Figure 11 (left).

Formerly, CATO dismissively characterized *Anaconda* as irrelevant in Figure 11 (left). It emphasized the importance of the fact that in *Mason*, plaintiff disclosed its product information in negotiations with defendant (factor F1 Disclosure-in-Negotiations). It characterized *Mason* as a case where defendant used legitimate means to obtain its information and *Anaconda* as a very different scenario where defendant employed improper means (abstract factor F120).

Given the modification in *Mason*, CATO can no longer make this argument on behalf of the defendant. Now the modified *Mason* target problem also may be characterized as an improper means case by virtue of defendant's procuring plaintiff's information by deceit (factor F26, Deception).

One side's loss is the other side's gain. A similar kind of reasoning gives CATO a new way of downplaying factor F16, Info-Reverse-Engineerable, on behalf of plaintiff, where formerly it had none. Some changes in CATO's ar-

gument are subtler. CATO still emphasizes factor F16, but instead of drawing a contrast between the cases in terms of abstract factor F120, Info-Not-Legitimately-Obtained-Or-Obtainable, it makes a more modest contrast in terms of abstract factor F108, Info-Available-Elsewhere. Here the algorithm has selected a path from the factor-to-be-emphasized, F16, to a different issue in the Factor Hierarchy. In Figure 5, abstract factor F120 is associated with the issue of Improper Means. By contrast, abstract factor F108 involves the issue of whether the information is a trade secret (Figure 4).

CATO's use of the Factor Hierarchy to represent the significance of factors abstractly in terms of the law's purposes and its criteria for selecting focal points, implemented in its algorithms for downplaying and emphasizing differences and for making issue-based arguments, enable it to select the right path and level of abstraction for characterizing the case flexibly in the argument context. (Aleven 1997, pp. 78–82; Ashley and Aleven 1997; Aleven and Ashley 1996).

As a result, CATO addresses Berman's and Hafner's challenge to develop case-based reasoning systems that can "explain why a factor favors the plaintiff or defendant, and why that factor is considered legally relevant" (Berman and Hafner 1993, p. 56).[18] Although the Factor Hierarchy's architecture is different, it is somewhat reminiscent of these authors' proposed scheme linking features and legal purposes and including oppositional links (Berman and Hafner 1993, Figure 2, p. 56). The issue-based, downplaying and emphasizing arguments in Figures 11 and 12 above, are very close to the manually produced arguments in Berman and Hafner (1993, Section 4, pp. 55–57). They draw abstract analogies between the target and source cases that emphasize the legal significance of the shared factors in terms of the purposes of trade secret law, such as protecting confidential relationships and discouraging unfair means of competition. They provide ways for a plaintiff or defendant to "bolster his response by stressing the importance", or lack of importance, of a factor (Berman and Hafner 1993, p. 57). On the other hand, while Berman and Hafner provided neither an algorithm nor an implementation for generating such arguments, CATO provides both (Aleven 1997).

4. Evaluating the CATO model from a jurisprudential viewpoint

CATO's computational model of case-based argument captures some important aspects of Brewer's and Sunstein's jurisprudential models of analogical legal reasoning. At the same time, it avoids some of the major problems for which Alexander criticizes the jurisprudential models.[19]

4.1. COMPARING CATO WITH SUNSTEIN'S AND BREWER'S MODELS

CATO's compositions of argument moves into issue-based arguments play the role Sunstein and Brewer see for legal analogical reasoning. Its arguments facilitate a legal analysis under conditions of uncertainty. CATO's model accounts

for analogizing and distinguishing, generates something very like analogy- (and disanalogy-) warranting rules and rationales, defines relevance in terms of analogy-warranting rules that are like low-level legal principles, and teaches a cognitive process similar to abductively formulating and testing an analogy-warranting rule.

4.1.1. *AWRs v. CATO's rule-like analogical explanations*

CATO's analogizing and distinguishing moves conform more-or-less closely to Brewer's specifications involving analogy- and disanalogy-warranting rules. CATO's source case analogies are always presented with what is, in effect, a rule intended succinctly to specify, as Brewer says, "in what [the source case's] exemplariness consists" (Brewer 1996, p. 975). CATO generates and teaches students to construct these succinct expressions in the analogizing move. Its analogy-explaining rules state that where certain conditions are satisfied, a side should win a specified legal claim.

On the other hand, there are some differences between Brewer's AWRs and CATO's analogy-explaining rules. The conclusion of an AWR looks like this: "If anything that has F and G also has H, then *everything that has F and G also has H*".

AWRs are more general than CATO's analogizing/conflict-resolution template (Figure 8) in two senses. First, the conclusions of AWRs deal with general legal properties. Of all the types of legal properties H to which an AWR's conclusion may refer, CATO's analogy-explaining rules deal with only a subset. The conclusion of the argument by analogy in Brewer's example is the appropriate standard of liability. CATO's examples deal only with the ultimate conclusion, who should win a legal claim.

Second, CATO's explanatory rules appear to apply only to the issue of who should win a legal claim in the context of the specific target problem. CATO does not emphasize the generality of the explanatory rule by calling it, for instance, a "general rule". Nor does the form of its rules contain a universal quantifier like "for all" or "for any" or otherwise emphasize a broader scope of application. Nevertheless, CATO's analogy-explaining rules can be seen as general legal propositions following a canonical form, which have a wider application than just the target problem. Logically, they would apply to any target problems that satisfy the antecedents. CATO's explanatory rules also are very similar to the hypotheses, which students learn to pose, test, and modify as described in Section 3.2. The CATO instruction clearly identifies these hypotheses as candidate general rules. Using the hypothesis-testing process of Figure 6, it is a simple matter to construct a query to test a decision rule against all the cases in CATO's database.

Another possible difference is that, as discussed above, the antecedent's conditions in a CATO analogy-explaining rule are expressed in terms of factors shared by the source case and target problem. The factors are used as abstract ways of describing stereotypical factual features strengthening one side or the other. Would Brewer and Sunstein accept a factor representation? Given the goal of building an

AI model, perhaps both authors would accept the need for using factors as a kind of simplifying representation, at least for certain legal domains.

A factor representation seems compatible with Brewer's example in Figure 1. To be sure, he does not describe the factual features referred to in the antecedent's conditions as strengths of one side or the other in an argument about the standard of liability. CATO's analogy-explaining rules, however, appear to be expressed at more or less the same level of abstraction as an AWR, judging from the antecedent in Brewer's example. In the legal domain of responsibility for theft losses of a travelling client's property, it would be reasonable to speak in terms of factors like F, Procuring-Room-For-Reasons-Of-Privacy and G, Vulnerability-Of-Clients-Property-To-Plunder that typically would strengthen the Client's argument for applying a strict liability standard. It is harder to characterize Sunstein's extended example in Section 2.2 (Sunstein 1993, pp. 759–767) as involving something like factors. Since the argument is about the constitutionality of alternative formulations of a statute prohibiting cross-burning, the argument and responses occur at a higher level of abstraction.

CATO's model complies with Brewer's requirement, as discussed above in Section 2.1, that a compelling analogy must have an analogy-warranting rationale (AWRa). The rationale should explain why, in the "eyes of the law", "the logical relation among the characteristics articulated by the analogy-warranting rule either does obtain or should obtain" (Brewer 1996, p. 965). CATO produces analogy-warranting rationales in its issue-based arguments and its arguments downplaying distinctions. In relating individual factors to the higher-level legal issues, these arguments more abstractly describe the reasons why a factor has (or does not have) legal significance in terms of the law's purposes and provide a rationale for the proposed decision. The reasons, in effect, are drawn from CATO's Factor Hierarchy. In its arguments emphasizing a distinction, CATO also produces something similar to Brewer's disanalogy-warranting rationales.

CATO's distinguishing move is comparable to the first of Brewer's modes of distinguishing, adding exclusionary conditions to an AWR to underscore the disanalogy between source and target. It creates a DWR of the form "if anything that has F and G also has H, then everything that has F and G also has H unless it also has not E". As explained above, CATO's template for distinguishing leads students, in effect, to propose a more specific version of the analogy-explaining rule offered in the analogizing move so that it no longer applies to the target. It makes the rule more specific by proposing two kinds of additional terms specifying (1) the presence of the source case's extra strengths and (2) the absence of the target problem's extra weaknesses.

Provided there are factors and a Factor Hierarchy with which to represent a distinction, CATO can also distinguish in Brewer's second sense, arguing that cases from a different line of authority do not satisfy the AWR's sufficient conditions for inferring a characteristic. In CATO's terms, "a contrary case from a different line of authority" is distinguishable, if at all, in the same way as any case cited by

an opponent. CATO uses the distinguishing move template (supplemented where appropriate with arguments emphasizing the distinctions). For instance, if there were factors and a Factor Hierarchy to represent the following distinction: "An open booth in a train does not afford the privacy of a room and thus does not present the railroad owner with a tempting opportunity for fraud and plunder of the client", CATO would be able to distinguish a case involving a steamship operator from that of a railroad operator as in Brewer's example of distinguishing (Brewer 1996, pp. 1014f).

Fully assessing whether CATO's implied analogy-explaining rules are really like AWRs raises the question, "What are these rules really for?" Given Alexander's belief in a rule model of precedent, the role of the rule of a precedent was very clear. It will be applied deductively in subsequent cases, whose courts may either follow the rule or overrule the precedent. For Brewer, too, an AWR must satisfy the entailment requirement so that it can be applied deductively. Both Brewer and Alexander[20] believe a judge is obligated to state the rule of the case from which her decision follows. Of course, Brewer's additional concern is that the AWR impart rational force to the analogy.

CATO's analogy-explaining rules are also intended to capture a kind of rational force underlying an analogy (albeit perhaps not the same kind of rational force that Brewer regards as essential). Brewer maintains that an analogy-warranting rule, which embodies the relevance criteria and satisfies the entailment requirement, serves a "basic aspirational ideal of the rule of law" (Brewer 1996, p. 1001). He argues that "the rule of law ideal norms of clarity, notice, and accountability presuppose that legal commands – including those embedded in legal analogies – are deductively applicable, and that vague norms – of the sort with which one is left if legal commands are not deductively applicable – are inconsistent with those basic values" (Brewer 1996, pp. 992f). The reason CATO employs rules in its analogizing moves is simpler, to encourage students in argument succinctly to describe the essence of the analogy. That succinctness, however, is compatible with Brewer's aspirational ideal. CATO's explanatory rules have a canonical logical form, and satisfy the entailment requirement. In its canonical form, it is a deductively applicable rule structure (subject to the discussion below of clause [b], the "even though clause"). It can serve as a premise, which when applied to the target case (or the source case) deductively entails the conclusion that the target case (source case) has the desired characteristic.

4.1.2. *Conflict resolution*

In terms of conflict resolution, CATO's analogy-explaining rules may do a better job than Brewer's AWRs of specifying "in what [a source case's] exemplariness consists". A conflict-resolution move involves citing a source case that resolved some of the same conflicting strengths and weaknesses as in the target. In a conflict-resolution move, CATO's rule antecedents are not of the same exact form as an AWR. As the conflict-resolution move template (Figure 8) indicates, CATO's rule

contains an "even though" clause. The clause is intended to capture the source case's significance as an example of a court's having resolved the same conflicting strengths and weaknesses as presented in the target problem. That would seem to be an important part of the analogy's rational force.

Given Brewer's emphasis on deductively applicable analogy-warranting rules, the question arises whether and how to try to accommodate the "even though" clause of CATO's analogy-explaining rule in Brewer's framework. On the one hand, his AWRs seem to be too general. Consider what happens if the rule's antecedent mentions only the shared plaintiff's strengths. Assuming that the "If" in Brewer's AWRs and the "Where" in a CATO explanatory rule invoke the so-called "material conditional" interpretation of "If", the following is a valid inference:

If A then C; so if A&B then C.

This is an example of left monotonicity. From the first inference (If A then C), any more specific inference of the form "A&B then C" is a valid inference. As a result, from "If ⟨shared plaintiffs strengths⟩ then plaintiff should win", one may validly infer "If ⟨shared plaintiff's strengths⟩ and ⟨any plaintiff's weakness⟩ then plaintiff should win". In other words, the conclusion follows no matter what weaknesses may hurt the plaintiff's position in the target problem. That seems to be too general a rule given the source case's facts and raises the problem of the defeasibility of legal rules.

As discussed in Section 2.1, Brewer recognizes that legal rules are defeasible, and that the legal system needs techniques to deal with the phenomenon of defeasibility. CATO's conflict-resolution move and the inclusion in its explanatory rule of the "even though" clause, Clause [b], is an attempt to make its rule less defeasible. Given the source case's decision and facts, clause [b] flags those weaknesses, which the strengths have already been held to overcome.

How best to achieve that goal is the question. It would be odd to translate CATO's analogy-explaining rule into: "If ⟨shared plaintiff's strengths⟩ and ⟨shared plaintiff's weaknesses⟩ then plaintiff should win claim". Such a rule is less general. It more specifically describes the shared circumstances by virtue of its additional conditions. The shared plaintiff's weaknesses, however, are reasons why plaintiff should not win the claim. This rule seems too specific. And in any event, the "and" fails to capture adequately the significance of the shared strengths' having overcome the shared weaknesses in the source case. It treats the weaknesses like any other condition, indeed, like the strengths. In both formulations, moreover, the conditionals (i.e., "Where" in the CATO rule and "If" in the AWR) seem not to capture the positive relation between the strengths and the conclusion.

Prakken and Sartor propose a formalism for representing the resolution of the conflicting factors in a case (Prakken and Sartor 1997). Their logical formalism includes strict rules and defeasible rules, represented with a $\Rightarrow$. "Strict rules are beyond debate; only defeasible rules can make an argument subject to defeat". They say, "Since we want to read a rule *antecedent* $\Rightarrow$ *consequent* as saying that

the antecedent is a reason for the consequent, we cannot express a dominance of, say, the pro factors over the con factors by conjoining pro and con reasons in the antecedent of a rule with the pro consequent". Instead, they represent the "rule" of a case as three rules. For instance, the rule of the source case, *Forest Laboratories*, in Figure 8, might be represented:
r1: IF F6 F15 F21 $\Rightarrow$ *P* (conclusion for Plaintiff)
r2: IF F1 $\Rightarrow$ not *P*
r3: *antecedent* $\Rightarrow$ r1 defeats r2
Here *antecedent* expresses the conditions under which F6, F15, and F21 outweigh F1. The *antecedent* resembles Brewer's analogy-warranting rationale; it may be derived from other rules such as legal principles and policies as applied in the circumstances of the case. The authors also formalize a dialectical context in which participants can invoke rules of cases, like the above, in a turn-taking argument game. The formalism has been extended to construct theories for explaining a decision (Bench-Capon and Sartor 2001). The formalization appears not to have been actually implemented in a computer program, however.

4.1.3. *Relevance*

CATO's computational model of relevance incorporates some important features of Brewer's and Sunstein's theories of relevance. As discussed in Section 2.3, the latter focus on principles as informing relevance. For Sunstein, low or intermediate level principles supply the relevance criteria. For Brewer, it is the analogy-warranting rules, which Brewer sees as closely connected to principles or norms (Brewer 1996, p. 979).

CATO's model of relevance employs shared factors, which it fashions into explanatory rules very like analogy-warranting rules. Since factors are stereotypical factual strengths and weaknesses, they are more abstract than mere facts. While CATO's computational model of relevance does not involve principles per se, it does extend in the direction of incorporating more abstract legal concepts. CATO generates more abstract rationales like AWR invoking higher-level normative concepts (i.e., the legal issues associated with a claim and the laws' purposes they implement.) Its Factor Hierarchy models a connection between factors and these more abstract concepts.

The connections in the Factor Hierarchy are somewhat analogous to those between facts and principles. For instance, a given factor may be connected to more than one abstract concept. Factors may support a conclusion that a normative concept applies and is satisfied or may support the contrary conclusion. In any given case, factors may favor conflicting sides and so may the conclusions concerning the applicable normative concepts. An interesting feature of the teleological model in Berman and Hafner (1993, p. 56) is their recognition that prior cases can be found to be "relevant if they involved the same pairs of competing purposes, even if they shared few (or even no) shared features". The Factor Hierarchy makes a start at computationally implementing this extended sense of relevance.

Extending CATO's model of relevance to include even more abstract principles is an as yet unrealized goal. I am still exploring methods to represent abstract legal principles computationally and computational mechanisms to reason with them. The Factor Hierarchy and algorithms for applying it in interpreting cases is a first step toward a more robust model of the connections between more abstract norms or principles and lower-level facts. A number of techniques for representing ethical principles and their connections to concrete facts in the domain of engineering ethics are presented in Ashley and McLaren (1995), McLaren (1999).

4.1.4. *Cognitive processes of formulating AWR*

As discussed above, Brewer and Sunstein both emphasize the importance of the cognitive processes of formulating and testing the rule, which epitomizes an analogy, and the role of case comparisons in that process.

CATO's model and instruction also emphasize a cognitive process of formulating and testing rules, which captures some important aspects of the process Brewer and Sunstein describe.

First, CATO's Workbook teaches and CATO technologically supports a process of testing and modifying hypotheses about a domain of law. The hypothesis-testing is similar to the kind of testing of legal hypotheses illustrated in Sunstein's example, albeit simpler. The four-step hypothesis-testing process is similar to the process of reflective adjustment as described by Brewer. The hypotheses, framed in terms of factors, are closely related to analogy-explaining rules and to AWRs. The techniques for retrieving cases for testing a hypothesis and for interpreting whether the retrieved cases are consistent with the hypothesis are closely related to those for preparing an argument.

Second, in making multi-case, issue-based arguments, CATO in effect performs a kind of reflective adjustment. The process CATO performs and teaches involves generating more abstract, rule-like summaries of multiple cases, seeking conflict-resolution and "not fatal" cases, and citing and distinguishing counterexamples.

Third, the analogizing/resolve-conflicts move focuses students on comparing cases to abduce or invent an analogy-explaining rule in which the source case helps resolve conflicts among factors. In a complementary way, the distinguishing move focuses students on considering disanalogies, and, in effect, modifying the explanatory rule. The process of formulation is a straightforward iteration of the target problem's and source case's shared factual strengths and weaknesses, guided by the analogizing (or distinguishing) move recipe. Still, the simplification makes a complex process clearer to students.

Finally, Sunstein's and Brewer's cognitive process explicitly involves a kind of "conceptual ascent", a kind of bottom up reasoning of a sort similar to the lower levels of reflective equilibrium. CATO's Factor Hierarchy implements some aspects of this kind of reasoning. Emphasizing and downplaying moves focus students on more abstractly characterizing a case to rationalize or critique the importance of distinctions.

4.2. COMPARING CATO TO A RESULT MODEL OF LEGAL PRECEDENT

While accounting for important features of Brewer's and Sunstein's theories of analogical legal reasoning, CATO's model also avoids some of the problems identified by their most vociferous critic. In particular, CATO eschews schemes for assigning weights to facts or principles, a problem Larry Alexander associates with result models of legal precedent.

4.2.1. *Critique of result models of legal precedent*

Ironically, the jurisprudential account of analogical legal reasoning most similar to CATO's model and its use of factors is Alexander's account of "result models of legal precedent" in a paper on which his critique of analogical legal reasoning's supposed bad beginnings is based (Alexander 1989). Alexander classifies a family of models of legal precedent as an "*a fortiori* case formulation of the pure result model of precedent" (Alexander 1989, pp. 29–30). The irony is that Alexander identifies them only to criticize them. It appears, however, that CATO avoids some of the problems about which Alexander complains. According to Alexander, result models focus on the comparative strengths of target problem and source cases:

> To follow precedent, a constrained court must decide its case for the party analogous to the winner in the precedent case if the constrained case is as strong or stronger a case for that result than the precedent case was for its result .... Conversely, however, the constrained court may depart from the precedent court's result if the constrained case is a weaker case for that result than was the precedent case, even when the stated rule of the precedent covers the constrained case and demands a similar result (Alexander 1989, pp. 29–30).

According to Alexander, a result model of precedent attempts to capture a sense of how precedents operate. In particular, it provides "a way of giving meaning to the injunction 'treat like cases alike"'.

> [It] translates into 'reach the same result as in the precedent case in any case that is as or more morally compelling for the result reached in the precedent case than was the precedent case itself, even if that result is not on balance compelling in either case'. In other words, a case is 'like' the precedent case if the facts point at least as strongly toward a decision analogous to the decision in the precedent case. (Alexander 1989, p. 30)

Like a result model, CATO's model of factors attempts to capture a sense of how precedents operate. Just as a result model's *a fortiori* criteria implement a sense of "treat[ing] like cases alike", so do factors in CATO and their forebears, HYPO's Dimensions. Factors define criteria of relevant similarity and difference in terms of a pragmatic sense of when one case presents factually stronger or weaker reasons for a decision than another.

CATO's model, of course, is not a model of precedent, per se. Its purpose is to model case-based legal argument. Also, as discussed below, CATO eschews weighting schemes and does not assign weights, natural or otherwise.

As a model of case-based argument, however, CATO does make assumptions about when it is plausible to invoke a precedent as a justification, and when it is reasonable to say that a target case is factually as strong as or stronger than a source precedent.

In the CATO model, a source precedent may plausibly be invoked as a justification for deciding the target problem (i.e., the precedent is citable) when it shares at least one factor with the target and that shared factor favors the precedent's result. A target problem can be as strong as, or stronger than, a citable precedent if and only if there are no distinctions. If there are no distinctions, a target problem is stronger for the plaintiff than the source precedent to the extent that the target also has unshared pro-plaintiff factors, or the precedent also has unshared pro-defendant factors.

CATO's model is weaker than the result model Alexander describes in at least two senses. First, in CATO a side may plausibly invoke a citable precedent, even though the target problem is weaker than the precedent. However, in responding to that argument, the opponent will point out the distinctions.[21] Second, as a consequence of the fairly restrictive assumptions in the CATO model described above, it often occurs that one cannot tell whether a target problem is stronger or weaker than a source precedent. The two cases' unshared factors make their relative strengths impossible to assess.

Since CATO's case evaluation criteria prefer more on-point, less distinguishable cases, the net effect, however, is somewhat similar to Alexander's description of the result model. In other words, CATO's model is similar to a result model in the following sense. In CATO, a target problem is 'like' the source case to the extent that the facts (literally, the factors) point as strongly or more strongly toward a decision analogous to the decision in the precedent case.

The good news about result models, according to Alexander, is that such jurisprudential luminaries as Edward Levi, Steven Burton, Joseph Raz, and perhaps even Ronald Dworkin have endorsed a pure result model of precedent "in one form or another" (Alexander 1989, p. 29).

The bad news, according to Alexander, is that such models in all their formulations "face either indeterminacy or incoherence" and lack "any normative appeal whatsoever" (Alexander 1989, p. 44). Alexander criticizes result models of precedent because they do not allow one correctly to determine if a wrongly decided precedent controls a factually distinct case (Alexander 1989, p. 33). He maintains that the proposed methodologies for doing so (i.e., weighting schemes and reflective equilibrium) careen between indeterminacy and incoherence.

Alexander reports that one possible methodology to determine the effect of a wrongly decided precedent employs a scheme of weights (Alexander 1989, p. 35). To summarize the point of his example, shown in Figure 14, if facts A, B, and

> Assume the precedent case was composed of facts A, B, and C on the plaintiff's side and facts X, Y, and Z on the defendant's side. A, B, and C, correctly weighted, outweigh X, Y, and Z. Therefore, the precedent case should have been decided for the plaintiff. It was, however, mistakenly decided for the defendant. In the constrained case, there are again facts A, B, and C for the plaintiff. For the defendant, however, there are facts W, Y, and Z, plus fact N, the natural weight of the precedent decision (reliance and perhaps equality). If fact W is a stronger reason for defendant than fact X, then facts A, B, and C might still outweigh facts W, Y, A, and N, but by less than they outweighed X, Y, and Z alone. Because the balance tips in favor of plaintiff by less in the constrained case than in the precedent case, the constrained case is an *a fortiori* case for the defendant. (Alexander 1989, p. 34).

Figure 14. Alexander's example of result model with weights.

C outweigh facts W, Y, A, and N, the target problem (i.e., the constrained case) "should" have been decided for the plaintiff. By virtue of the incorrectly decided precedent, however, the target must be decided (incorrectly in terms of the weights) for the defendant, just as the precedent was.

The normative weights of facts that inform the correct decisions to these incorrectly-decided cases, Alexander assumes for purposes of argument, come from underlying principles (Alexander 1989, pp. 34–37). If that is true, however, he complains that the weighting schemes are unworkable whether the target and source cases involve only a single principle or multiple incommensurable principles.

According to Alexander, "If only a single principle is involved in both cases, or if one principle is lexically prior to the others, . . . any case that damages that principle less than the [wrongly-decided] precedent case did is an *a fortiori* case . . . " (Alexander 1989, pp. 36f).

For a variety of reasons, however, Alexander rejects any scheme of common weights, one, which assumes, "a single master principle that assigns weights in a common currency to various facts".

> Facts that help the plaintiff or the defendant might represent distinct principles or policies that a decision for a particular party will further. A decision in the constrained case that parallels the decision in the precedent case may further some principles or policies to a greater extent than in the precedent case and others to a lesser extent. Under such circumstances, therefore, what is an *a fortiori* case? (Alexander 1989, p. 35).

Furthermore, Alexander argues, if there are multiple, incommensurable principles, it is not clear how one resolves whether a case is controlled by the precedent. "[I]f the principles at stake are multiple and are not lexically ordered or reducible to a common master principle, determining whether the constrained case is an *a fortiori* case is impossible". (The only exception would be if the constrained case is an *a fortiori* case under every principle.) (Alexander 1989, pp. 36f).

Beyond rejecting weighting schemes, Alexander criticizes reflective equilibrium as a doomed attempt to deal with the same problem. In Alexander's view, reflective equilibrium founders from indeterminacy and incoherence. Not only are

there multiple, incommensurable principles, but some of those principles are based on the inevitable erroneously decided precedents (Alexander 1989, pp. 31–32, 37–39).

4.2.2. *CATO, weighting schemes, and the problem of incommensurable principles*

In the design of CATO (and HYPO before that), I have consciously eschewed weighting schemes for reasons related to Alexander's concerns. CATO has no means, using weights or otherwise, of combining conflicting strengths and weaknesses quantitatively into a decision. Instead, it generates alternative arguments about how the conflicting strengths and weaknesses should be resolved.

Any legal weighting scheme must be sensitive to the factual context. Attorneys may agree that one factor or principle is usually more important than another, but they are always mindful that in some factual contexts the opposite may be true. Premature commitment to one scheme of weights might cut off potentially fruitful lines of inquiry and argument. Some recent work in AI has focused on developing context-sensitive weighting schemes, although not in a legal context and not involving argumentation (Aha 1998).

Any legal weighting scheme must also be authoritative. Judges and other authorities, however, do not speak as though they were employing numerical weights in resolving conflicting strengths and justifying decisions. In general, it would not be appropriate to make statistical arguments as to the relative weights of strengths and weaknesses. Judges do not regard such arguments as persuasive, probably in part because the body of litigated cases decided by courts is such a biased sample (omitting, for instance, the vast majority of disputes which are settled.)

Instead, judges expect adversaries to make non-numerical arguments about how to resolve conflicts. CATO's conflict-resolution moves exemplify such arguments. The competing sides draw qualitative analogies to source cases that resolve the conflicting strengths and weaknesses (Ashley 1992).

Similarly, schemes for weighting ethical principles are impracticable. In practical ethics, "the metaphor of the 'weight' of a principle of duty has not proven amenable to precise analysis" (Beauchamp and McCullough 1984, p. 16). What would the weight of a principle, represented by a number, mean and how would one compute it? If represented by some hierarchy of principles (or lexical ordering as Alexander suggests and rejects), the assignment of weights would not be sufficiently sensitive to factual context. "Principle A may usually be more important than principle B, but there are always some sets of circumstances in which Principle B should win out" (Ashley 1992, p. 198).

As a result of this design choice, CATO does not assign any weights to precedents or their facts nor does it combine weights. When confronted with an example like Alexander's in Figure 14, the CATO model would not be able to represent such quantitative assertions as: "[F]act W is a stronger reason for defendant than fact X" or that some set of facts (A, B, and C) outweigh some other set of facts "by

Analogize (for Defendant): Where Y and Z, defendant should win just as it did in the source case even though A, B, and C.

Distinguish (for Plaintiff): The source case is distinguishable. X favored defendant in the source case. This is not so in the target problem.

Downplay (for Defendant): Lack of X in the target problem is not an important distinction. W is in the target case and W is a strong compensating reason for defendant because of Legal Concern 1 (which X and W share).

Emphasize (for Plaintiff): Lack of X in the target problem is an important distinction. X favored defendant in the context of X, Y, Z, A, B, C because of Legal Concern 2. W may help defendant in the target problem but not because of Legal Concern 2. Legal Concern 2 is not an issue in the target problem.

Figure 15. CATO-style argument exchange for the example in Figure 14.

less than they outweighed X, Y, and Z alone" or that "the balance tips in favor of plaintiff by less in the constrained case than in the precedent case".

Where the target problem's strengths and weaknesses conflict, CATO will not combine them quantitatively. Instead, CATO's arguments elaborate reasons for and against a decision. If the conflicting factors involve common reasons, CATO will assert that the strengths compensate for the weaknesses. In addition, if a source case involving a similar conflict is available, CATO will make an argument that the conflict should be resolved in the same way as in this precedent. CATO responds to this argument by distinguishing, and then it may downplay and emphasize distinctions.

In terms of CATO's model, the case-based argument in Alexander's example in Figure 14 is plausible and could be rendered as shown in Figure 15. Opening with an analogizing move, the defendant would argue that the target case should be decided for defendant, as was the source case. CATO (with some modest extensions) could also generate a plausible series of responses and replies. In the downplaying move, defendant asserts W should compensate for the lack of X in the target problem. In the emphasizing move, plaintiff argues that W does not compensate for the lack of X because each feature is significant for different reasons, and the reason underlying X is not important in the context of the target problem. (Such an argument depends, of course, on the particular connections in the Factor Hierarchy. The downplaying and emphasizing moves assume a certain configuration of connections among W, X, Y, Z, A, B, and C. Arguments about the meaning of a factor are possible only where the given factor may be linked to more than one abstract factor. For instance, if Legal Concerns 1 and 2 are the same, the attempt at abstract distinguishing will not succeed.)

In other words, even without a weighting scheme, a Factor Hierarchy can help generate plausible arguments and responses in examples like Alexander's. The argument citing the precedent can be seen as raising a kind of presumption. Someone arguing for a conclusion opposite to the one in the precedent has the burden of meeting the presumption. The subsequent comparison of the target and source case in the distinguishing, downplaying, and emphasizing moves drives a kind of "conceptual ascent", to use Sunstein's phrase, in analyzing the conflicting strengths and

weaknesses. Like the result models Alexander describes, CATO's design assumes that the significance of these factual features is tied to more abstract normative concepts. In the CATO model these more abstract concepts are normative concerns associated with a legal claim, which like principles, are very much more abstract than the facts.

Unlike reflective equilibrium, of course, the CATO model stops well short of attempting to resolve such arguments and should escape Alexander's further criticism. In the next section, however, I argue that CATO may contribute to jurisprudential efforts to explore phenomena like reflective equilibrium.

5. Conclusions

In designing case-based computational models of legal argument, AI and Law researchers have much to learn from jurisprudential models of analogical legal reasoning. Sunstein exhorts designers to build systems that make "the good normative arguments that underlie assessments of analogousness" (Sunstein 1993, p. 774, n.116.) His example of the conceptual work legal practitioners perform with analogical reasoning, discussed in Section 2.2, illustrates the goal: to design systems that can explore a legal issue in the context of a specific problem, flexibly interrelate discussions of the facts of the problem and relevant precedents with low-level legal principles, generate hypotheses to explain proposed decisions in light of the precedents and principles, test the hypotheses against the past cases and hypothetical scenarios, modify the inquiry accordingly, and fashion the resulting analysis into persuasive arguments. In short, the goal is to create systems that can engage in robust reflective adjustment and even that can work toward achieving a kind of reflective equilibrium.

AI and Law models of case-based legal argument such as Taxman II (McCarty and Sridharan 1981), GREBE (Branting 1991), BankXX (Rissland et al. 1996), the dialogue game in Prakken and Sartor (1997) with extensions for theory construction in (Bench-Capon and Sartor 2001), HYPO and CATO have all modeled aspects of reflective adjustment. Taxman II used concept-defining prototypes and associated deformations to align a proposed decision of an early Supreme Court tax case with the alternate groups of conflicting precedents. GREBE assessed relevance of precedents in terms of how completely the program could match the explanations of past cases to the problem's facts. BankXX used argument-strengthening factors to guide the program's search through a network of interrelated precedents and legal theories for the best cases and arguments. Theory construction in a dialogue game promises a way to relate conflicting case facts to alternative principles, values, and theories of a case.

Don Berman's and Carole Hafner's challenge to design case-based systems that explain "why a factor favors the plaintiff or defendant, and why that factor is considered legally relevant" is a prerequisite for achieving "good normative arguments". Part of that challenge involves being able to choose legal reasons in a

way that is sensitive both to the side and positions one represents and to the specific facts of a case and the opportunities they afford. Just as a law student in writing an examination question, or a judge in writing an opinion, must take account of how the problem's facts affect the strength of a claim, the applicability of precedents, and the applicability and force of legal principles, so must the case-based model of legal argument.

CATO's significance lies in this regard. It implements an algorithmic model of the interpretive re-characterization of source cases in terms of abstract legal reasons, one that is sensitive to the argument side, the problem facts, and those of the precedents. It generates multiple interpretations of the significance of differences among similar cases. Its algorithms for downplaying and emphasizing distinctions guide a kind of "conceptual ascent" as CATO decides which paths upward through the Factor Hierarchy to pursue and how high to go in selecting the legal reasons why factors matter in the context of the particular problem and argument. This interpretive capability, its analogy-explaining rules, issue-based arguments, and ability to help students learn to pose and test legal hypotheses are closely related to the tasks of abducing, testing, and modifying rules with which to epitomize a legal analogy. It can perform these tasks for any of the 150 cases in its database.

CATO's arguments still fall a good deal short of being good normative arguments. It does not include a representation of legal principles, nor do its algorithms relate such principles to its analysis of problems. In this respect, one must look to the jurisprudential sources for guidance as to exactly what role normative principles play in legal reasoning. Although I must keep reading, my sense is that even the jurisprudential sources are not entirely clear on that point.

In defending their models against a reductionist critique, for instance, Brewer and Sunstein both appealed to complex processes such as abduction of rules and principles from a comparison of cases and confirmation of the abduced rules and principles in a process of reflective adjustment shading upward into reflective equilibrium. Judges, it is said, (and other systematic decision-makers such as practical ethicists and policy-makers) take pains to reconcile a proposed decision horizontally with prior decisions (i.e., precedents) and vertically with general principles. They may concern themselves with synthesizing only a handful of "near-by" cases (i.e., local consistency) or with all precedents in this and "near-by" domains (i.e., global consistency).

Complex processes like these, with both cognitive and normative components, are difficult to describe. Scholars invoke intuitively suggestive descriptions like "conceptual ascent", "reflective adjustment", "reflective equilibrium", "fitting morally acceptable principles", "principled consistency with respect to individual cases and low-level principles", "local" versus "global" consistency. But even the standards of consistency and coherence appear to be more intuitively evocative than well defined in the context of such complex processes. Scholars also illustrate the processes with textually described extended examples like Sunstein's. The examples, while illuminating, only begin to illustrate the processes.

The formulations of the processes, moreover, remain controversial. Scholars dispute the nature and degree by which the processes are constrained. Alexander, Brewer, and Sunstein all take different views of the possible degrees of freedom involved: whether one must treat a precedent as a fixed point or as something to be rejected if mistaken, whether one must either follow a precedent's rule or overrule it, or whether one may follow a precedent's decision while rejecting its rule, or even whether one may distinguish a case. Their different models, moreover, have very different implications for the kinds of arguments it is reasonable to make.

Conceivably, AI models provide a means for exploring more systematically and explicitly how these processes may work. Don Berman prized AI and Law models as a valuable pedagogical tool for teaching law students about legal reasoning; the models illuminate the assumptions, limitations, and uncertainties inherent in our conceptualizations of legal reasoning. The same virtues make them good investigative tools for jurisprudential analysis, especially into complex cognitive processes. Computational models of case-based legal argument can exemplify certain aspects of these processes more robustly than pencil-and-paper examples and over a wider range of conditions. One can perturb the model systematically as a way of investigating the underlying phenomenon. Where the model fails to capture important aspects of expert legal arguments, its concrete detail enables one to identify which changes need to be made. The empirical, descriptive, and reproducible methodology of AI research may help investigate systematically the cognitive aspects of such phenomena as the dialectical interaction between abducing legal rules or normative principles and testing them on concrete cases. The choice of case representations, the design of a process of case comparison, and the formulation of standards of reasonable arguments all appear to play important roles in achieving a coherent account of processes like these.

It is, of course, not likely that some Herculean graduate student will soon build an AI model of reflective equilibrium. The advantages of an AI methodology have a substantial cost in terms of time and complexity. Nevertheless, one could use an AI and Law model like Taxman II, GREBE, BankXX, CATO, or an implementation of a dialogue game, to begin a systematic cognitive task analysis of reflective adjustment. It would be instructive, for instance, to collect protocols of experienced instructors and practicing attorneys using CATO to fashion a multi-case, multi-issue argument with a target problem and source cases selected from its database. By carefully choosing cases or constructing hypotheticals, one could focus them on fact situations involving conflicting factors and principles. One might examine how they edit (or teach others to create or edit) such arguments. A program's data collection techniques would help the investigation. CATO, for instance, can log subjects' commands and help record think-aloud protocols of subjects' problem-solving and argument-making activity. By analyzing the protocols, one might learn how to improve the computational model's algorithms so that it can more effectively synthesize coherent accounts of cases and principles in arguments, and present better

examples of synthesis for teaching students a process of reflective adjustment. The algorithms would frame an iterative process of improvements involving:

- Crafting a decision rule or AWR more explicitly,
- Achieving flexibility in selecting the level(s) of abstraction for expressing the decision rule and implementing criteria for making this choice intelligently,
- Achieving greater flexibility in selecting issues and cases to emphasize, perhaps by:
- Identifying the legal principles at issue and using them to characterize the target problem abstractly,
- Improving criteria for selecting and synthesizing the multiple best source cases to emphasize, and
- Developing techniques for integrating more detailed factual descriptions from the source cases and relating them to the applicable principles.

An AI and Law approach, for instance, may help concretize a model of coherence. Recently, Horty has proposed a deontic and non-monotonic logical reformulation of HYPO's account of case-based legal argument, one which clearly defines the concept of a coherent interpretation of a case base as applied to a problem situation (Horty 1999). It would be interesting to investigate this formulation in the context of CATO's case base, or even the set of cases that underlie a casebook writers' treatment of a particular legal issue. In the context of a casebook chapter, for instance, one might investigate how the author's choices of cases and order of presentation support students in taking coherence into account as a goal. How do legal writing instructors teach students a process of synthesizing even a small group of cases into a coherent account in a brief? When a precedent is inconsistent with an account, how do they help students address the problem of deciding whether the account needs to be modified or the precedent was wrongly decided?

Computational models may also help to identify cases, which are arguably erroneously decided. As discussed in Section 2.4, Alexander's critique of reflective equilibrium challenges its proponents to propose a theory of error. Sunstein recognized that inevitably, some cases will have been wrongly decided. Lawyers routinely make arguments that a prior case was wrongly decided in connection with fitting a proposed decision into a coherent theory of the decided cases and relevant principles. Computational models may be of some assistance in developing a theory of error. For instance, they may help to conceptualize an area of law as a multi-dimensional "space" of problems. Actual cases can be assigned locations in that space, and algorithms may be applied to cluster those cases into categories. The clustering may help to identify decisions that are exceptions or outliers, information that can assist in determining whether such decisions have been wrongly decided.

It will be a long road, but Don Berman's challenge points the way.

Notes

[1] Sunstein describes John Rawl's notion of reflective equilibrium as reasoning which "entails an effort to produce both general theories and ... considered judgments about particular disputes .... [T]he particular views are adjusted to conform to the general theory and vice-versa. Through this process, we hope finally to reach a form of equilibrium" (Sunstein 1993, p. 751).

[2] This DWR is not in Brewer's example.

[3] Although (Aleven 1997, p. 19) lists eight argument moves, I prefer to consolidate into one move the two counterexample moves, citing a more-on-point counterexample and an as-on-point counterexample.

[4] The role of an AI knowledge representation is to capture selected aspects of a complex reality faithfully enough to enable a computer program to manipulate its representations in a meaningful way. Jurisprudential scholars may encounter problems of knowledge representation when they formalize or schematize legal knowledge. For instance, in order to present a "paradigmatic exemplary argument", Brewer offers an "interpretive reconstruction of [the] structure of [the judge's reasoning]" in the case of the innkeeper (Brewer 1996, p. 1004.) Brewer calls the Judge's explanation of the distinction between an inn and a sleeping car "complex and somewhat obscure" (Brewer 1996, pp. 1013–1014). In selecting those of the Judge's reasons, which are sufficiently factual and clear to represent in his paradigmatic example of the form of distinguishing, Brewer makes knowledge representation choices.

[5] In addition to factors (and dimensions), AI and Law programs have employed a wide range of knowledge representation techniques: Logical rules (Sergot et al. 1986), heuristic rules (Waterman and Peterson 1981), prototypes and deformations (i.e., template-like descriptions of legal concepts such as taxable income and a set of possible mappings from one description into other possible ones) (McCarty and Sridharan 1981; McCarty 1977, 1989), Augmented Transition Network (ATNs) (i.e., a kind of grammar of legal of rules for "parsing" events having to do, for instance, with offer and acceptance (Gardner 1987; Yoshino et al. 1998) and property settlements (Zeleznikow et al. 1995–1996)), semantic networks (McCarty 1977; Branting 1991), and connectionist networks (i.e., a system of many nodes connected to other nodes by weighted links, where the weights may be adjusted pursuant to a training rule) (Rose and Belew 1991; Zeleznikow et al., 1995–1996).

[6] Factors are based on dimensions, a knowledge representation device I invented for use in my Ph.D. dissertation program, HYPO, a program that also performed case-based legal reasoning in the domain of trade secret misappropriation law (Ashley 1990). As compared to factors each dimension contained additional structure including (1) tests for deciding if a dimension applies to a case or is a near-miss and (2) a focal slot for specifying the magnitude of a dimension in a case (see *infra*, note 7.) HYPO also had heuristics based on dimensions for posing meaningful hypothetical variations of target problems in order to strengthen or weaken the arguments for or against plaintiff's claim. CATO cannot pose hypotheticals.

[7] CATO's factors are boolean; they either apply or they do not apply. By contrast, dimensions had different possible ranges, designed to represent the magnitude of a dimension in a given case. A dimension's magnitude was not a measure of its weight. Instead, it indicated how extreme an example of the dimension the case was. HYPO's heuristics also enabled it to pose hypothetical variations of magnitudes to demonstrate a kind of slippery slope (Ashley 1990, pp. 148–154).

[8] RESTATEMENT (FIRST) OF TORTS, 757 cmt. b (1939).

[9] RESTATEMENT (FIRST) OF TORTS, 757 (1939).

[10] Uniform Trade Secrets Act, 14 U.L.A. 438 (1990).

[11] Forest Laboratories, Inc. v. Formulations, Inc., 299 F. Supp. 202 (E. D. Wis. 1969), *rev'd in part*, 452 F.2d 621 (7th Cir. 1971).

[12] By substituting "defendant" for "plaintiff" and vice versa, the templates may be modified to accommodate instances of defendant's analogizing and plaintiff's distinguishing.

[13] Television Telecommunication Systems, Inc. v. Saindon, 522 N.E.2d 1359 (Ill. App. 2d Dist. 1988).

[14] Although we hoped that students would recognize the connection between the hypotheses they tested and modified earlier in the curriculum and these template-guided succinct expressions of analogies and distinctions, the instruction did not use the terminology of logical rules, which, we feared, would tend to confuse first year students.

[15] Anaconda Co. v. Metric Tool & Die Co., 485 F.Supp. 410 (E.D. Pa., 1980).

[16] Mason v. Jack Daniel Distillery, 518 So.2d 130 (Ala. Civ. App. 1987).

[17] Anaconda Co. v. Metric Tool & Die Co., 485 F. Supp. 410 (E.D. Pa., 1980).

[18] When I last saw Don Berman at a workshop at Chicago-Kent School of Law, he told me that he used the Factor Hierarchy in his Intellectual Property class to help teach students about trade secret law.

[19] We began to design CATO in 1991; its curriculum did not take shape in its current form until 1995. The experimental evaluation of CATO took place in February, 1996. Neither CATO's model nor curriculum were designed with Sunstein's 1993 or Brewer's 1996 models of analogical legal reasoning in mind, nor with Alexander's 1996 critique. The principle influence on (and inspiration for) CATO was the account of reasoning with cases in Levi's 1949 classic, *An Introduction to Legal Reasoning.* In connection with designing the curriculum, we consulted Llewellyn's *The Bramble Bush* (1986, 1960) and later Burton's *An Introduction to Law and Legal Reasoning* (1995).

[20] Alexander allows that the rule may be implied from the "material facts"' of the case as long as it does not depend on what any other courts do (Alexander 1989, pp. 18f). Karl Branting has argued after Cross (1979), that "any adequate model [jurisprudential or computational] of *ratio decidendi* must include every warrant necessary for the justification of a precedent" (Branting 1994). Since Branting maintains "that the ratio decidendi of a case is determined, at least in part, by subsequent cases"' (Branting 1994, p. 7), his model would not qualify as a rule model of precedent under Alexander's criteria. CATO offers warrants drawn from the Factor Hierarchy, for instance, in downplaying or emphasizing distinctions.

[21] For example, in HYPO, for a target problem with but one disclosure to outsiders (see Factor F10, Secrets-Disclosed-Outsiders (d)), defendant could still cite a pro-defendant source case where disclosures had been made to a thousand outsiders. In distinguishing, however, plaintiff would point out how much stronger for the defendant the source case is along the dimension than in the target problem.

References

Aha, D. W. (1998) Feature Weighting for Lazy Learning Algorithms. In Liu, H. and Motoda, H. (eds.) Feature Extraction, Construction and Selection: A Data Mining Perspective, Kluwer.

Aleven, V. (1996). CATO *Workbooks: Workbook I* <http://adson.lrdc.pitt.edu/wb1.html>, *Workbook I, Reference Pages*, Section 7 <http://adson.lrdc.pitt.edu/ wb1.html# theorytesting>. *Workbook II* <http://adson.lrdc.pitt.edu/wb2.html>, *Workbook 2*, Section 3 <http://adson.lrdc.pitt.edu/wb2.html#importanceofdistinctions>. *Workbook III* <http://adson.lrdc.pitt.edu/wb3.html>, *Workbook IV* <http://adson.lrdc.pitt.edu/wb4.html>.

Aleven, V. (1997). Teaching Case-Based Argumentation Through a Model and Examples, Ph.D. dissertation, University of Pittsburgh Graduate Program in Intelligent Systems, (unnumbered technical report, Learning Research and Development Center).

Aleven, V. and Ashley, K. (1996). How Different is Different? Arguing about the Significance of Similarities and Differences. In Smith, I. and Faltings, B. (eds.) Proceedings 3d European Workshop on Case-Based Reasoning (EWCBR-96), 1–15. Springer Lecture Notes in Artificial Intelligence Series No. 1168.

Aleven, V. and Ashley, K. (1997). Teaching Case-Based Argumentation Through a Model and Examples: Empirical Evaluation of an Intelligent Learning Environment. In du Boulay, B. and Mizoguchi, R. (eds) Artificial Intelligence in Education, Proceedings AI-Ed 97 World Conference on Artificial Intelligence in Education, 87–94. (August) IOS Press.

Alexander, L. (1989). Constrained By Precedent, Southern California Law Review 63: 1.

Alexander, L. (1996). Bad Beginnings, University of Pennsylvania Law Review 145: 57–86.

Allen, L. E. and Saxon, C. S. (1987). Some Problems in Designing Expert Systems to Aid Legal Reasoning. In Proceedings of First International Conference on Artificial Intelligence and Law, 94–103. ACM Press.

Ashley, K. (1990). Modeling Legal Argument: Reasoning with Cases and Hypotheticals, The MIT Press.

Ashley, K. (1992). Case-Based Reasoning and Its Implications for Legal Expert Systems, Artificial Intelligence and Law 1: 150.

Ashley, K. and Aleven, V. (1996). Casebook Chapter on the Law of Trade Secret Misappropriation <http://adson.lrdc.pitt.edu/casebook1.html>.

Ashley, K. and Aleven, V. (1997). Reasoning Symbolically About Partially Matched Cases. In Proceedings International Joint Conference on Artificial Intelligence, IJCAI-97, 335–341. Morgan Kaufmann.

Ashley, K. and McLaren, B. (1995). Reasoning with Reasons in Case-Based Comparisons. In Veloso, M. and Aamodt, A. (eds) Proceedings of the First International Conference on Case-Based Reasoning (ICCBR-95), 133–144. Springer Lecture Notes in Artificial Intelligence Series No. 1010.

Beauchamp, T. and McCullough, L. B. (1984). Medical Ethics: The Moral Responsibilities of Physicians. Prentice-Hall.

Bench-Capon, T. and Sartor, G. (2001). Theory Based Explanation of Case Law Domains. In Proceedings 8th Int'l Conference on Artificial Intelligence and Law, 12–21. Association for Computing Machinery.

Berman, D. and Hafner, C. (1993). Representing Teleological Structure in Case-Based Legal Reasoning: The Missing Link in Proceedings of the Fourth International Conference on Artificial Intelligence and Law, 50–59. Association for Computing Machinery.

Berman, D. and Hafner, C. (1996). Obstacles to the Development of Logic-Based Models of Legal Reasoning. In Walter, C. (ed.) Computer Power and Legal Language. Greenwood Press.

Branting, L.K. (1991). Building Explanations from Rules and Structured Cases, Int'l Journal of Man-Machine Studies 34: 797–837.

Branting, K. (1994) A Computational Model of Ratio Decidendi, Artificial Intelligence and Law 2: 1–31.

Brewer, S. (1996). Exemplary Reasoning: Semantics, Pragmatics, and the Rational Force of Legal Argument by Analogy, Harvard Law Review 109: 923–1028.

Burton, S. (1995). An Introduction to Law and Legal Reasoning, 2d ed. Little, Brown.

Cross, R. (1979). Precedent in English Law, 3d ed. Oxford University Press.

Gardner, A. vdL. (1987). An Artificial Intelligence Approach to Legal Reasoning, MIT Press.

Goldstein, P. (1990). Copyright, Patent, Trademark and Related State Doctrines, 3d ed. Foundation Press.

Horty, J. F. (1999). Precedent, Deontic Logic, and Inheritance. In Proceedings, Seventh International Conference on Artificial Intelligence and Law, 63–71. Association of Computing Machinery.

Levi, E. (1949). An Introduction to Legal Reasoning, University of Chicago Press.

Llewellyn, K. (1960). The Bramble Bush. On Our Law and Its Study. Oceana Publications 1986.

McCarty, L. T. (1977).Reflections on TAXMAN: An Experiment in Artificial Intelligence and Legal Reasoning. Harvard Law Review 90: 837–893.

McCarty, L. T. (1989). A Language for Legal Discourse: I. Basic Features. In Proceedings of Second International Conference on Artificial Intelligence and Law, 180–189. Association of Computing Machinery.

McCarty, L. T. and Sridharan, N. S. (1981). The Representation of an Evolving System of Legal Concepts: II. Prototypes and Deformations. LRP-TR-11. Lab. for Computer Science Res. Rutgers U. New Brunswick, NJ.

McLaren, B. (1999). Assessing the Relevance of Cases and Principles Using Operationalization Techniques. Ph.D. dissertation, University of Pittsburgh Graduate Program in Intelligent Systems.

Milgrim, R. M. (1995). Business Organizations, Milgrim on Trade Secrets, vol. 12. Matthew Bender.

Minsky, M. L. (1968) Semantic Information Processing. MIT Press.

Prakken, H. and Sartor, G. (1997). Reasoning with Precedents in a Dialogue Game. In Proceedings of the Sixth International Conference on Artificial Intelligence and Law, 1–9. Association for Computing Machinery.

Rissland, E. L. (1990). Artificial Intelligence and Law: Stepping Stones to a Model of Legal Reasoning, Yale Law Journal 99: 1957–1981.

Rissland, E. L., Skalak, D. B. et al. (1996). BankXX: Supporting Legal Arguments Through Heuristic Retrieval, Artificial Intelligence and Law 4: 19.

Rose, D. E. and Belew, R. K. (1991). A Connectionist and Symbolic Hybrid for Improving Legal Research, International Journal of Man-Machine Studies 35: 1–33.

Schaffner, K. F. (1990). Presentation entitled "Case-Based Reasoning in Law and Ethics" at the Foundations of Bioethics Conference. Hastings Center. December.

Sergot, M. J. et al. (1986). The British Nationality Act as a Logic Program, Comm. of the ACM 29: 370–386.

Sunstein, C. R. (1993). On Analogical Reasoning, Harvard Law Review 106: 741–791.

Waterman, D. A. and Peterson, M. (1981). Models of Legal Decision-making. Technical Report R-2717-1CJ. Rand Corporation. Santa Monica, CA.

Yoshino, H. (1998). Logical Structure of Contract Law System, J. Adv. Comp. Intelligence 2: 2–11.

Zeleznikow, J. Stranieri, A. et al. (1995–1996). Project Report: Split-Up – A Legal Expert System which Determines Property Division Upon Divorce, Artificial Intelligence and Law 3: 267–275.