Common Law Reasoning

John Horty

www.umiacs.umd.edu/users/horty

*** DRAFT #1 ***

March 21, 2013

Contents

1 Overview

The goal of this paper is to advance our formal understanding of the common law—particularly the nature of the reasoning involved, but also, though to a lesser extent, its point, or justification. As so often in discussions of the modern common law, I will focus special attention on the doctrine of precedent, according to which later courts are constrained by the decisions of earlier courts, but also granted some degree of freedom to respond in creative ways to fresh circumstances. The central challenge is to understand this balance between constraint and freedom.

In order to locate the account to be developed here, it will be useful to begin with a few words about the surrounding territory. There is, of course, the position, often associated with the American legal realists, that precedential constraint is nothing but a fiction, at least from a normative perspective. On this view, the interesting questions concern, not what courts ought to decide in particular cases, or what they are constrained to decide, but only what they will decide; these questions are then to be answered using empirical methods and tools from the social sciences, rather than ideas from moral or legal theory.[1] There are also a number of coherence theories of precedential constraint, perhaps the most familiar of which is due to Ronald Dworkin. This view is complex, but very roughly, constraint is supposed to be determined in three stages: a court facing a particular case should, first, survey the legal materials bearing on that case, including previously decided cases as well as any statutory or constitutional provisions; next, the court should identify the best moral theory that is able to explain the existing pattern of decisions and legislation; constraint then results, finally, through an application of this favored theory to the particular case at hand.[2]

[1] A useful summary of legal realism is provided by Leiter (2005).

[2] The account summarized here can be found in Dworkin (1975), an early paper; it has been refined in

Although theories falling into these categories, both the nihilistic theories of the legal realists and the various coherence theories, have a good deal of interest, I will set them aside in this paper to begin, instead, with the more familiar position that the common law is based on rules—that a court facing a particular case either invokes a previous common law rule or articulates a new one to justify its decision in that case, and that this rule then constrains the decisions that might be reached in any future case to which it applies. This idea can be developed in two ways, depending on the nature of the rules involved. Some writers feel that the rules introduced by courts are best seen as defeasible, or prima facie. This line of development is one with which I have a good deal of sympathy, but again, not one I discuss here.[3] Instead, I will concentrate on the position that common law rules should be taken as strict, or exceptionless, just like the "if ... then ..." rules we study in introductory logic, where we typically learn to express these rules using the ordinary material conditional.

In fact, the idea that common law rules state exceptionless generalizations can itself be developed in two ways, depending this time, not on the meaning of the rules themselves, but on the set of conventions within which they are thought to be embedded. Some writers argue that, once a common law rule has been introduced in an earlier case, it must then govern any later case to which it applies, unless the court in the later case wishes to *overrule* the earlier decision and has the authority to do so.[4] Other writers, however, favor a more flexible interpretation, according to which, although only certain courts, depending on their place in the judicial hierarchy, have the authority to overrule earlier decisions, all courts have

various ways over the years, while still retaining, I believe, the same overall shape.

[3]The analysis of common law reasoning using defeasible rules has been developed in great detail within the field of artificial intelligence and law; see Prakken and Sartor (1998) for a seminal paper. I believe that the recent work presented in Holton (2010) and (2011) can likewise be seen as falling within this category.

[4]This view has been developed with great force by Alexander (1989), and then, especially, by Alexander and Sherwin (2001) and (2009). We will return to consider aspects of this view later in the paper.

the power of *distinguishing*—the power, that is, of identifying important differences between the facts present in certain later cases and those of earlier cases, and so modifying the rules set out in earlier cases in order to avoid inappropriate application to these later cases.

The idea that common law rules are malleable in this way—that the rules introduced by earlier courts can be modified in light of later cases—is, arguably, the most prevalent view among legal theorists, and provides what I refer to as the *standard model* of common law reasoning.[5] This model has a number of virtues, and is often thought to offer the most accurate picture of ordinary, incremental, legal development, through the gradual modification of common law rules in response to new situations.[6] Nevertheless, and in spite of its virtues, the standard model faces an immediate problem in accounting for precedential constraint. For if the constraints imposed by the decisions of earlier courts are supposed to be carried by rules, but later courts are free to modify the rules set out by earlier courts, then it is hard to see how these rules can impose any real constraints at all: how can courts be constrained by rules that they are free to modify at will?

This problem could be illustrated by tracing the development of an actual common law doctrine, but also, and perhaps more clearly, with an artificial example. Suppose, then, that Laura and Ron are the parents of two children—Emma, who has just turned nine, and Max, age twelve—and that they have agreed to respect each other's decisions concerning the children, treating these decisions, in effect, as precedents. And imagine that one night, Emma, who has completed her homework but did not finish dinner, asks Laura if she can stay up and watch TV. This is like a legal case: a situation is presented to an authority,

[5] Versions of this view have been developed by Eisenberg (1988), Levi (1949), Raz (1979), Schauer (1989) and (1991), and Simpson (1961), along with many others.

[6] See, for example, Levi (1949, pp. 8–27) for a discussion of the development, within the standard model, of common law rules characterizing certain objects as "inherently dangerous."

Laura, who must make a decision and, ideally, provide a rationale for her decision. Suppose that Laura resolves the case by granting the request, stating that Emma can stay up to watch TV since she is now nine years old. This decision can be seen as introducing a household version of a common law rule—perhaps, "Children age nine or greater can stay up and watch TV"—fashioned in response to a particular fact situation, but applicable to future situations as well.

Now imagine that, the next day, Max, who has neither completed his homework nor finished dinner, asks Ron whether he can stay up to watch TV, but in this case Ron refuses, on the grounds that Max has not completed his homework. Max might reasonably appeal this decision with the complaint, "Ah, but given the precedent established last night, in the case of Emma, our household is now governed by a rule according to which children age nine or greater can stay up and watch TV." According to the standard model of precedential constraint, however, Ron can defend his decision by distinguishing the two cases, arguing that the previous rule should not apply exactly as formulated to the new case of Max, since this case, unlike the previous case of Emma, presents the additional feature that the child in question has not completed his homework. The overall effect of Ron's decision, then, is that the set of household rules should now be seen as having been changed in two ways. It should be seen, first of all, as containing a new rule to justify Ron's decision in the case of Max—perhaps the rule, "Children who have not completed their homework cannot stay up and watch TV." And second, the rule previously set out by Laura in the case of Emma must now be taken as having been modified, in order to avoid an unwanted application in the latter case—perhaps now reading, "Children age nine or greater can stay up and watch TV, unless they have failed to complete their homework."

The literature on the standard model contains various proposals about how, exactly,

Ron's modification of Laura's previous rule might be justified—perhaps Ron feels that his modified rule provides a better representation than Laura's original formulation of what she had in mind to begin with, or that the new rule is the one Laura would have set out if only she had envisioned the new situation; or perhaps Ron simply feels that the overall regulatory system of the household is sufficiently improved by his modification of Laura's rule. Regardless of how the question of justification is to be settled, however, the fact remains that the standard model does allow Ron to avoid an unwanted application of Laura's original rule by means of a reformulation, and that is enough to generate the problem at hand: as long as Ron—or a common law court—can modify an earlier rule at will to avoid an unwanted application in a latter case, it is hard to see how either of them could be thought of as constrained by that earlier rule.

As it happens, the standard model allows for a response to this problem. The response was first set out explicitly by Joseph Raz, although, as Raz notes, it owes much to the previous work of A. W. B. Simpson.[7] The central idea is that, although later courts are indeed free to modify the rules set out by earlier courts, they are not free to modify these rules entirely at will. Any later modification of an earlier rule must satisfy two conditions: first, the modification can consist only in the addition of further restrictions, which narrow the original rule; and second, the modified rule must continue to yield the original outcome in the case in which it was introduced, as well as in any further cases in which this rule was appealed to as a justification.

The force of these *Raz/Simpson conditions* on rule modification can be illustrated by returning to our example, where Laura's initial rule, "Children age nine or greater can stay up and watch TV," introduced in the case of Emma, was later modified by Ron to read

[7]See Raz (1979, pp. 180–209) and Simpson (1961).

"Children age nine or greater can stay up and watch TV, unless they have failed to complete their homework," in order to block applicability to Max. Here, it is easy to see that Ron's modification of the rule satisfies both of the two Raz/Simpson conditions: it simply narrows Laura's original rule with a further requirement for applicability, and it yields the same result as the original in the case in which the original was introduced, that Emma can watch TV. Suppose, by contrast, that Ron had modified Emma's original rule to read, "Children who are female can stay up and watch TV." Although this replacement would succeed in blocking applicability to Max, it violates the first of the two Raz/Simpson conditions; the new rule is not simply a narrowing of Laura's original rule, but instead, applies in some situations where the original rule would not. Or suppose Ron had modified the original rule to read, "Children age nine or greater can stay up and watch TV, unless they have not finished their dinner." The modification would again block applicability to Max, since he did not finish his dinner, but in this case it violates the second of the two conditions; it fails to justify the original outcome in the original case of Emma, since she did not finish her dinner either.

If we understand the standard model as including the Raz/Simpson conditions on rule modification, then, we are able to fashion a response to our original problem concerning the possibility of constraint: even though later courts are free to modify the rules set out by earlier courts, they are still constrained by these rules, since they can modify them only in certain ways, those satisfying the two conditions. But responding to the original problem in this fashion leads at once to another, first highlighted, I believe, by Grant Lamond.[8] Presumably, even if a particular modification of an earlier rule satisfies the Raz/Simpson conditions, a later court would still not choose to modify the rule in that way unless the court believed that it could actually improve the rule by doing so. But if a later court

[8]See Lamond (2005, pp. 11–15).

believes that it can improve an earlier rule through modification, why should it limit itself to modifications that satisfy these conditions? Why should the court not be free to modify the rule in any way at all that leads to an improvement—in short: what is the justification for the Raz/Simpson conditions on rule modification?[9]

In recent work, motivated in part by research from the field of artificial intelligence and law, as well as by a previous proposal due to Lamond, I developed a model according to which precedential constraint is not a matter of rules at all, but of reasons.[10] On this view—which I refer to as the *reason model*—what matters about a precedent case is the previous court's assessment of the balance of reasons presented by that case; later courts are then constrained, not to follow some rule set out by the previous court, or even to modify this rule only in certain ways, but simply to reach decisions that are consistent with the earlier court's assessment of the balance of reasons. The development of the common law is to be pictured, then, not as the elaboration over time of an increasingly complex system of rules, but instead as the gradual construction of an ordering relation on reasons, reflecting importance, or priority.

What I show in the current paper is that, although this reason model of precedential constraint was explicitly developed as an alternative to the standard model, it turns out that the reason model can be used, also, to support the standard model, by providing a kind of semantic justification for the Raz/Simpson conditions on rule modification; indeed, I show that the standard model and the reason model are, in a straightforward sense, equivalent.

[9]Lamond (2005) puts the concern somewhat differently; he sees no justification for the Raz/Simpson conditions, and therefore argues that these conditions on rule modification are "haphazard" (p. 11) and "arbitrary" (p. 15).

[10]See Horty (2011) for the initial presentation of this model, and then Horty and Bench-Capon (2012) for a development of the model within the context of related research from artificial intelligence and law; see Lamond (2005) for his earlier proposal.

This equivalence between the two models of precedential constraint, though surprising, is also reassuring—in the way that it is always reassuring when two different analyses of a concept, starting from different initial points and relying on different ideas, agree in their outcome. Nevertheless, and in spite of the equivalence between these two models, I will argue that the reason model of precedential constraint provides a better picture than the standard model, for three reasons: it offers, first of all, representational advantages, second, a satisfying picture of common law reasoning more broadly, and third, resources for responding to some important criticisms of the standard model.

The paper is organized as follows. I begin, in the next section, by setting out the basic representational framework at work throughout the paper. Section 3 then presents, for the first time, I believe, a precise formulation of the standard model of precedential constraint—not just the Raz/Simpson conditions on rule modification, but the resulting constraints on decisions by later courts. Section 4 reviews the reason model presented in my earlier work. Section 5 then establishes the equivalence between these two models of precedential constraint, and begins the project of comparing them by describing the representational advantages of the reason model. This comparison is continued in Section 6, which describes the picture of common law reasoning underlying the reason model, and in Section 7, which shows how the reason model can be used in responding to criticisms of the standard model. There are two appendices. Appendix A collects together some abstract cases that are introduced as examples throughout the paper, while Appendix B verifies an observation underlying the central equivalence result.

2 Basic concepts

We will suppose that the situation presented to the court in a legal case can usefully by represented as a set of *factors*, where a factor is a legally significant fact or pattern of facts. The concept can be illustrated by returning to our domestic example. Here, the legal, or quasi-legal, issue at hand is whether a child can stay up and watch TV, and the factors involved might reasonably include those already considered—whether the child has reached the age of nine, whether dinner was eaten, homework completed—as well as countless others, such as: whether chores were done on time, whether good behavior was exhibited throughout the day, or whether the child has recently been subjected to any indignities that might merit special compensation, such as a nasty trip to the dentist.

But the factor-based representation of legal situations is not restricted only to everyday examples of this kind. In fact, this style of representation has been used to analyze case-based reasoning in a number of complex legal domains within the field of artificial intelligence and law, where it originated.[11] In the domain of trade secrets law, for example, where the factor-based analysis has been developed most extensively, a case will typically concern the issue of whether the defendant has gained an unfair competitive advantage over the plaintiff through the misappropriation of a trade secret; and here the factors involved might turn on, say, questions concerning whether the plaintiff took measures to protect the trade secret, whether a confidential relationship existed between the plaintiff and the defendant, whether the information acquired was reverse-engineerable or in some other way publicly available,

[11] The analysis of legal cases in terms of factors, initially taken only as points along legally significant dimensions, was first introduced by Rissland and Ashley (1987); see Ashley (1990) for a canonical treatment, and Rissland (1990) for an overview of research in artificial intelligence and law that places this work in a broader perspective. A useful criticism of the factor-based approach, along with further perspective, can be found in McCarty (1997).

and the extent to which this information did, in fact, lead to a real competitive advantage for the defendant.[12] Of course, the mere ability to understand a case in terms of the factors involved itself requires a significant degree of legal expertise, which is presupposed here. The theory thus begins with cases to which we must imagine that this expertise has already been applied, so that they can be represented directly in terms of the factors they present.

Many factors can naturally be taken to have polarities, favoring one side or another. In our domestic example, being older than nine or exhibiting good behavior throughout the day strengthens the child's claim, as plaintiff, that he or she should be allowed to stay up and watch TV; failing to finish dinner or to complete homework strengthens the parents' claim, as defendants, that the child should go to bed immediately. In the domain of trade secrets law, the presence of security measures likewise favors the plaintiff, since it strengthens the claim that the information secured was a valuable trade secret; reverse-engineerability favors the defendant, since it suggests that the product information might have been acquired through proper means. The present paper is based on the simplifying assumption, not just that many, or even most, factors have polarities, but that all factors are like this, favoring one particular side. In addition, we rely on the further assumption, also a simplification, that the reasoning under consideration involves only a single step, proceeding directly from the factors present in a case to a decision—in favor of the plaintiff or the defendant—rather

[12] Aleven (1997) has analyzed 147 cases from trade secrets law in terms of a factor hierarchy that includes 5 high-level issues, 11 intermediate-level concerns, and 26 base-level factors. The resulting knowledge base is used in an intelligent tutoring system for teaching elementary skills in legal argumentation, which has achieved results comparable to traditional methods of instruction in controlled studies; see Aleven and Ashley (1997). The formal treatment sketched in the present paper abstracts considerably from this detailed representational work, and in particular, the idea that legal factors are organized into a hierarchy is missing entirely.

than moving through a series of intermediate legal concepts.[13]

Formally, then, we begin by postulating a set F of legal factors. A *fact situation* X, of the sort presented in a legal case, can then be defined as some particular subset of these factors: $X \subseteq F$. We will let $F^\pi = \{f_1^\pi, \ldots, f_n^\pi\}$ represent the set of factors favoring the plaintiff and $F^\delta = \{f_1^\delta, \ldots, f_m^\delta\}$ the set of factors favoring the defendant. Given our assumption that each factor favors one side of the other, we can suppose that the entire set of legal factors is exhausted by those favoring the plaintiff together with those favoring the defendant, so that $F = F^\pi \cup F^\delta$.

A *precedent case* will be represented as a fact situation together with an outcome and a rule through which that outcome is reached. Such a case, then, can be defined as a triple of the form $c = \langle X, r, s \rangle$, where X is a fact situation containing the legal factors presented by the case, r is the rule of the case, and s is its outcome.[14] We define three functions—*Factors*, *Rule*, and *Outcome*—to map cases into their component parts, so that, in the case c above, for example, we would have $\mathit{Factors}(c) = X$, $\mathit{Rule}(c) = r$, and $\mathit{Outcome}(c) = s$.

Given our assumption that reasoning proceeds in a single step, we can suppose that the *outcome* s of a case is always either a decision in favor of the plaintiff or a decision in favor of the defendant, with these two outcomes represented as π or δ respectively; and where s is a particular outcome, a decision for some side, we suppose that $\overline{s}$ represents a decision for the opposite side, so that $\overline{\pi} = \delta$ and $\overline{\delta} = \pi$. Where X is a fact situation, we let X^s represent the factors from X that support the side s; that is, $X^\pi = X \cap F^\pi$ is the set of factors from X

[13] These simplifications, and others, are discussed further in the final section of Horty (2011).

[14] We suppose, as a further simplification, that the rule underlying a court's decision is plain, ignoring the extensive literature on methods for determining the rule, or *ratio decidendi*, of a case; we suppose also that a case always contains a single rule, ignoring situations in which a judge might offer several rules for a decision, or in which a court reaches a decision by majority, with different judges offering different rules, or in which a court might simply render a decision in a case without setting out any general rule at all.

favoring the plaintiff, while $X^\delta = X \cap F^\delta$ is the set of factors from X favoring the defendant.

Rules will be defined in terms of reasons, where a *reason for a side* is some set of factors favoring that side. A *reason* can then be defined as a set of factors favoring one side or another. To illustrate: $\{f_1^\pi, f_2^\pi\}$ is a reason favoring the plaintiff, and so a reason, while $\{f_1^\delta\}$ is a reason favoring the defendant, and likewise a reason; but the set $\{f_1^\pi, f_1^\delta\}$ is not a reason, since the factors it contains do not uniformly favor one side or another. Reasons of this kind are to be interpreted conjunctively, so that, for example, the reason $\{f_1^\pi, f_2^\pi\}$ can be identified with the proposition that both the factors f_1^π and f_2^π are present.

A statement of the form $X \models R$ indicates that the reason R *holds* in the fact situation X, or using more technical language, that the fact situation *satisfies* the reason; this idea can be defined by stipulating that

$$X \models R \text{ if and only if } R \subseteq X,$$

and then extended in the usual way to statements ϕ and ψ formed by closing the sets of factors representing reasons under conjunction and negation:

$$X \models \neg\phi \text{ if and only if it fails that } X \models \phi,$$

$$X \models \phi \wedge \psi \text{ if and only if } X \models \phi \text{ and } X \models \psi.$$

To illustrate, consider the fact situation $X_1 = \{f_1^\pi, f_2^\pi, f_3^\pi, f_1^\delta, f_2^\delta, f_3^\delta, f_4^\delta\}$. It is easy to see that the reason $\{f_1^\pi, f_2^\pi\}$ holds in this fact situation, since $\{f_1^\pi, f_2^\pi\}$ is a subset of X_1, but that the reason $\{f_5^\delta\}$ fails to hold, since it is not a subset; as a result, we have $X_1 \models \{f_1^\pi, f_2^\pi\}$ and $X_1 \models \neg\{f_5^\delta\}$, so that $X_1 \models \{f_1^\pi, f_2^\pi\} \wedge \neg\{f_5^\delta\}$.

Based on this notion of a reason, a rule can now be defined as a pair whose conclusion is an outcome, a decision for one side or the other, and whose premise is a conjunction containing a reason favoring that side together with statements indicating that certain reasons favoring the opposite side are not present. More exactly, where R^s is a reason for the side s and

$R_1^{\overline{s}}, \ldots R_i^{\overline{s}}$ are zero or more reasons for the opposite side, a *rule for the side s* has the form

$$R^s \wedge \neg R_1^{\overline{s}} \wedge \ldots \wedge \neg R_i^{\overline{s}} \rightarrow s,$$

and a *rule* is simply a rule for one side or the other. Our rules, then, are required to take a limited syntactic form, but one that is, also, a very natural form. The idea is that, when the reason R^s favoring s holds in some situation, and none of the reasons $R_1^{\overline{s}}, \ldots R_i^{\overline{s}}$ favoring the opposite side hold, then r requires a decision for the side s.

When r is a rule of the form displayed above, we define functions *Premise*, *Premise*^{s}, and *Conclusion* picking out its premise, the positive part of its premise, and its conclusion, as follows:

$$\textit{Premise}(r) = R^s \wedge \neg R_1^{\overline{s}} \wedge \ldots \wedge \neg R_i^{\overline{s}},$$

$$\textit{Premise}^s(r) = R^s,$$

$$\textit{Conclusion}(r) = s.$$

We can then say that r *applies* in a fact situation X just in case $X \models \textit{Premise}(r)$. And when a decision for the side s is justified by the rule r, we will refer to $\textit{Premise}^s(r) = R^s$ as the *reason for the decision*—the positive consideration favoring that outcome.[15]

Let us return, now, to the concept of a precedent case $c = \langle X, r, s \rangle$, containing a fact situation X along with a rule r leading to the outcome s. In order for this concept to make sense, we impose two coherence constraints. The first is that the rule contained in the case must actually apply to the facts of the case, or that $X \models \textit{Premise}(r)$. The second is that the conclusion of the precedent rule must match the outcome of the case itself, or that $\textit{Conclusion}(r) = \textit{Outcome}(c)$.

[15] Some writers argue that the entire premise of the rule r should be taken as a reason for the decision—that is, $\textit{Premise}(r) = R^s \wedge \neg R_1^{\overline{s}} \wedge \ldots \wedge \neg R_i^{\overline{s}}$, the positive consideration favoring the outcome together with statements indicating that various considerations favoring the other side are not present. This issue is discussed at length in Dancy (2004, pp. 38–52), and also in Horty (2012, pp. 53–59, 141–146).

These various concepts and constraints can be illustrated through the case $c_1 = \langle X_1, r_1, s_1 \rangle$, where $X_1 = \{f_1^\pi, f_2^\pi, f_3^\pi, f_1^\delta, f_2^\delta, f_3^\delta, f_4^\delta\}$ is the underlying fact situation, with three factors favoring the plaintiff and four favoring the defendant, where $r_1 = \{f_1^\pi, f_2^\pi\} \wedge \neg\{f_5^\delta\} \rightarrow \pi$ is the rule of the case, and where $s_1 = \pi$ is its outcome, a decision for the plaintiff. Since we have both $X_1 \models \mathit{Premise}(r_1)$ and $\mathit{Conclusion}(r_1) = \mathit{Outcome}(c_1)$, it is clear that the case satisfies our two coherence constraints: the rule of the case applies to the fact situation, and the conclusion of the rule matches the outcome of the case. This particular case, then, is one in which the court decided for the plaintiff by applying or introducing a rule according to which the presence of the factors f_1^π and f_2^π, together with the absence of the factor f_5^δ, leads to decision for the plaintiff. The reason for the decision is $\mathit{Premise}^s(r_1)$, or $\{f_1^\pi, f_2^\pi\}$.

With this notion of a case in hand, we can define a *case base* simply as a set Γ of cases—a set of fact situations presented to various courts, together with their outcomes and the rules justifying these outcomes. It is a case base of this sort that will be taken to represent the common law in some area, and to constrain the decisions of future courts.

3 The standard model

We now turn to the standard model of precedential constraint, formulated in terms of rules that can be modified in accord with the Raz/Simpson conditions. I proceed by tracing three simple examples of legal development in accord with the standard model, generalizing from these examples, and then characterizing the standard model itself in terms of this generalization.

To begin with, then, suppose that the background case base is $\Gamma_1 = \{c_2\}$, containing only the single precedent case $c_2 = \langle X_2, r_2, s_2 \rangle$, where $X_2 = \{f_1^\pi, f_1^\delta\}$, where $r_2 = \{f_1^\pi\} \rightarrow \pi$, and where $s_2 = \pi$; this precedent represents a situation in which a prior court, confronted with

the conflicting factors f_1^π and f_1^δ, decided for the plaintiff on the basis of f_1^π. Now imagine that, against the background of this case base, a later court is confronted with the new fact situation $X_3 = \{f_1^\pi, f_2^\delta\}$, and takes the presence of the new factor f_2^δ as sufficient to justify a decision for the defendant. Of course, the previous rule r_2 applies to the new fact situation, apparently requiring a decision for the plaintiff. But according to the standard model, the court can decide for the defendant all the same by distinguish the new fact situation from that of the case in which r_2 was introduced—pointing out that the new situation, unlike that of the earlier case, contains the factor f_2^δ, and so declining to apply the earlier rule on that basis.

The result of this decision, then, is that the original case base is changed in two ways. First, by deciding the new situation for the defendant on the basis of f_2^δ, the court supplements this case base with the new case $c_3 = \langle X_3, r_3, s_3 \rangle$, where X_3 is as above, where $r_3 = \{f_2^\delta\} \rightarrow \delta$, and where $s_3 = \delta$. And second, by declining to apply the earlier r_2 to the new situation due to the presence of f_2^δ, the court, in effect, modifies this earlier rule to accommodate its decision, so that the earlier rule now carries the force of $r_2{}' = \{f_1^\pi\} \wedge \neg\{f_2^\delta\} \rightarrow \pi$. Note that the modification conforms to the Raz/Simpson conditions, both narrowing the earlier rule, and narrowing it in such a way that it continues to support the earlier conclusion. The new case base is thus $\Gamma_1{}' = \{c_2{}', c_3\}$, with $c_2{}' = \langle X_2{}', r_2{}', s_2{}' \rangle$ as a modification of the previous c_2, where $X_2{}' = X_2$, where $r_2{}'$ is as above, and where $s_2{}' = s_2$; and with c_3 as above.

The process could continue, of course. Suppose that, against the background of the modified case base $\Gamma_1{}' = \{c_2{}', c_3\}$, a court is confronted with the further fact situation $X_4 = \{f_1^\pi, f_3^\delta\}$, and again takes the new factor f_3^δ as sufficient to justify a decision for the defendant, in spite of the fact that even the modified rule $r_2{}'$ requires a decision for

the plaintiff. Once again, this decision changes the current case base in two ways: first, supplementing this case base with a new case representing the current decision, and second, further modifying the previous rule to accommodate this new decision as well. The resulting case base is therefore $\Gamma_1{''} = \{c_2{''}, c_3, c_4\}$, with $c_2{''} = \langle X_2{''}, r_2{''}, s_2{''}\rangle$ as a modification of the previous $c_2{'}$, where $X_2{''} = X_2{'}$, where $r_2{''} = \{f_1^\pi\} \wedge \neg\{f_2^\delta\} \wedge \neg\{f_3^\delta\} \rightarrow \pi$, and where $s_2{''} = s_2{'}$; with c_3 as above; and with $c_4 = \langle X_4, r_4, s_4\rangle$ representing the current decision, where X_4 is as above, where $r_4 = \{f_3^\delta\} \rightarrow \delta$, and where $s_4 = \delta$,

As our second example, suppose that the background case base is $\Gamma_2 = \{c_2, c_5\}$, with c_2 as above, and now with $c_5 = \langle X_5, r_5, s_5\rangle$, where $X_5 = \{f_1^\pi, f_2^\delta\}$, where $r_5 = \{f_1^\pi\} \rightarrow \pi$, and where $s_5 = \pi$. This case base represents a pair of prior decisions for the plaintiff on the basis of f_1^π, in spite of the conflicting factor f_1^δ in one case, and f_2^δ in the other. Now imagine that, against this background, a later court confronts the new situation $X_6 = \{f_1^\pi, f_1^\delta, f_2^\delta\}$, and decides that, although earlier cases favored f_1^π over the conflicting f_1^δ and f_2^δ presented separately, the combination of f_1^δ and f_2^δ together outweighs f_1^π, and so justifies a decision for the defendant. Again, this decision supplements the existing case base with a new case $c_6 = \langle X_6, r_6, s_6\rangle$, where X_6 is as above, where $r_6 = \{f_1^\delta, f_2^\delta\} \rightarrow \delta$, and where $s_6 = \delta$. But here, the rules from both of the existing cases, c_2 and c_5, must be modified to block application to situations in which f_1^δ and f_2^δ appear together, though continuing to allow application to situations in which those factors appear separately; the rules will thus carry the force of $r_2{'} = r_5{'} = \{f_1^\pi\} \wedge \neg\{f_1^\delta, f_2^\delta\} \rightarrow \pi$. The case base resulting from this decision is $\Gamma_2{'} = \{c_2{'}, c_5{'}, c_6, \}$, with $c_2{'} = \langle X_2{'}, r_2{'}, s_2{'}\rangle$, where $X_2{'} = X_2$, where $r_2{'}$ is as above, and where $s_2{'} = s_2$; with $c_5{'} = \langle X_5{'}, r_5{'}, s_5{'}\rangle$, where $X_5{'} = X_5$, where $r_5{'}$ is as above, and where $s_5{'} = s_5$; and with c_6 as above.

Finally, suppose the background case base is $\Gamma_3 = \{c_2, c_7\}$, again with c_2 as above, and

with $c_7 = \langle X_7, r_7, s_7 \rangle$, where $X_7 = \{f_2^\pi, f_2^\delta\}$, where $r_7 = \{f_2^\pi\} \rightarrow \pi$, and where $s_7 = \pi$. This case base represents a pair of previous decisions for the plaintiff, one on the basis of f_1^π in spite of the conflicting f_1^δ, and one on the basis of f_2^π in spite of the conflicting f_2^δ. Now imagine that a later court confronts the new situation $X_8 = \{f_1^\pi, f_2^\delta\}$, containing the two factors that have not yet been compared, and concludes that f_2^δ is sufficient to justify a decision for the defendant, in spite of the conflicting f_1^π. As before, the earlier rule r_2 must be taken to have the force of $r_2{}' = \{f_1^\pi\} \wedge \neg\{f_2^\delta\} \rightarrow \pi$, in order to accommodate the current decision. In this case, however, the new rule cannot be formulated simply as $\{f_2^\delta\} \rightarrow \delta$, but must now have the form $r_8 = \{f_2^\delta\} \wedge \neg\{f_2^\pi\} \rightarrow \delta$, in order to accommodate the decision for the plaintiff previously reached in c_7. This scenario, then, is one in which modifications are forced in both directions: a previous rule must be modified to avoid conflict with the current decision, while at the same time, the rule set out in the current case must be hedged to avoid conflict with a previous decision. The resulting case base is $\Gamma_3{}' = \{c_2{}', c_7, c_8, \}$, with $c_2{}' = \langle X_2{}', r_2{}', s_2{}' \rangle$, where $X_2{}' = X_2$, where $r_2{}'$ is as above, and where $s_2{}' = s_2$; with c_7 as above; and with $c_8 = \langle X_8, r_8, s_8 \rangle$, where X_8 is as above, where r_8 is as above, and where $s_8 = \delta$.

Each of these examples describes a scenario in which a sequence of fact situations is confronted, decisions are reached, rules are formulated to justify these decisions, and rules are modified in accord with the Raz/Simpson conditions to accommodate later, or earlier, decisions. It is interesting to note that, as long as all decisions from a case base can be accommodated in this way, with rules modified appropriately, then the order in which these cases are confronted is irrelevant. To put this point precisely, recall from the previous section that, where $c = \langle X, r, s \rangle$ is a precedent case decided for the side s, the reason for this decision is $\mathit{Premise}^s(r)$, the positive consideration from the premise of the case rule;

and suppose that a case base has been constructed through the process of considering fact situations in some particular sequence and, on each occasion, reaching a decision for some particular reason and modifying other rules accordingly. It then follows that the same case base will be constructed, with all rules modified in the same way, regardless of the sequence in which the fact situations are considered, as long as the same decisions are reached in each case, and for the same reasons. Indeed, the fact situations need not be considered in any sequence at all: if the set of decisions in these situations is capable of being accommodated through appropriate rule modifications, then all the rules can be modified at once, through the process of *case base refinement*.

Definition 1 (Refinement of a case base) Where Γ is a case base, its refinement—written, Γ^+—is the set that results from carrying out the following procedure. For each case $c = \langle X, r, s\rangle$ belonging to Γ:

1. Let
$$\Gamma_c = \{c' \in \Gamma_c : c' = \langle Y, r', \overline{s}\rangle \ \&\ Y \models \mathit{Premise}^s(r)\}$$
2. For each case $c' = \langle Y, r', \overline{s}\rangle$ from Γ_c, let
$$d_{c,c'} = \neg\, \mathit{Premise}^{\overline{s}}(r')$$
3. Define
$$D_c = \bigwedge_{c' \in \Gamma_c} d_{c,c'}$$
4. Replace the case $c = \langle X, r, s\rangle$ from Γ with $c'' = \langle X, r'', s\rangle$, where r'' is the new rule
$$\mathit{Premise}^s(r) \wedge D_c \rightarrow s$$

This process of transforming a case base Γ into its refinement Γ^+ can be described informally as follows. First, for each case c belonging to Γ, decided for some side and for

some particular reason, collect together into Γ_c all of the cases from Γ in which that reason holds, but which were decided for the other side. Second, for each such case c' from Γ_c, let $d_{c,c'}$—the consideration that distinguishes c from c'—be the negation of the reason for which c' was decided. Third, conjoin these various distinguishing considerations together into a single statement D_c, which then distinguishes c from all cases in Γ_c at once. And fourth, replace the original rule r from c with the new rule r'' whose premise is formed by conjoining $\mathit{Premise}^s(r)$, the reason for the original decision, with the distinguishing statement D_c—resulting in a rule that will no longer apply to any other case in which the reason for the original decision holds, but which was decided for the opposite side.

The reader is invited to verify that, in each of the three examples we have considered, the case bases resulting from our suggested modifications are identical with those that would have resulted simply from deciding the same fact situations for the same reasons, including the new decisions along with the cases already decided, and the refining the result. Focusing on the first of our examples, we can see, for instance, that ${\Gamma_1}' = (\Gamma_1 \cup \{c_3\})^+$, and then that ${\Gamma_1}'' = ({\Gamma_1}' \cup \{c_4\})^+$—or, considering the cases simultaneously, that ${\Gamma_1}'' = (\Gamma_1 \cup \{c_3, c_4\})^+$.

In the examples we have considered, then, where a decision can be accommodated against the background of a case base by means of an appropriate modification of rules, the same outcome can be achieved, with all rules modified in the same way, simply by supplementing the background case base with that decision and then refining the result. My suggestion is that what holds in these particular examples also holds in general: development of a case base in accord with the standard model amounts to supplementing the existing case base with new decisions, and then accommodating these new decisions by refining the supplemented case base.

But what if a particular decision cannot be accommodated against the background of a

case base? What does the process of refinement lead to then? The answer is that, when a case base is supplemented with a decision that cannot be accommodated through appropriate rule modification, the formal process of refinement will then alter some of the underlying rules in such a way that they fail to apply to their corresponding fact situations—so that, strictly speaking, the result will not be a case at all.[16] The linkage between accommodation and refinement therefore works in both directions, and can provide us with a formal explication of the concept of accommodation: a decision can be *accommodated* against the background of a case base, with rules modified appropriately, just in case the result of supplementing the case base with that decision can itself be refined into a case base.

We are now in a position to define the notion of precedential constraint, according to the standard model. The intuitive idea is that a court is constrained to reach a decision that can be accommodated within the context of a background case base through an appropriate modification of rules—or, given our formal explication of this concept, a decision that can be combined with the background case base to yield a result whose refinement is itself a case base.

Definition 2 (Precedential constraint: the standard model) Let Γ be a case base and X a new fact situation confronting the court. Then the standard model of precedential constraint requires the court to base its decision on some rule r leading to an outcome s such that $(\Gamma \cup \{\langle X, r, s\rangle\})^+$ is a case base.

This definition can be illustrated by taking as background the case base $\Gamma_4 = \{c_9\}$, containing the single case $c_9 = \langle X_9, r_9, s_9\rangle$, where $X_9 = \{f_1^\pi, f_2^\pi, f_1^\delta, f_2^\delta\}$, where $r_9 = \{f_1^\pi\} \rightarrow \pi$, and where $s_9 = \pi$. Now suppose the court confronts the new situation $X_{10} = \{f_1^\pi, f_1^\delta, f_2^\delta, f_3^\delta\}$,

[16] Recall from the previous section that a case base is defined as a set of cases, and that a case is subject to the coherence condition that the rule of a case must apply to its fact situation.

and considers finding for the defendant on the basis of f_1^δ and f_2^δ, leading to the decision $c_{10} = \langle X_{10}, r_{10}, s_{10}\rangle$, where X_{10} is as above, where $r_{10} = \{f_1^\delta, f_2^\delta\} \rightarrow \delta$, and where $s_{10} = \delta$. According to our proposed definition, this decision would be ruled out by the standard model of precedential constraint, since the result of supplementing the background case base Γ_4 with the new decision c_{10}—that is $\Gamma_4 \cup \{c_{10}\}$—cannot itself be refined into a case base. Indeed, the refinement of this supplemented case base is the set $(\Gamma_4 \cup \{c_{10}\})^+ = \{c_9{'}, c_{10}{'}\}$, with $c_9{'} = \langle X_9{'}, r_9{'}, s_9{'}\rangle$, where $X_9{'} = X_9$, where $r_9{'} = \{f_1^\pi\} \wedge \neg\{f_1^\delta, f_2^\delta\} \rightarrow \pi$, and where $s_9 = \pi$; and with $c_{10}{'} = \langle X_{10}{'}, r_{10}{'}, s_{10}{'}\rangle$, where $X_{10}{'} = X_{10}$, where $r_{10}{'} = \{f_1^\delta, f_2^\delta\} \wedge \neg\{f_1^\pi\} \rightarrow \delta$, and where $s_{10} = \delta$. But this set is not a case base at all: neither $c_9{'}$ nor $c_{10}{'}$ is a case, in our technical sense, since the rule $r_9{'}$ fails to apply to the fact situation $X_9{'}$, and the rule $r_{10}{'}$ fails to apply to the fact situation $X_{10}{'}$.

4 The reason model

Having provided a formal reconstruction of what I take to be the standard model of precedential constraint, in terms of rules that can be modified, I now want to review my own previous proposal, the reason model, developed in terms of an ordering relation on reasons.

In order to motivate this proposal, let us return to the case $c_9 = \langle X_9, r_9, s_9\rangle$—where again $X_9 = \{f_1^\pi, f_2^\pi, f_1^\delta, f_2^\delta\}$, where $r_9 = \{f_1^\pi\} \rightarrow \pi$, and where $s_9 = \pi$—and ask what information is actually carried by this case; what is the court telling us with its decision? Well, two things, at least. First of all, by appealing to the rule r_9 as justification, the court is telling us that the reason for its decision—that is, $Premise^\pi(r_9)$, or $\{f_1^\pi\}$—is sufficient to justify a decision in favor of the plaintiff. But second, with its decision for the plaintiff, the court is also telling us that this reason must be stronger than the strongest reason presented by the case in favor of the defendant.

To put this precisely, let is first stipulate that, if X and Y are reasons favoring the same side, then Y is *at least as strong* a reason as X for that side whenever $X \subseteq Y$. Returning to our example, then, where $X_9 = \{f_1^\pi, f_2^\pi, f_1^\delta, f_2^\delta\}$, it is clear that the strongest reason present for the defendant is $X_9^\delta = \{f_1^\delta, f_2^\delta\}$, containing all those factors from the original fact situation that favor the defendant. Since the c_9 court has decided for the plaintiff on the grounds of the reason $Premise^\pi(r_9)$, even in the face of the conflicting X_9^δ, it seems to follow as a consequence of the court's decision that the reason $Premise^\pi(r_9)$ for the plaintiff should be assigned a higher priority than the reason X_9^δ for the defendant—that is, that $\{f_1^\pi\}$ should be assigned a higher priority than $\{f_1^\delta, f_2^\delta\}$. If we let $<_{c_9}$ represent the priority relation on reasons that is derived from the particular case c_9, then this consequence of the court's decision can be put more formally as the claim that $\{f_1^\delta, f_2^\delta\} <_{c_9} \{f_1^\pi\}$, or equivalently, that $X_9^\delta <_{c_9} Premise^\pi(r_9)$.

As far as the priority ordering goes, then, the earlier court is telling us at least that $X_9^\delta <_{c_9} Premise^\pi(r_9)$, but is it telling us anything else? Perhaps not explicitly, but implicitly, yes. For if the reason $Premise^\pi(r_9)$ for the plaintiff is preferred to the reason X_9^δ for the defendant, then surely any reason for the plaintiff that is at least as strong as $Premise^\pi(r_9)$ must likewise be preferred to X_9^δ, and just as surely, $Premise^\pi(r_9)$ must be preferred to any reason for the defendant that is at least as weak as X_9^δ. As we have seen, a reason Z for the plaintiff is at least as strong as $Premise^\pi(r_9)$ if it contains all the factors contained in $Premise^\pi(r_9)$—that is, if $Premise^\pi(r_9) \subseteq Z$. And we can conclude, likewise, that a reason W for the defendant is at least as weak as X_9^δ if it contains no more factors than X_9^δ itself—that is, if $W \subseteq X_9^\delta$. It therefore follows from the earlier court's decision in c_9, not only that $X_9^\delta <_{c_9} Premise^\pi(r_9)$, but that $W <_{c_9} Z$ whenever W is at least as weak a reason for the defendant as X_9^δ and Z is at least as strong a reason for the plaintiff as $Premise^\pi(r_9)$—

whenever, that is, $W \subseteq X_9^\delta$ and $Premise^\pi(r_9) \subseteq Z$. To illustrate: from the court's explicit decision that $\{f_1^\delta, f_2^\delta\} <_{c_9} \{f_1^\pi\}$, we can conclude also that $\{f_1^\delta\} <_{c_9} \{f_1^\pi, f_3^\pi\}$, for example.

This line of argument leads to the following definition of the preference relation among reasons that can be derived from a single case.

Definition 3 (Priority relation derived from a case) Let $c = \langle X, r, s\rangle$ be a case, and suppose W and Z are reasons. Then the relation $<_c$ representing the priority on reasons derived from the case c is defined by stipulating that $W <_c Z$ if and only if $W \subseteq X^{\overline{s}}$ and $Premise^s(r) \subseteq Z$.

Once we have defined the priority relation derived from a single case, we can introduce a priority relation $<_\Gamma$ derived from an entire case base Γ in the natural way, by stipulating that one reason has a higher priority than another according to the entire case base whenever that priority relation is supported by some particular case from the case base.

Definition 4 (Priority relation derived from a case base) Let Γ be a case base, and suppose W and Z are reasons. Then the relation $<_\Gamma$ representing the priority relation on reasons derived from the case base Γ is defined by stipulating that $W <_\Gamma Z$ if and only if $W <_c Z$ for some case c from Γ.

And we can then define a case base as reason inconsistent if it provides conflicting information about the priority among reasons—telling us, for any two reasons, that each has a higher priority than the other—and reason consistent otherwise.

Definition 5 (Reason consistent case bases) Let Γ be a case base with $<_\Gamma$ its derived priority relation. Then Γ is reason inconsistent if and only if there are reasons X and Y such that $X <_\Gamma Y$ and $Y <_\Gamma X$. Γ is reason consistent if and only if it is not reason inconsistent.

Given this notion of reason consistency, we can now turn to the concept of precedential constraint itself, according to the reason model. The intuition is simply that, in deciding a case, a constrained court is required to preserve the consistency of the background case base. Suppose, more exactly, that a court constrained by a background case base Γ is confronted with a new fact situation X. Then the court is required to reach a decision on X that is itself consistent with Γ—that is, a decision that does not introduce inconsistency.

Definition 6 (Precedential constraint: the reason model) Let Γ be a case base and X a new fact situation confronting the court. Then the reason model of precedential constraint requires the court to base its decision on some rule r leading to an outcome s such that the new case base $\Gamma \cup \{\langle X, r, s\rangle\}$ is reason consistent.

This idea can be illustrated by assuming as background the previous case base $\Gamma_4 = \{c_9\}$, containing only the previous case c_9, supposing once again that, against this background, the court confronts the fresh situation $X_{10} = \{f_1^\pi, f_1^\delta, f_2^\delta, f_3^\delta\}$ and considers finding for the defendant on the basis of f_1^δ and f_2^δ, leading to the decision $c_{10} = \langle X_{10}, r_{10}, s_{10}\rangle$, where X_{10} is as above, where $r_{10} = \{f_1^\delta, f_2^\delta\} \rightarrow \delta$, and where $s_{10} = \delta$. We saw in the previous section that such a decision would fail to satisfy the rule constraint, and we can see now that it fails to satisfy the reason constraint as well. Why? Because the new case c_{10} would support the priority relation $\{f_1^\pi\} <_{c_{10}} \{f_1^\delta, f_2^\delta\}$, telling us that the reason $\{f_1^\delta, f_2^\delta\}$ for the defendant outweighs the reason $\{f_1^\pi\}$ for the plaintiff. But Γ_4 already contains the case c_9, from which we can derive the priority relation $\{f_1^\delta, f_2^\delta\} <_{c_9} \{f_1^\pi\}$, telling us exactly the opposite. As a result, the augmented case base $\Gamma_4 \cup \{c_{10}\}$ would be reason inconsistent.

5 Comparing the models

The two models now before us—the standard model from Section 3 and the reason model from Section 4—offer strikingly different pictures of precedential constraint, and of the process of common law development.

According to the standard model, what is important about a background case base is the set of rules it contains, together with the facts of the cases in which these rules were previously formulated or applied. In reaching a decision concerning a new fact situation, the court is then required to modify the existing set of rules, if necessary, in order to accommodate its decision. Precedential constraint is due to the fact that such accommodation is not always possible; legal development results from the modification of existing rules, together with the introduction of new rules from new decisions. According to the reason model, what is important about a background case base is, not the set of rules it contains, but instead, a priority ordering on reasons derived from decisions in previous cases. In confronting a new fact situation, a court is required only to reach a decision that is consistent with this existing priority ordering. Constraint is due to the fact that not all possible decisions are consistent; legal development results from supplementing the existing priority ordering with the new priorities generated by decisions in new cases, which then strengthen the overall priority ordering on reasons.

Given the very different pictures associated with these two models of precedential constraint, it is surprising to note that the two models are equivalent, in the following sense: against the background of a fixed case base, any decision in a new case, for a particular side and on the basis of a particular reason, satisfies the standard model of precedential constraint just in case it satisfies the reason model. This observation—the central result of the paper—follows at once from a preliminary observation, proved in Appendix B, linking

the technical concepts underlying the two models of constraint:

Observation 1 Let Γ be a case base. Then Γ is reason consistent if and only if its refinement Γ^+ is itself a case base.

Our central result can now be stated as follows:

Observation 2 Let Γ be a case base, and let X be a new fact situation. Then a decision in the situation X, on the basis of a rule r leading to an outcome s, satisfies the standard model of precedential constraint if and only if that same decision satisfies the reason model.

And the result can be verified very simply. Suppose first that, against the background of the case base Γ, a decision in the new fact situation X, on the basis of r and leading to s, satisfies the standard model of precedential constraint. What this means, by Definition 2, is that $(\Gamma \cup \{\langle X, r, s\rangle\})^+$—the refinement of Γ once it is supplemented with the new decision $\langle X, r, s\rangle$—must itself be a case base. It then follows from the equivalence established in Observation 1 that the supplemented case base $\Gamma \cup \{\langle X, r, s\rangle\}$ is reason consistent, and so from Definition 6, that the same decision satisfies the reason model. Next, suppose that a decision in the situation X, on the basis of r and leading to s, satisfies the reason model of precedential constraint. What this means, by Definition 6, is that the supplemented case base $\Gamma \cup \{\langle X, r, s\rangle\}$ must be reason consistent. From Observation 1, again, it then follows that the refinement $(\Gamma \cup \{\langle X, r, s\rangle\})^+$ of this supplemented case base is itself a case base, and so from Definition 2, that the same decision satisfies the standard model.

This result shows how the reason model of precedential constraint can be used, as suggested earlier, to supply a semantic justification for the Raz/Simpson conditions on rule modification. Suppose that, from a perspective that takes the reason model as fundamental, we imagine a model based on rule modification instead, such as the standard model, and

search for conditions on rule modification that will guarantee its equivalence with the reason model; what our central result then shows is that it is exactly the Raz/Simpson conditions that do the job. Still, and especially in light of this equivalence between the reason model and the standard model, now taken together with the Raz/Simpson conditions, it is natural to ask why we should view the reason model as more fundamental at all—why not take the standard model as fundamental, so that it is the reason model that needs to be justified, or more plausibly, why not simply take the two models of precedential constraint as two different accounts of the same phenomenon that happen to agree, without supposing that either is necessarily more fundamental than the other? This is the question I focus on throughout the remainder of the paper, beginning in this section with a technical point.

Whatever else a case base might be, whatever other legal or social roles it might serve, it is at least a repository of legal information, and indeed a dynamic repository, recording the results in previous cases and routinely updated with information from new decisions. It is therefore appropriate that our overall account of common law reasoning should be evaluated, at least in part, according to the standards developed within the subfield of computer science known as "knowledge representation," which studies the most efficient ways of representing, reasoning with, and maintaining a given body information, or "knowledge."[17] One of the most important of these standards is the requirement that the representation of information should be *modular*, so that it can proceed piece by piece. In the present setting, if we let f be the function that maps a case base into its formal representation, and where Γ and Γ' are case

[17]From a philosophical standpoint, the use of the term "knowledge" in the designation of this subject is unfortunate, since there is no requirement that the information represented should be true—"belief representation," or simply "information representation," would have been better, but the term "knowledge" is now established. Some classic readings in the area are collected together in Brachman and Levesque (1985); a contemporary textbook is Brachman and Levesque (2004).

bases, then modularity can be expressed as the condition that $f(\Gamma \cup \Gamma') = f(\Gamma) \cup f(\Gamma')$—that a compound case base can be represented simply by joining together the representations of its parts.

What modularity requires, quite generally, is that the representation of information must be atomistic, rather than holistic—or in the present setting, that the representation of a particular portion of the case base cannot be affected by what is present elsewhere. From a technical standpoint, this kind of atomism is especially important in a dynamic "knowledge base," or store of information. Any such system must be able to handle updates in some efficient way; but update can be a very complicated process if the representation of information is holistic, so that the addition of new information triggers changes in the representation of information already present. The modularity condition can be given a philosophical reading as well, if we take the formal representation of an item of information as capturing, in some sense, its real meaning, as far as the relevant domain of reasoning is concerned. In the present setting, then, modularity, and the associated atomism, can also be seen as guaranteeing that a case should have whatever meaning it does—that it should constrain our reasoning in the same way—regardless of the overall case base in which that case occurs.

Now, how do our two models of precedential constraint, standard and reason, fare with respect to the requirement of modularity? The idea underlying the standard model is that a case base should be represented through its refinement—or, again taking f as the function mapping a case base Γ into its representation, that $f(\Gamma) = \Gamma^+$. It is, therefore, clear that this model fails to satisfy modularity; we do not in general have

$$(\Gamma \cup \Gamma')^+ = \Gamma^+ \cup \Gamma'^+.$$

The reader is invited to verify this point by returning to our initial example from Section 3 and noting that $(\Gamma_1 \cup \{c_3\})^+$ does not coincide with $\Gamma_1^+ \cup \{c_3\}^+$, since the case base $\Gamma_1^+ \cup$

$\{c_3\}^+$, where the components are refined separately and then joined together, is identical to $\Gamma_1 \cup \{c_3\}$, while the joint refinement $(\Gamma_1 \cup \{c_3\})^+$ is not; in the latter, the case c_2 already present in Γ_1 must be modified to accommodate the new c_3. Indeed, as we have seen, it is a central feature of the standard model that cases already present in the case base are modified to accommodate new decisions, so that the representational strategy is holistic, rather than atomistic. As a result, updating a case base can be a complex operation, and the representation of the same decision can vary from one case base to another, depending what other cases might be present.

The reason model, by contrast, presents a much smoother picture. All that matters about a case base, according to this model, is the priority relation among reasons it generates—so that, where f is the function mapping a case base Γ into its representation, we have $f(\Gamma) = <_\Gamma$. Modularity then follows at once from Definition 4, which tells us that

$$<_{\Gamma \cup \Gamma'} \; = \; <_\Gamma \cup <_{\Gamma'} .$$

Because the representation of legal information associated with the reason model is, therefore, atomistic rather than holistic, updating an existing case base with a new case requires only that the priority relation derived from the existing case base should be strengthened with that derived from the new case, without any modification of existing priorities at all; and the representation of a single decision will remain the same, regardless of the case base in which it occurs.

This is one consideration, then, that favors taking the reason model of precedential constraint as fundamental, rather than the standard model, and rather than placing the two models on equal footing: only the reason model provides us with a modular account of the information carried by a case base, allowing for a simple treatment of case base update, and also for an atomistic view of the meaning of a legal case. But I do not mean for my argument

to hinge only on technical concerns of this kind. Instead, I want to show, in the remainder of this paper, that the reason model of precedential constraint suggests a new picture of common law reasoning more broadly, and also that it is able to deflect some important criticisms of the standard model.

6 Constraint and freedom

One of the things that makes common law reasoning so difficult to understand is that it appears to slip between two familiar models of judicial decision making, with familiar advantages and disadvantages.

The first of these is decision making based on what I will call, following Larry Alexander and Emily Sherwin, the *serious rule model*—decision making, that is, based on rules that cannot be modified once they are introduced, but must simply be applied as stated.[18] Of course, even decision making with serious rules is not entirely unproblematic. The predicates involved in these rules will require interpretation, a process that may itself involve a form of case-based reasoning. And there are difficulties posed by gaps and gluts: at times, a decision may need to be reached in situations in which it appears that no rule at all is applicable; at other times, multiple rules, supporting conflicting results, may apply in the same situation. But at least in situations in which some rule is applicable—conflict aside, and modulo interpretation—the serious rule model of decision making is one according to which results are determined by the rules alone.

This form of decision making offers several advantages, which have been discussed at length by a number of writers, and of which I mention only a few of the most central here.[19] It is, first of all, particularly simple, involving nothing more than a straightforward

[18]The term is due to Alexander and Sherwin (2009).

[19]For further discussion, and a guide to the literature, see especially Alexander and Sherwin (2009) and

application of rules, and so leading to the advantage of *efficiency*.[20] It possesses, in addition, the advantage described by Melvin Eisenberg as *replicability*, according to which individuals who are affected by the judgments of decision makers can replicate the reasoning of these decision makers.[21] Many other virtues follow from replicability. For example, as Eisenberg notes, those who can replicate, and so understand, the reasoning of decision makers are in a better position to appreciate their competence, and so more likely to comply with the resulting decisions; or in cases of incompetence, those who can replicate the underlying reasoning are better able to question that reasoning at the appropriate points.

Most important, however, replicability implies a degree of *predictability*, since a model according to which individuals can replicate a court's reasoning in previous cases is likely to be one in which they can, in the same way, predict its reasoning in future cases. And if these individuals are able to anticipate the decisions a court might reach in future cases, they can plan their actions accordingly, leading to the further advantage of *social coordination*. Imagine, for example, that Jack is considering the construction of a shopping center on his

Schauer (1991).

[20]Raz (1979, pp. 181–182) seems to question this advantage, arguing that even the straightforward application of rules can be difficult: he defines a "regulated dispute" as one governed by rules whose application does not require interpretation, claims that "regulated cases can be complex and more difficult to decide than unregulated cases," and illustrates this claim by noting that the "difficulty in solving a complex tax problem according to law may be much greater than that of solving a natural justice problem according to moral principles." I agree with this, of course—regulatory problems can be made arbitrarily complex. But it is important to bear in mind that the difficulties we find here are of a special sort, reflecting our own information processing limitations as much as anything else. In the field of artificial intelligence and law, where the focus is on machines with a very different pattern of cognitive limitations, pure rule-based reasoning is relatively unproblematic even in complex regulatory domains; see, for example, Bench-Capon (1991) and Schild and Herzog (1993).

[21]See Eisenberg (1988, pp. 10–12); the idea is also discussed by Lamond (2005, p. 7).

property, while, at the same time, Jo is searching for a site for her vacation home; and suppose that decisions about what can be built where are determined by a set of serious zoning rules. Then, since reasoning with serious rules is replicable, and so predictable, Jack can apply these rules to the situation at hand to conclude that a shopping center is allowed, or his representatives can do so.[22] Jack can therefore proceed with construction, without worrying that a neighbor might be able to convince a court that it is best, all things considered, to halt the project. And Jo, applying the same rules, will be able to predict the same result—that she cannot convince a court halt Jack's construction—and conclude, therefore, that she should not buy property adjacent to Jack, unless she is willing to accept the possibility of a shopping center next door. Coordination is thus achieved, with minimal judicial involvement.

Decision making with serious rules, then, has these advantages—efficiency, replicability, predictability, social coordination—as well as many others. Indeed, its sole disadvantage seems to be that, by screening off from consideration all features of a particular situation, no matter how important, apart from those that trigger the application of existing rules, this form of decision making can lead, at times, to suboptimal results. We are all familiar with situations in which, for example, an important and deserved benefit is denied because of a minor violation of rules—perhaps a form was filed containing a trivial error, or a correct form filed just slightly past deadline.[23] Here, the direct application of rules designed to

[22]Eisenberg emphasizes (1988, p. 11) that replicability, and so predictability, is what allows for so many legal issues to be resolved by a professional class of lawyers, taking pressure off the judicial system.

[23]Schauer (2009, p. 10) discusses the case of *United States v. Locke*, 471 U.S. 84 (1985), in which the Supreme Court decided that an individual should be denied a benefit because he filed a form on December 31 while the relevant statute read that the form should be filed "prior to December 31"—even though everyone seemed to agree that the actual language of the statute reflected a drafting error, and that what Congress had intended to say was that the form should be filed "on or prior to December 31."

promote the bureaucratic goals of order and efficiency interferes with achievement of the different, but arguably more important, goals of beneficence and equity. The suboptimality of decision making based on serious rules is even more striking when their direct application interferes with the achievement, not of different goals, even of more important goals, but of the very goals that the system of rules was introduced to advance—as when, for example, speed limits are put in place to assure that traffic flows smoothly and with minimal risk, but I find, in some situation, that I can both improve traffic flow and dramatically reduce risk by adjusting my speed to that of the surrounding vehicles, thus joining my fellow drivers in breaking the law.[24]

Returning to the balance between constraint and freedom, it is clear that the emphasis of the first of our two models of decision making, the serious rule model, falls entirely on the side of constraint, with no freedom at all granted to decision makers to adjust rules in order to avoid suboptimal outcomes in particular situations. The justification for this model, of course, lies in the hope that the benefits arising from a uniform application of serious rules will outweigh the costs of an occasional suboptimal decision.

The second of the familiar model of decision making considered here is one I will call, again following Alexander and Sherwin, the *natural model*—though I have also heard this form of decision making described as pure reason-based decision making, as all-things-considered decision making, or as decision making that is open-ended, unconstrained, or particularistic. An agent reaching a decision in accord with the natural model, in some particular situation, will proceed simply by surveying all the reasons that seem to bear on that situation, assigning these reasons the weights, or priorities, they seem to deserve, and then reaching whatever results these reasons together with their assigned priorities seem to support.

[24]The difficulties pointed out in this paragraph are, of course, standard objections to one version of rule utilitarianism.

Two immediate comments are in order. It is apparent, to begin with, that this description of the natural model contains some important gaps, since there are no clear, generally accepted answers to the questions of how the range of reasons bearing on some situation is to be surveyed, how priorities are be assigned to those reasons, or how reasons along with their priorities are supposed to support the results they do. Although I have explored answers to some of these questions elsewhere, I propose, in this paper, simply to live with the gaps, taking the idea of decision making purely on the basis of reasons as sufficiently well-understood to ground the natural model—this is, after all, how most of us make most of the decisions we do.[25] Second, we must note that, among the features of a situation that might be relevant to a decision, is the existence of certain rules. If there are rules, there may also be expectations that these rules will be followed; and both the rules and the resulting expectations may lead to reasons for one decision or another. There is nothing in the natural model that prevents rules from being taken into account in this way. All the natural model requires is that the rules themselves do not determine the resulting decision. Instead, the facts that there are rules, and that these rules may lead to expectations, constitute reasons, which can then be assigned priorities and weighed along with other reasons in supporting one decision or another. Even if the resulting decision then conforms to the rules, it is the reasons generated by these rules that determine the decision, not the rules themselves.

The central advantage of the natural model is that it is guaranteed to lead to an optimal decision. How do we know this? By stipulation, or very nearly so. When we say that the direct application of serious rules leads to a suboptimal result in some situation, what we mean by this is that it leads to a result different from that we would have endorsed if we had considered all the reasons bearing on that situation, with each assigned its proper weight,

[25]My own account of how priorities are assigned to reasons, and of how reasons along with their priorities support the results they do, can be found in Horty (2012).

or priority, and arrived at a decision on that basis. But of course, the natural model of decision making simply is the model according to which decisions are arrived at through the consideration of all relevant reasons, with each assigned its proper weight. This model, therefore, can be taken as defining what it means for an outcome to be optimal.

The disadvantage of the natural model is that, in allowing for full consideration of all reasons bearing on a particular situation, it sacrifices the advantages associated with the model of decision making based on serious rules. The form of reasoning recommended by the natural model is not particularly efficient: while the direct application of rules is straightforward, reflection on reasons, their priorities, and the decisions they support can be slow and agonizing. Nor is this form of reasoning, by and large, replicable: different individuals facing the same situation may well identify different reasons as relevant, and even if they agree on the relevant reasons, they are likely to assign them different priorities, so that different outcomes will be supported when these reasons are considered along with their priorities.[26] Since those affected by the judgments of decision makers will no longer be able to replicate the reasoning underlying these judgments, they will have less confidence in the judgments themselves, and will be less likely to comply; and if they wish to question the underlying reasoning, they will find it more difficult to do so in a useful way.

Most important, however, if the reasoning underlying a court's previous decisions cannot be replicated, it is unlikely that the court's reasoning in future situations could be predicted either—with the result that individuals will no longer be able to anticipate judicial decisions with any degree of certainty, and the advantages of social coordination will be lost. Recalling our earlier example, just imagine, if you can, a system in which decisions about what can be

[26]If we accept the account proposed in Horty (2012) of the way in which reasons support outcomes, there could be a further source of indeterminacy, since one version of this account allows that even the same set of reasons, with the same assignment of priorities, might at times lead to different, and conflicting, results.

built where were made, not in accord with a system of serious zoning regulations, but by a number a different courts, each assigning different weights, or priorities, to the various reasons in play. The result would be chaos, or paralysis. Jack could never begin building his shopping center without worrying that a court that happens to place great weight on environmental concerns, or on the rights of Jack's neighbors to enjoy their property, might force him to halt construction. Jo could never buy property for her vacation home without worrying that a court favoring commercial development would allow construction of a shopping center next door.

It is plain that, in the balance between constraint and freedom, the emphasis of the natural model lies entirely on the side of freedom—the freedom to give proper consideration to all reasons bearing on a particular situation, along with their priorities, in order to reach an optimal decision. Courts, on this model, are not constrained at all by the rules articulated in previous decisions, except to the extent that those rules may themselves provide reasons, to be balanced against others in reaching a judgment. Any justification for the natural model would have to involve the claim that, by reasoning about cases in an unconstrained and fully particular way, courts would then be able to reach decisions of high enough quality to compensate for the resulting loss of predictability and social coordination.

Now, as I suggested earlier, the correct view of common law reasoning often appears to lie between the two models of judicial decision making presented here, the serious rule model and the natural model, with their attendant advantages and disadvantages—often appearing to allow more freedom than the serious rule model, but to enforce more constraint than the natural model.[27] And there are, in the literature, two reactions to this suggestion. The

[27] Alexander (1989, p. 28) makes this point by alluding to a children's story, writing of readers presented only with these two models that, "Like Goldilocks and the bowls of porridge and beds, they will complain that the natural model of precedent is too weak to capture their sense of how precedents operate and that

first is to deny that there is, in fact, any defensible middle ground lying between these two familiar models of decision making, so that the theorist is forced to assimilate common law reasoning to one of these models or the other, with no further options. This hard-headed position is advocated most forcefully by Alexander and Sherwin, who go on to argue that it is best for everyone, both theorists and practitioners, to understand common law reasoning in accord with the model of reasoning with serious rules.[28] The second reaction to the idea that common law reasoning appears to lie in a middle ground between our familiar models is, of course, to try to define this middle ground, and argue that it is defensible. This project has been pursued by a number of writers, in different ways; but a representative, and very attractive, proposal can be found in Frederick Schauer's "presumptive positivism."[29] According to this proposal, common law reasoning proceeds in accord with the model of serious rules in the vast run of cases, even in cases in which the direct application of serious rules leads to moderately suboptimal outcomes. There is thus a strong presumption in favor of rule application, and so constraint. The exception, according to Schauer, is that, in cases in which suboptimality threatens to become extreme, the application of serious rules can then be discarded in favor of reasoning in accord with the natural model, allowing courts the freedom to avoid the most egregious outcomes.

My own account—reflected in the reason model—is different. It is not an attempt to locate common law reasoning somewhere in the territory between the model of decision making based on serious rules and the natural model; indeed, there is no appeal to rules at all. Instead, I view decision making in the common law as entirely reason-driven, just

the rule model of precedent is too strong."

[28]See Alexander (1989) throughout, and then Alexander and Sherwin (2001, pp. 136–156) and (2009, pp. 27–127).

[29]See Schauer (1989, p. 117n, pp. 196–206) and (1991, pp. 469–471).

like the form of decision making set out in the natural model, but with the sole difference that a common law decision maker, constrained by precedent, must adapt his or her own priority ordering on reasons so that it coheres with the priority ordering derived from the background case base.

To describe this proposal more precisely, recall that, according to the natural model, an agent's decision in some situation depends on two things: first, the reasons that the agent sees as bearing on that situation, and second, the weights, or priorities, that the agent assigns to those reasons. We can simplify by imagining that the reasons bearing on some situation X are clear, and let us suppose that $<_i$ is the priority ordering among those reasons as assigned by the agent i—so that, according to the natural model, it is this priority ordering that would guide the agent's decision. On my view, then, all that differs when the agent is a common law decision maker, reasoning about the situation under precedential constraints derived from a background case base Γ, is that the agent's own priority ordering $<_i$ on reasons must be revised to cohere with the ordering $<_\Gamma$ derived from the background case base—leading to, let us say, $<_{i/\Gamma}$ as a revised ordering—and that it is this revised ordering, rather than the original $<_i$, that guides the agent's decision.

How is the revised ordering on reasons $<_{i/\Gamma}$ to be determined, given the agent's original ordering $<_i$ together with the ordering $<_\Gamma$ derived from the background case base? We can assume, as an idealization, that the revised ordering must be consistent—that we cannot have both $W <_{i/\Gamma} Z$ and $Z <_{i/\Gamma} W$, for reasons W and Z. And if the agent's reasoning is actually supposed to be constrained by the derived priority ordering $<_\Gamma$, then it is natural to require also that the revised ordering $<_{i/\Gamma}$ should extend $<_\Gamma$—that we should have $W <_{i/\Gamma} Z$ whenever $W <_\Gamma Z$, so that the revised ordering tells us that Z has higher priority than W whenever this relation can be derived from the background case base, no matter how the

agent's original ordering might have ranked these reasons.

Beyond these two conditions, I believe there is very little of a systematic nature to be said. It may be tempting, from a conservative perspective, to imagine that the revised ordering $<_{i/\Gamma}$ should represent some minimal modification of the agent's original ordering $<_i$—or more exactly, that $<_{i/\Gamma}$ should result from combining the ordering $<_\Gamma$ derived from the background case base with some maximal subset of the agent's original ordering $<_i$ that is consistent with $<_\Gamma$. But I can think of no justification for such a strong requirement. Of course, a particular agent might take such a resolute, unyielding stance toward its own ordering on reasons that the agent is unwilling to accept any modifications at all, apart from those strictly necessary for reconciling this ordering with that derived from the background case base. But it is also possible for an agent to adopt a more open-minded, or receptive, attitude. The agent might follow a procedure something like that sketched by Dworkin, for example—first extrapolating from the actual decisions contained in the background case base to the best moral theory that explains those decisions, and then, in light of this theory, modifying its own priority ordering on reasons in ways that may go well beyond those strictly necessary for incorporating the ordering derived from the case base into its own. Which attitude the agent adopts, and how, exactly, the agent's original ordering on reasons is modified to cohere with that derived from the case base might depend on a number of variables—including psychological facts about the agent, structural facts about the relation between the agent's original ordering and that derived from the case base, and substantive facts about the nature of the reasons under consideration.

However it is accomplished, then, let us suppose that, in deliberating about the situation X under precedential constraints provided by the background case base Γ, the agent i has managed to revise its own original ordering on reasons $<_i$ so that it coheres with the ordering

$<_\Gamma$ derived from the case base, leading to the new ordering $<_{i/\Gamma}$. My suggestion, once again, is simply that, to reach a decision in the situation X, the agent should be thought of as reasoning in the natural way about this situation, just as in the natural model, except with the priority ordering on reasons given by the revised $<_{i/\Gamma}$ rather than the original $<_i$. Suppose that, reasoning in this way, the agent arrives at a decision in favor of the side s, say, on the basis of the rule r. The case base Γ will then, as we have seen, be supplemented with this new decision, leading to the richer case base $\Gamma' = \Gamma \cup \{\langle X, r, s\rangle\}$, so that the next judicial agent j, confronting the next situation Y, will be forced to revise its initial ordering on reasons $<_j$ with the stronger ordering $<_{\Gamma'}$ derived from this new case base. This agent will, therefore, base its deliberations about the situation Y on the revised priority ordering $<_{j/\Gamma'}$, rather than its own original $<_j$, eventually reaching some decision, updating the case base accordingly, and the process will continue. At each stage, the case base will be supplemented, and the derived priority ordering strengthened, in a way that depends on the particular decision maker's own ordering among the reasons bearing on the situation at hand, but only after this ordering has been revised to cohere with the priority ordering derived from the existing case base.

What of the balance between constraint and freedom, and of the respective advantages and disadvantages associated with the serious rule model and the natural model? Unlike the rule model, which provides constraint without freedom, and the natural model, which allows freedom without constraint, the current suggestion seems to offer an appropriate balance. In considering some new situation against the background of an existing case base, a decision maker is, first of all, constrained by the requirement that the reasons bearing on this situation must be evaluated, not in accord with the priorities that the decision maker would naturally assign to them, but instead, in accord with a priority ordering that has been

revised to cohere with that derived from the existing case base. Once this requirement has been satisfied, however, the decision maker is then free to engage in an open-ended process of deliberation that brings all reasons bearing on that situation into play, and that proceeds in the natural way.

Although there is this balance between constraint and freedom, it is important to see that the balance shifts as the case law in some area is developed. The priority relation among reasons endorsed by an individual agent i is likely to be relatively complete—since people are opinionated—so that the ordering $<_i$ is likely to be strong. But at least initially, as the law in an area is just beginning to be developed, the background case base Γ will contain very few decisions, so that the derived priority ordering $<_\Gamma$ will be weak. As a result, this derived priority ordering will have relatively little impact on the revised ordering $<_{i/\Gamma}$, so that reasoning in accord with this revised ordering will approximate reasoning in accord with the agent's original ordering $<_i$. At this early stage of legal development, the model of reasoning described here will share the advantages and disadvantages of the natural model: common law decision makers, even reasoning under the constraints of precedent, will have a good deal of freedom to reach solutions they consider to be optimal, but their reasoning will not necessarily be predictable by those affected by these decisions, and certain advantages of social coordination will be sacrificed.[30]

As the law is developed, so that the case base Γ becomes more and more populated with decisions, the ordering $<_\Gamma$ derived from this case base will grow increasingly stronger. Because the revised priority ordering $<_{i/\Gamma}$ must extend this derived ordering, it follows that the derived ordering will have an increasing influence on the nature of the revised ordering, so

[30]In the extreme case, of course, where the law in some area is entirely undeveloped, so that the case base Γ is empty, and so $<_\Gamma$ empty as well, the revised ordering $<_{i/\Gamma}$ will coincide with the original ordering $<_i$, for each agent i, and the model of reasoning described here will coincide with the natural model.

that reasoning in accord with this revised ordering will diverge more and more significantly from reasoning in accord with the agent's original ordering $<_i$. In these later stages of legal development, the model of reasoning developed here will come to share the characteristics of the serious rule model: common law decision makers, reasoning under increasingly strong precedential constraints, will have less freedom to reach decisions they view as optimal, but their decisions will become more predictable, supporting social coordination.[31]

7 Objections

Although the reason model of precedential constraint, summarized in Section 4 of this paper, is equivalent to the standard model developed in Section 3, I have argued that the reason model should be taken as fundamental, both because it offers a modular representation of the information contained in a case base and because it supports an attractive picture of common law reasoning. In this final section, I show how the reason model helps us understand two features of the common law that may seem puzzling, or objectionable, when viewed only from the perspective of the standard model. These two features are the distinction between distinguishing and overruling previous decisions, and the constraining effect of precedent.

The distinction between distinguishing and overruling is generally taken as central to the common law. All courts are thought to have the power of distinguishing previous decisions, through which they carry out a process of gradual, incremental, adaptive legal development. Overruling a previous decision, by contrast, is viewed as a more radical operation, generally available only to courts either above or, sometimes, at the same level as that which decided

[31] Considering the extreme case once again, if we now imagine that the case base Γ contains enough decisions that every possible comparison between reasons is already settled by $<_\Gamma$, it follows that the revised ordering $<_{i/\Gamma}$ will have to coincide with the derived priority ordering $<_\Gamma$, so that the agent will be entirely constrained, and so predictable, in its reasoning, with no freedom remaining at all.

the case to be overruled. Even then, this option is avoided whenever possible, since the resulting legal transformations can be abrupt and extreme: when a precedent case is overruled, it is as if the case were completely "wiped off the slate," or removed "root and branch."[32]

The distinction itself is best illustrated with a concrete example, such as our initial example from this paper. The example centered around a situation in which Laura and Ron have two children: Emma, age nine, who did not finish dinner but did completed her homework, and Max, age fourteen, who has neither finished dinner nor completed his homework. Both children would like to stay up and watch TV. We imagined that Emma first asked Laura, who granted the request, justifying her decision with the rule, "Children age nine or greater can stay up and watch TV." Next, we imagined, Max asked Ron, who denied the request, distinguishing this case from that of Emma by appeal to the fact that Max failed to complete his homework—thus, both introducing the new rule "Children who have not completed their homework cannot stay up and watch TV" and also modifying Laura's previous rule to read, "Children age nine or greater who have completed their homework can stay up and watch TV." Since Ron's modification of Laura's rule satisfies the Raz/Simpson conditions—merely narrowing the rule, and doing so in a way that continues to support the previous decision—it can be taken as a legitimate case of distinguishing.

This scenario can be represented in our framework by taking f_1^π as the factor that a child is over nine years old, and then f_1^δ and f_2^δ as the respective factors that the child failed to finish dinner and homework. The initial situation confronting Laura, then, was $X_{11} = \{f_1^\pi, f_1^\delta\}$, which she decided on the basis of the rule $r_{11} = \{f_1^\pi\} \rightarrow \pi$, leading to the case base $\Gamma_5 = \{c_{11}\}$, containing only the single case $c_{11} = \langle X_{11}, r_{11}, s_{11}\rangle$, where X_{11} and r_{11} are as above, and where $s_{11} = \pi$. Next, Ron was confronted by the situation $X_{12} = \{f_1^\pi, f_1^\delta, f_2^\delta\}$,

[32]The first phrase is due to Cross (1968, p. 119), who attributes it to Lord Dunedin; the second is due to Raz (1979, p. 189).

which he decided on the basis of the rule $r_{12} = \{f_2^{\delta}\} \rightarrow \delta$—both supplementing the existing case base with the new case $c_{12} = \langle X_{12}, r_{12}, s_{12} \rangle$ where X_{12} and r_{12} are as above and where $s_{12} = \delta$, and also modifying Laura's original rule to carry the force $r_{11}{}' = \{f_1^{\pi}\} \wedge \neg\{f_2^{\delta}\} \rightarrow \pi$. The new case base is thus $\Gamma_5{}' = \{c_{11}{}', c_{12}\}$, with $c_{11}{}' = \langle X_{11}{}', r_{11}{}', s_{11}{}' \rangle$ as a modification of the previous c_{11}, where $X_{11}{}' = X_{11}$, where $r_{11}{}'$ is as above, and where $s_{11}{}' = s_{11}$; and with c_{12} as above.

Suppose, however, that Ron had disagreed with Laura's original decision, which downplays the significance of failing to finish dinner, and chose to use the case of Max to reaffirm the importance of this factor. We can then imagine that, rather than proceeding as in the initial scenario, Ron had instead chosen to justify his decision with the new rule "Children who have not finished their dinner cannot stay up and watch TV," represented here as $r_{13} = \{f_1^{\delta}\} \rightarrow \delta$, and so leading to the new case $c_{13} = \langle X_{13}, r_{13}, s_{13} \rangle$, where $X_{13} = X_{12}$, where r_{13} is as above, and where $s_{13} = \delta$. This new rule would no longer satisfy the Raz/Simpson conditions—it is neither a narrowing of Laura's original rule, nor does it support the previous decision—and so Ron would now have to be taken, not simply as distinguishing, but as overruling Laura's decision. How can this operation be modeled in the present framework? If an overruled case is indeed to be "wiped from the slate," then it is natural to suppose that one logical effect of this operation is that the overruled case should be removed from the case base entirely.[33] On this view, Ron can be seen as both supplementing the background case

[33]There may be other logical effects as well. Perhaps, in overruling a case, a court should be taken as removing from the case base, not only that particular case, but every other case that shares the same rule; or perhaps there is a temporal dimension, so that the court should be taken as removing every other case sharing the same rule as the original that was decided at a later date. Overruling can be a complex operation, but there is no need to consider its complexities here, since our example contains only a single case to be overruled.

base with his new decision and then, not modifying, but simply removing Laura's previous decision, leading to

$$\begin{aligned} \Gamma_5'' &= (\Gamma_5 - \{c_{11}\}) \cup \{c_{13}\} \\ &= \{c_{13}\} \end{aligned}$$

as the updated case base, with c_{13} as above.

From an intuitive standpoint, this example seems to support the standard distinction between distinguishing and overruling a previous decision. It does seem, for example, that the rule set out by Ron in the second scenario represents a much more radical challenge to Laura's decision than Ron's rule in the initial scenario; and we can understand why, in a stable legal system, only certain courts should be able to challenge earlier decisions in such a radical way. Nevertheless, in spite of the intuitive force of the distinction between distinguishing and overruling, Alexander and Sherwin—writing from the standpoint of the serious rule model—argue that this distinction is both confused and dangerous. On confusion, they write:

> the rule model does not and cannot distinguish between overruling precedent rules and modifying or "distinguishing" them. When a judge makes an exception to a rule to accommodate a particular case, the judge is effectively eliminating the precedent rule and announcing a new rule in its place.[34]

And on danger:

> The practice of distinguishing precedent rules is dangerous to the stability of rules because it creates an illusion of modesty. Judges may intervene more often

[34] Alexander and Sherwin (2009, pp. 114–115; see also pp. 59, 84). Schauer voices a similar concern when he argues that a rule that can be modified at the moment of application is "in an important way not a rule at all" (1991, p. 117).

> when they believe they are merely modifying, rather than overruling, established rules. This belief is mistaken because modifying or distinguishing precedent rules just *is* overruling them.[35]

When they are viewed from the perspective of the standard model, I believe these objections to the distinction between distinguishing and overruling make good sense. To begin with, we must agree with Alexander and Sherwin that, even in instances of legal development typically classified as distinguishing, where the Raz/Simpson conditions are satisfied, a later court is not, strictly speaking, modifying an earlier rule at all, but instead, removing that rule from consideration and introducing one or more new rules; these new rules may have more or less similarity to the original, but they are nevertheless different rules, exhibiting different syntactic structures and yielding different results in a variety of situations, including the case at hand. In the first of our two scenarios, for example, Ron removes Laura's original rule $r_{11} = \{f_1^{\pi}\} \rightarrow \pi$ and introduces the new rules $r_{11}{}' = \{f_1^{\pi}\} \wedge \neg\{f_2^{\delta}\} \rightarrow \pi$ and $r_{12} = \{f_2^{\delta}\} \rightarrow \delta$; in the second, Ron again removes Laura's original rule and introduces the new rule $r_{13} = \{f_1^{\delta}\} \rightarrow \delta$. In each scenario, then, Laura's original rule is eliminated entirely, and none of the new rules support the same result as the original in the new fact situation presented by Max. Why, then, should we think of distinguishing as any less radical than overruling—why should we think that Ron's challenge to Laura's original decision in the first scenario is any less radical than his challenge in the second?

This problem for the standard model, with its emphasis on rules, has a happy solution when the matter is viewed from the perspective of the reason model, which allows a clear semantic distinction between distinguishing and overruling to be drawn in terms of the priority ordering on reasons derived from a background case base. Since both decisions that

[35] Alexander and Sherwin (2009, p. 124).

distinguish and decisions that overrule change the case base, both kinds of decisions change the derived priority ordering as well; but they do so in strikingly different ways. A court that distinguishes a previous decision merely expands the existing case base with a new decision, with the result that the derived priority ordering on reasons is strengthened. But a court that overrules a previous decision both expands the existing case base with a new decision and contracts it through the removal of a previous case, with the result that the derived priority ordering is strengthened in some ways but weakened in others, and so incomparable to the original.

This point can be illustrated by returning to our two scenarios. In the first scenario, when Ron distinguishes Laura's earlier decision, moving from the original case base Γ_5 to the new case base $\Gamma_5{}'$, the derived priority ordering is strengthened: it is easy to see that $W <_{\Gamma_5} Z$ implies $W <_{\Gamma_5{}'} Z$ for any reasons W and Z, and the new ordering yields $\{f_1^\pi\} <_{\Gamma_5{}'} \{f_2^\delta\}$ while the original ordering did not support the corresponding $\{f_1^\pi\} <_{\Gamma_5} \{f_2^\delta\}$. But in the second scenario, when Ron overrules Laura's earlier decision, now moving from the original case base Γ_5 to the new $\Gamma_5{}''$, the resulting derived priority ordering is incomparable to the original: it is stronger in some ways, since we have $\{f_1^\pi\} <_{\Gamma_5{}''} \{f_1^\delta\}$ but not $\{f_1^\pi\} <_{\Gamma_5} \{f_1^\delta\}$, but it is also weaker in some ways, since we have $\{f_1^\delta\} <_{\Gamma_5} \{f_1^\pi\}$ but not $\{f_1^\delta\} <_{\Gamma_5{}''} \{f_1^\pi\}$. Overruling is thus both a constructive and a destructive operation, adding new information to the existing priority ordering but also removing information that is already present. The operation of distinguishing can likewise appear to be destructive from the perspective of the standard model, since it involves the removal of existing rules; but from the perspective of the reason model, we can see that this operation is entirely constructive, merely strengthening the existing priority ordering.

I now turn to a second objection to the standard model: that, as long as the rules

set out by courts can be distinguished—even if the modifications involved are required to satisfy the Raz/Simpson conditions—common law decisions can have no real constraining effect on future courts at all, since there will always be features available for these courts to use in distinguishing the situations they face from those confronted earlier. Again, this objection is set out forcefully by Alexander and Sherwin, who illustrate the problem with their story of the ocelot and the alligator.[36] We are asked to imagine that an earlier court has already considered the question whether a certain individual could keep an ocelot at home and arrived at a negative decision, justifying this decision with the rule, "Wild animals in residential neighborhoods are nuisances"—where we can assume it is already settled that animals can be kept at home just in case they are not nuisances. A later court now faces the question whether another individual can keep an alligator. This court is sympathetic to the alligator, wishes to arrive at a positive decision in the case at hand, but is aware that it must distinguish the previous rule in order to do so. The court therefore notes that ocelots but not alligators are furry, and proceeds to distinguish on that basis, modifying the previous rule to read, "Furry wild animals in residential neighborhoods are nuisances," and, we might as well suppose, justifying its decision with the new rule, "Animals without fur are not nuisances."

By modifying the earlier rule in this way, the later court has rendered it inapplicable to the case of the alligator, giving itself the freedom to reach whatever decision it wishes. And as Alexander and Sherwin point out, this instance of rule modification satisfies the Raz/Simpson conditions, merely narrowing the previous rule, and narrowing it in such a way that the modified rule continues to support the decision arrived at in the previous case. The example thus highlights the fact that any two cases can be differentiated in any number

[36] See Alexander and Sherwin (2009, pp. 84–86).

of ways, even if many of these differences are only incidental—that one dangerous wild animal but not the other is furry, for example, or that the defendant in one case but not the other has freckles, or plays the harmonica, or has an aunt living in Idaho. And if all a court needs to do in order to shield the decision it wishes to reach from some previous rule is to narrow the rule by appeal to one of these incidental differences, then it really is hard to see, from the perspective of the standard model alone, how the decisions reached in earlier cases constrain later decisions.

But let us look at the example from the standpoint of the reason model, focusing in particular on the account of common law reasoning accompanying this model. Suppose the court considering the case of the alligator reasons in accord with this account—that is, suppose the court deliberates through a process very much like natural reasoning, but with the weight it would naturally assign to certain reasons modified to cohere with the priority ordering derived from the background case base, including the case of the ocelot. And imagine that the court, reasoning in this way, and reasoning sincerely, really does conclude that the alligator's lack of fur is a reason against classifying it as a nuisance, and indeed a stronger reason than that provided for the opposite conclusion by the fact that the alligator is a wild animal. In that case, I believe it would be right, at least from an internal perspective, for the court to reach exactly the decision described in the example—that the previous case should be distinguished, and the alligator allowed because it has no fur. The court, after all, has an obligation to reach the decision it sincerely thinks is best, taking into account both the reasons bearing on this situation and the priorities it sees among these reasons, once these priorities have been adjusted to cohere with those derived from the background case base.

What is so odd about this scenario, and what gives the example its force, is not some

problem with the idea of distinguishing, but simply the assumption that the court might actually conclude, in all sincerity, that this particular decision is best—that any court, reasoning in the natural way, could ever conclude that lack of fur is a consideration that bears on the situation at all, let alone a reason strong enough to outweigh important considerations favoring the other side. Surely any such court would be criticizable, in the same way that those who engage in poor natural reasoning in any other domain are criticizable. There is also the suggestion in Alexander and Sherwin's presentation of the example that the court, by introducing lack of fur as a reason, is being disingenuous. But this would be criticizable as well—not in the way that poor natural reasoning is criticizable, but in the way that we might criticize a person who is misleading or dismissive about a matter that should be taken seriously. And it is by focusing on this general idea of criticizability, I think, that we can locate a response to Alexander and Sherwin's objection concerning the possibility of constraint: earlier courts constrain later courts, not by preventing these later courts from reaching certain decisions, but by limiting the resources available to these courts for arriving at or justifying decisions in ways that are not criticizable.

This suggestion can be illustrated with a slightly more elaborate version of the original example. So suppose that, in both the initial case of the ocelot and the later case of the alligator, there are two considerations with real bearing on the question whether the animal at hand is a nuisance: both animals are wild, but both are kept in sturdy pens. And also of course, there are a number of incidental factors: one animal has fur while the other does not, for example. It is natural to imagine that, in deciding whether the ocelot should be classified as a nuisance, the initial court weighed the inherent danger of keeping a wild animal at all against the security derived from keeping that animal in a sturdy pen. Either decision would have been understandable, but we can assume, as in the original scenario, that the

court eventually concluded that the ocelot should be classified as a nuisance because it is a wild animal. As a result of this decision, the priority ordering derived from the case base will—according to the reason model—contain the information that the property of being wild carries more weight in favor of the conclusion that an animal is a nuisance than the property of being kept in a sturdy pen carries for the opposite conclusion.

Now against this background, how should the court reach its decision in the later case of the alligator? Just as before, there are good reasons favoring each conclusion: the animal is wild, but it is kept in a sturdy pen. And we might even suppose that the judge in this case, reasoning as an individual, would assign greater weight to the latter of these considerations—so that, if the initial question of the ocelot had come before this court, it would have been found not to be a nuisance, since it is kept in a sturdy pen. Nevertheless, that question has now been decided for the other side, and as a result of this decision, it has now been established that the case base supports the opposite priority relation on the relevant reasons. According to our account of common law reasoning, therefore, the judge, deliberating in an official capacity at least, must now revise his or her own individual priority relation to cohere with that derived from the background case base, so that being wild will be assigned greater weight than being kept in a sturdy pen.

The later court will, therefore, no longer able to arrive at, or justify, a decision that the alligator is not a nuisance on the basis of the reasonable consideration that it is kept in a sturdy pen, since the consideration that it is a wild animal, supporting the other side, has already been given greater weight. Suppose, then, that the court reaches, or justifies, its decision in favor of the alligator by appealing instead to an incidental feature of the situation, such as the fact that the alligator lacks fur. In that case, as we have seen, there are two possibilities. If the court reaches this conclusion sincerely—if it really does assign

this consideration sufficient weight to override the danger posed by the alligator as a wild animal—then the court can be criticized for its poor natural reasoning, and for its very odd prioritization among reasons. On the other hand, if the court realizes that its revised priority ordering or reasons no longer provides any real basis for its desired conclusion, but introduces lack of fur as a consideration supporting this conclusion in a misleading or disingenuous way, then the court can be criticized on other grounds. In either case, the earlier decision regarding the ocelot deprives the later court of resources for reaching, or justifying, the conclusion that the alligator is not a nuisance in a reasonable way, leaving open only paths to this conclusion that are more questionable, and more easily criticized.

From the perspective of the standard model, then, it may appear that the constraint of common law is minimal, since a court can distinguish a previous rule on the basis of any consideration at all, as long as the modification of that rule satisfies the purely formal Raz/Simpson conditions. But the reason model allows us to see that there is more to it than that. Whenever a court distinguishes a previous rule, it offers, at the same time, a claim about the priority relation among reasons—that lacking fur is a more important consideration than being a dangerous wild animal, for example—which is itself subject to evaluation.[37] Common law constrains because each decision settles the priority relations among certain reasons, and so limits the ability of later courts to claim otherwise. After a sufficient number of decisions, the priorities among all the important reasons in some domain will be well enough understood that a later court can then distinguish an earlier rule only by offering further claims of priority that are unlikely to withstand evaluation.

[37]How are claims about priority relations among reasons to be evaluated? Two different accounts of the way in which these claims can be evaluated by appeal to further reasons are offered in Schroeder (2007, pp. 123–145) and Horty (2012, pp. 111–121); the latter explores an example (pp. 119–121) in which priority relations among legal reasons, in particular, are established by appeal to further legal reasons.

A Table of cases

1. $c_1 = \langle X_1, r_1, s_1 \rangle$, where $X_1 = \{f_1^\pi, f_2^\pi, f_3^\pi, f_1^\delta, f_2^\delta, f_3^\delta, f_4^\delta\}$, where $r_1 = \{f_1^\pi, f_2^\pi\} \wedge \neg\{f_5^\delta\} \rightarrow \pi$ and where $s_1 = \pi$.

2. $c_2 = \langle X_2, r_2, s_2 \rangle$, where $X_2 = \{f_1^\pi, f_1^\delta\}$, where $r_2 = \{f_1^\pi\} \rightarrow \pi$, and where $s_2 = \pi$.

3. $c_3 = \langle X_3, r_3, s_3 \rangle$, where $X_3 = \{f_1^\pi, f_2^\delta\}$, where $r_3 = \{f_2^\delta\} \rightarrow \delta$, and where $s_3 = \delta$.

4. $c_4 = \langle X_4, r_4, s_4 \rangle$, where $X_4 = \{f_1^\pi, f_3^\delta\}$, where $r_4 = \{f_3^\delta\} \rightarrow \delta$, and where $s_4 = \delta$.

5. $c_5 = \langle X_5, r_5, s_5 \rangle$, where $X_5 = \{f_1^\pi, f_2^\delta\}$, where $r_5 = \{f_1^\pi\} \rightarrow \pi$, and where $s_5 = \pi$.

6. $c_6 = \langle X_6, r_6, s_6 \rangle$, where $X_6 = \{f_1^\pi, f_1^\delta, f_2^\delta\}$, where $r_6 = \{f_1^\delta, f_2^\delta\} \rightarrow \delta$, and where $s_6 = \delta$.

7. $c_7 = \langle X_7, r_7, s_7 \rangle$, where $X_7 = \{f_2^\pi, f_2^\delta\}$, where $r_7 = \{f_2^\pi\} \rightarrow \pi$, and where $s_7 = \pi$.

8. $c_8 = \langle X_8, r_8, s_8 \rangle$, where $X_8 = \{f_1^\pi, f_2^\delta\}$, where $r_8 = \{f_2^\delta\} \wedge \neg\{f_2^\pi\} \rightarrow \delta$, and where $s_8 = \delta$.

9. $c_9 = \langle X_9, r_9, s_9 \rangle$, where $X_9 = \{f_1^\pi, f_2^\pi, f_1^\delta, f_2^\delta\}$, where $r_9 = \{f_1^\pi\} \rightarrow \pi$, and where $s_9 = \pi$.

10. $c_{10} = \langle X_{10}, r_{10}, s_{10} \rangle$, where $X_{10} = \{f_1^\pi, f_1^\delta, f_2^\delta, f_3^\delta\}$, where $r_{10} = \{f_1^\delta, f_2^\delta\} \rightarrow \delta$, and where $s_{10} = \delta$.

11. $c_{11} = \langle X_{11}, r_{11}, s_{11} \rangle$, where $X_{11} = \{f_1^\pi, f_1^\delta\}$, where $r_{11} = \{f_1^\pi\} \rightarrow \pi$, and where $s_{11} = \pi$.

12. $c_{12} = \langle X_{12}, r_{12}, s_{12}\rangle$ where $X_{12} = \{f_1^{\pi}, f_1^{\delta}, f_2^{\delta}\}$, where $r_{12} = \{f_2^{\delta}\} \rightarrow \delta$, and where $s_{12} = \delta$

13. $c_{13} = \langle X_{13}, r_{13}, s_{13}\rangle$, where $X_{13} = \{f_1^{\pi}, f_1^{\delta}, f_2^{\delta}\}$, where $r_{13} = \{f_1^{\delta}\} \rightarrow \delta$, and where $s_{13} = \delta$.

B Proof of Observation 1

Observation 1 Let Γ be a case base. Then Γ is reason consistent if and only if its refinement Γ^+ is itself a case base.

Proof The proof of this observation is divided into two parts.

Part I: If Γ is a reason consistent case base, then its refinement Γ^+ is a case base.

Proof of Part I: Suppose Γ is a reason consistent case base. Γ^+ is constructed from Γ by replacing each case $c = \langle X, r, s\rangle$ from Γ with the new $c'' = \langle X, r'', s\rangle$, where the new rule r'' has the form $Premise^s(r) \wedge D_c \rightarrow s$, as specified as in Definition 1. Since all of the new rules involved in moving from Γ to Γ^+ support the same outcomes as the original, we can verify that Γ^+ is a case base as well simply by establishing that, for each $c'' = \langle X, r'', s\rangle$ from Γ^+, the new rule r'' continues to be applicable to the fact situation X—that is, that $X \models Premise(r'')$, or that $X \models Premise^s(r) \wedge D_c$. We know, of course, that $X \models Premise^s(r)$, since Γ is a case base, and so need only show that $X \models D_c$.

It follows from Steps 2 and 3 of the construction that establishing that $X \models D_c$ amounts to showing, for each $c' = \langle Y, r', \overline{s}\rangle$ from Γ_c, where $c = \langle X, r, s\rangle$, that $X \models \neg\, Premise^{\overline{s}}(r')$. So suppose the contrary—that $X \not\models \neg\, Premise^{\overline{s}}(r')$, or $X \models Premise^{\overline{s}}(r')$, from which we can conclude that (1) $Premise^{\overline{s}}(r') \subseteq X^{\overline{s}}$. Since $c' = \langle Y, r', \overline{s}\rangle$ belongs to Γ_c, we know from Step 1 of the construction that $Y \models Premise^s(r)$, from which we can conclude that (2) $Premise^s(r) \subseteq Y^s$. From (1), we can then conclude by Definition 3 that (3) $Premise^{\overline{s}}(r') <_c$

Premise$^s(r)$, and from (2), that (4) *Premise*$^s(r) <_{c'}$ *Premise*$^{\overline{s}}(r')$. But since both c and c' belong to Γ, the combination of (3) and (4) contradicts the stipulation that Γ is reason consistent. Hence, our assumption fails, from which we can conclude that $X \models D_c$.

Part II: If Γ is a case base whose refinement Γ^+ is also a case base, then Γ is reason consistent.

Proof of Part II: Suppose Γ is a case base whose refinement Γ^+ is a case base, but that Γ itself is not reason consistent. Since Γ is not reason consistent, there are reasons A and B such that (1) $A <_c B$ and (2) $B <_{c'} A$ for cases $c = \langle X, r, s\rangle$ and $c' = \langle Y, r', \overline{s}\rangle$ from Γ. From (1) we have (3) $A \subseteq X^{\overline{s}}$ and (4) *Premise*$^s(r) \subseteq B$, and from (2) we have (5) $B \subseteq Y^s$ and (6) *Premise*$^{\overline{s}}(r') \subseteq A$. Together, (4) and (5), along with the fact that $Y^s \subseteq Y$, yield *Premise*$^s(r) \subseteq Y$, or (7) $Y \models$ *Premise*$^s(r)$. In the same way, (3) and (6), together with the fact that $X^{\overline{s}} \subseteq X$, yield *Premise*$^{\overline{s}}(r') \subseteq X$, or (8) $X \models$ *Premise*$^{\overline{s}}(r')$.

Γ^+ is constructed from the case base Γ by replacing each case $c = \langle X, r, s\rangle$ with the new $c'' = \langle X, r'', s\rangle$, where the new rule r'' has the form *Premise*$^s(r) \wedge D_c \rightarrow s$, as specified in Definition 1. Step 1 of this construction, together with (7), tells us that c' belongs to Γ_c, and then Steps 2, 3, and 4 allow us to conclude that $\neg$ *Premise*$^{\overline{s}}(r')$ is one of the conjuncts of D_c, and so of the new rule r''. From (8), however, we know that $X \models$ *Premise*$^{\overline{s}}(r')$, from which it follows that $X \not\models \neg$ *Premise*(r''). As a result, the rule of c'' does not apply to its facts, from which it follows that c'' is not a case, and so Γ^+ not a case base, contrary to our assumption. ∎

References

Aleven, Vincent (1997). *Teaching Case-Based Argumentation Through a Model and Examples*. PhD thesis, Intelligent Systems Program, University of Pittsburgh.

Aleven, Vincent and Ashley, Kevin (1997). Evaluating a learning environment for case-based argumentation skills. In *Proceedings of the Sixth International Conference on Artificial Intelligence and Law (ICAIL-97)*, pages 170–179. The Association for Computing Machinery Press.

Alexander, Larry (1989). Constrained by precedent. *Southern California Law Review*, 63:1–64.

Alexander, Larry and Sherwin, Emily (2001). *The Rule of Rules: Morality, Rules, and the Dilemmas of Law*. Duke University Press.

Alexander, Larry and Sherwin, Emily (2009). *Demystifying Legal Reasoning*. Cambridge University Press.

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

Bench-Capon, Trevor (1991). Practical legal expert systems: the relation between a formalisation of law and expert knowledge. In Bennun, M. and Narayanan, A., editors, *Computers, Law, and AI*, pages 191–201. Ablex Publishers.

Brachman, Ronald and Levesque, Hector, editors (1985). *Readings in Knowledge Representation*. Morgan Kaufmann Publishers.

Brachman, Ronald and Levesque, Hector, editors (2004). *Knowledge Representation and Reasoning*. Morgan Kaufmann Publishers.

Cross, Rupert (1968). *Precedent in English Law*. Oxford University Press. Second edition.

Dancy, Jonathan (2004). *Ethics Without Principles*. Oxford University Press.

Dworkin, Ronald (1975). Hard cases. *Harvard Law Review*, 99:1057–1109. Reprinted in Dworkin (1977), pp. 81–130.

Dworkin, Ronald (1977). *Taking Rights Seriously*. Harvard University Press.

Eisenberg, Melvin (1988). *The Nature of the Common Law*. Harvard University Press.

Holton, Richard (2010). The exception proves the rule. *The Journal of Political Philosophy*, 18:369–388.

Holton, Richard (2011). Modelling legal rules. In *Philosophical Foundations of Language in the Law*, pages 165–183. Oxford University Press.

Horty, John (2011). Rules and reasons in the theory of precedent. *Legal Theory*, 17:1–33.

Horty, John (2012). *Reasons as Defaults*. Oxford University Press.

Horty, John and Bench-Capon, Trevor (2012). A factor based definition of precedential constraint. *Artificial Intelligence and Law*, 20:181–214.

Lamond, Grant (2005). Do precedents create rules? *Legal Theory*, 11:1–26.

Leiter, Brian (2005). American legal realism. In Golding, Martin and Edmundson, William, editors, *The Blackwell Guide to the Philosophy of Law and Legal Theory*, pages 50–66. Blackwell Publishing Company.

Levi, Edward (1949). *An Introduction to Legal Reasoning*. The University of Chicago Press.

McCarty, L. Thorne (1997). Some arguments about legal arguments. In *Proceedings of the Sixth International Conference on Artificial Intelligence and Law (ICAIL-97)*, pages 215–224. The Association for Computing Machinery Press.

Prakken, Henry and Sartor, Giovanni (1998). Modelling reasoning with precedents in a formal dialogue game. *Artificial Intelligence and Law*, 6:231–287.

Raz, Joseph (1979). *The Authority of Law*. Oxford University Press.

Rissland, Edwina (1990). Artificial intelligence and law: stepping stones to a model of legal reasoning. *Yale Law Journal*, 99:1957–1981.

Rissland, Edwina and Ashley, Kevin (1987). A case-based system for trade secrets law. In *Proceedings of the First International Conference on AI and Law (ICAIL-87)*, pages 60–66. The Association for Computing Machinery Press.

Schauer, Frederick (1989). Is the common law law? *California Law Review*, 77:455–471.

Schauer, Frederick (1991). *Playing by the Rules: A Philosophical Examination of Rule-Based Decision-Making in Law and Life*. Oxford University Press.

Schauer, Frederick (2009). *Thinking Like a Lawyer*. Harvard University Press.

Schild, Uri and Herzog, Shai (1993). The use of meta-rules in rule based legal computer systems. In *Proceedings of the Fourth International Conference on AI and Law (ICAIL-93)*, pages 100–109.

Schroeder, Mark (2007). *Slaves of the Passions*. Oxford University Press.

Simpson, A. W. B. (1961). The *ratio decidendi* of a case and the doctrine of binding precedent. In Guest, A. G., editor, *Oxford Essays in Jurisprudence*, pages 148–175. Oxford University Press.