

Rules and Reasons in the Theory of Precedent

John F. Horty

Philosophy Department and
Institute for Advanced Computer Studies

University of Maryland

College Park, MD 20742

horty@umiacs.umd.edu

www.umiacs.umd.edu/users/horty

*** DRAFT ***

Comments Welcome

Version of: May 16, 2008

Contents

1	Introduction	1
2	Basic concepts	5
2.1	Factors, rules, and cases	5
2.2	Ordered considerations and binding rules	10
3	The theory	17
3.1	Constraint by precedent	17
3.2	Case base dynamics	22
4	Comparisons	31
4.1	The result model	31
4.2	Distinguishing	38
5	Conclusion	48
A	Cases and fact situations	54
B	Observations and proofs	56

1 Introduction

The doctrine of precedent, as it has evolved within the common law, is constituted by a system of conventions through which the decisions of earlier courts in particular cases somehow generalize to constrain the decisions of later courts facing different cases, while still allowing these later courts a degree of freedom in responding to fresh circumstances. Although the techniques for reasoning with precedents are taught early on in law schools, mastered with relative ease, and applied on a daily basis by legal practitioners, it has proved to be considerably more difficult to arrive at a theoretical understanding of the doctrine itself, a clear articulation of the underlying system of conventions.

My purpose in this paper is to describe a new framework within which we can address this problem. I concentrate on two of the most fundamental questions in the theory of precedent. First, how is it, exactly, that precedent cases constrain future decisions—what is the mechanism of constraint? And second, how is a balance then achieved between the constraints of precedent and the freedoms allowed to later courts for developing the law.

The view presented here will be contrasted with three other views, or models, of precedential constraint appearing in the literature. The first is the *rule model*. A precedent case normally contains, not only a description of the facts of the case along with a decision on the basis of those facts, but also some particular rule through which that decision was reached. According to the rule model, it is this rule that carries the precedential constraint. Constraint by precedent just is constraint by rules; a precedent case constrains the decision of a later court when the rule contained in that precedent applies to the fact situation confronting the later court.

A particularly strict version of the rule model has been advanced by Larry Alexander, and developed by Alexander and Emily Sherwin.¹ According to this strict rule model, the constraints carried by precedent rules are very nearly absolute, with the result that the freedoms given to later courts for developing the law are extremely limited. When a precedent rule applies to a later fact situation, then a court confronting that fact situation has, at most, two choices: the court must either follow the precedent rule, or, if that court has the authority to do so, it can overrule the precedent. There is, however, nothing in the rule model itself—the general idea that precedential constraint is carried through rules—that forces this very strict interpretation. A number of writers, while still adhering to this general idea, have argued for a more flexible version of the rule model according to which later courts have the power to develop the law by modifying, without overruling, the rules contained in precedent cases. The process through which these earlier rules might be modified, or *distinguished*, as well as the constraints on this process of distinguishing a precedent rule, are described with particular care by A.W.B. Simpson and Joseph Raz.²

The idea that precedential constraint is carried through rules surely reflects the popular conception of the matter, and at least the more flexible version of the rule model, according to which precedent rules can be distinguished, is perhaps the received view among legal theorists. Still, two further views deserve our attention. The second model of precedential

¹See Alexander (1989), many of the arguments from which are summarized in Alexander (1996); see also Alexander and Sherwin (2001). The term “rule model” is itself due to Alexander, who applies it only to what I describe here as the strict rule model, which reflects his own preferred view; I use the term more broadly to apply to any model in which precedential constraint is thought to be carried by rules.

²See Simpson (1961) and Chapter 10 of Raz (1979).

constraint to be considered here is the *result model*.³ According to this model, what matters about a precedent case is not the rule it contains, but, first of all, the result of the case, and second, the strength of that case for its result. Precedential constraint is then thought to be a simple matter of a fortiori reasoning: a later court is constrained to follow the ruling of a precedent case when the facts confronting the later court are at least as strong for the winning side of the precedent case as the facts of the precedent case itself. The result model of precedent depends, of course, on some definite way of measuring the strength of a case for one side or another. The idea that there might be such a measure has been criticized by Alexander as unattractive, and perhaps incoherent, but I have defended the idea as coherent, at least, and one that holds some attractions, even if it does not tell the whole story.⁴

The third view of precedential constraint to be considered here, recently introduced by Grant Lamond, may be termed the *reason model*.⁵ According to Lamond, what is most important about an earlier court's decision in a precedent case is, not the rule it contains, nor even the strength of the precedent case for one side or another, but instead, the earlier court's assessment of the balance of reasons presented by the facts of that case. The requirement of precedential constraint can then be defined as follows: unless it wishes to overrule the precedent, a later court is constrained to reach a decision that is consistent with the earlier court's assessment of the balance of reasons. My own view, as we will see, makes crucial use of this idea.

A central advantage of the account to be presented here is that it shows how these

³The term is again due to Alexander (1989).

⁴See Horty (2004); Alexander and Sherwin respond to some of my arguments in their (2007).

⁵See Lamond (2005).

three theoretical models of precedent—rule, result, and reason—can be, in a precise sense, unified; it helps us see what is correct in each of these views, and how they are related. The unification between the rule model and the reason model is achieved by interpreting the rules contained in precedent cases, not as strict rules, but as default, or defeasible, rules, while reasons are then defined as the premises of default rules. This analysis of reasons as the premises of default rules, which we will return to shortly, is not at all unnatural, nor is it ad hoc, introduced here only to establish a connection between two theoretical models of legal precedent; it is something that I have defended on independent grounds elsewhere.⁶ The unification between the rule model and the result model is achieved even more simply, by showing that the rule model is actually a generalization of the result model, or put another way, that the result model is simply a special case of the rule model, in which precedent rules are required to exhibit a certain form.

The key innovation of the present account is that it makes explicit what is generally only implicit in case law: a priority ordering representing the strength of the reasons—or more broadly, the considerations—underlying judicial decisions. Like the set of rules contained in precedent cases, this priority ordering on considerations is itself taken to be a part of the law, although, like the precedent rules themselves, the priority ordering is derived from the decisions reached in precedent cases, not defined independently. As it turns out, once this priority ordering on considerations has been made explicit, various aspects of precedential reasoning can then be simplified in substantial ways. The notion of precedential constraint itself can be defined by reference to the underlying priority ordering on considerations, and

⁶See Horty (2007).

likewise the effects of following or distinguishing a rule.

The paper is organized as follows. The next section presents basic ideas and notation. The third section presents the core theory of precedential constraint and then moves on to a discussion of case base dynamics—the effects of following or distinguishing a precedent rule. The fourth section relates the account presented here to two previous models of precedent, the result model and the rule model, and includes a comparison between the current picture of distinguishing and the traditional notion. The entire account relies, of course, on a number of simplifying assumptions; some of these are discussed in the conclusion.

2 Basic concepts

2.1 Factors, rules, and cases

Throughout the paper, I follow the work of Kevin Ashley and his colleagues in supposing that the situation presented to the court in a legal case can usefully be represented as a set of *factors*, where a factor represents a legally significant fact or pattern of facts.⁷ Cases in different areas of the law will be characterized by different sets of factors, of course. 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

⁷See Ashley (1989) and (1990) for an introduction to the model; see also, Rissland (1990) for an overview of research in Artificial Intelligence and Law that places this work in a broader context.

between the plaintiff and the defendant, whether the information acquired was reverse-engineerable or in some other way publicly available, and the extent to which this information did, in fact, lead to a real competitive advantage for the defendant.⁸

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 factors they present.

Many factors can naturally be taken to have polarities, favoring one side or another. In the domain of trade secrets law, again, the presence of security measures 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 than moving through a series of intermediate legal concepts. Both of

⁸Aleven (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). Of course, 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.

these assumptions will be reexamined toward the end of the paper.

Formally, then, let us 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 $F^\pi = \{f_1^\pi, \dots, f_n^\pi\}$ represent the set of factors favoring the plaintiff and $F^\delta = \{f_1^\delta, \dots, f_m^\delta\}$ 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: $F = F^\pi \cup F^\delta$.

A *precedent case* will be represented as a fact situation together with an outcome as well as 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 present in the case, r is the rule of the case, and s is its outcome. We define three functions—*Facts*, *Rule*, and *Outcome*—to map cases into their component parts, so that, in the case c above, for example, we would have $Facts(c) = X$, $Rule(c) = r$, and $Outcome(c) = s$.

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 \bar{s} represents a decision for the opposite side, so that $\bar{\pi} = \delta$ and $\bar{\delta} = \pi$. Where X is a fact situation, we let X^s represent the factors from X that support the side s ; that is, $X^\pi = X \cap F^\pi$ and $X^\delta = X \cap F^\delta$.

The *rule* r contained in a precedent case has the form $Y \rightarrow s$, where Y is some set of factors supporting s as an outcome. We define two functions—*Premise* and *Conclusion*—picking out the premise and the conclusion of a rule, so that, in the case of this particular

rule r , for example, we would have $Premise(r) = Y$ and $Conclusion(r) = s$. A precedent rule of this sort, once again, is to be interpreted as a defeasible rule, telling us that its premise entails its conclusion, not as a matter of necessity, but only by default. What the rule $Y \rightarrow s$ means, then, is that, if some fact situation contains all the factors from Y , then as a default, the court ought to reach a decision in this situation in favor of the side s —or perhaps more intuitively, that the factors from Y , taken together, provide the court with a reason for deciding in favor of the side s .

This connection between precedent rules and reasons—a guiding theme of the paper—can be illustrated by examining a different sort of normative rule, an ethical generalization, such as “If you make a promise, you ought to keep it.” Consider an instance of this generalization, such as “If I promise to have lunch with Alex, I ought to do so,” and suppose I have, in fact, promised to have lunch with Alex, so that the rule is applicable. What, then, is the force of this rule? It cannot mean that I ought to have lunch with Alex no matter what. Surely other, more important obligations might legitimately interfere; I might be called upon to save a life, for example. Instead, it is natural to interpret the rule as telling us that my promise, the premise of the rule, provides me with a *reason* for having lunch with Alex—presumably a very strong reason, since it is based on a promise, but still a reason that might be defeated by stronger reasons, or perhaps excluded from consideration entirely.⁹

The idea behind the current account is that precedent rules work in exactly the same

⁹This general picture is described in much more detail in Horty (2007), where I appeal to techniques from nonmonotonic logic to develop a detailed theory as reasons as the antecedents of default rules; the theory allows us to understand both how reasons can be defeated by stronger reasons and how they can be excluded from consideration, in the way suggested by Raz (1975).

way, so that the *legal reasons* present in the common law will simply be defined here as the premises of precedent rules. What the rule $Y \rightarrow s$ tells us, then, is that the factor set Y provides the court with a legal reason for deciding in favor of the side s . Just as in the case of ethical generalizations, however, the reason provided by this precedent rule may be defeated by a stronger reason favoring the opposite side, in a way that will be explained shortly.

Let us return, now, to the concept of a precedent case $c = \langle X, r, s \rangle$, containing a fact situation along with a rule r leading to the outcome s . In order for this concept to make sense, we impose three coherence constraints. First, the rule contained in the case must be applicable to the facts of the case, in the sense that the fact situation contains the factors required by the premise of the rule: $Premise(r) \subseteq X$. Second, each of the factors contained in the premise of the precedent rule must actually support its conclusion, not the opposite side: where $Conclusion(r) = s$, then, we require $Premise(r) \subseteq F^s$. And third, the conclusion of the precedent rule must match the outcome of the case: $Conclusion(r) = Outcome(c)$.

These various concepts and constraints can be illustrated through the concrete case $c_1 = \langle X_1, r_1, s_1 \rangle$, containing the fact situation $X_1 = \{f_1^\pi, f_2^\pi, f_3^\pi, f_1^\delta, f_2^\delta, f_3^\delta, f_4^\delta\}$, with three factors favoring the plaintiff and four favoring the defendant, where r_1 is the rule $\{f_1^\pi, f_2^\pi\} \rightarrow \pi$, and where the outcome s_1 is π , a decision for the plaintiff. Evidently, the case satisfies our three coherence constraints. The precedent rule is applicable to the fact situation, in the sense that $Premise(r_1) \subseteq X_1$. The various factors contained in the premise of this rule all support its conclusion, a decision for the plaintiff, in the sense that $Premise(r_1) \subseteq F^\pi$. And the conclusion of the precedent rule matches the outcome of the case, both favoring the plaintiff: $Conclusion(r_1) = Outcome(c_1)$. This particular precedent, then, represents a case in which the court decided for the plaintiff through the application or introduction of a rule according

to which the presence of the factors f_1^π and f_2^π lead, by default, to a decision for the plaintiff.

2.2 Ordered considerations and binding rules

With this notion of a precedent case in hand, we can now define a *case base* as a set Γ of precedent cases. It is a case base of this sort that is supposed to represent the common law in some area, and to constrain the decisions of future courts. But according to the present theory, these constraints depend more immediately on two additional concepts, both of which can be defined in terms of the case base.

The first is simply the set of rules derived from a case base, which is definable quite simply as the set containing any rule belonging to any case from that case base. The concept can be introduced formally by extending the function *Rule*, which extracts the rule from a single case, so that it applies also to an entire set of cases, yielding as a result the set of rules contained in those cases.

Definition 1 (Rules derived from a case base) Let Γ be a case base. Then the set $Rule(\Gamma)$ of rules derived from Γ is defined by taking $Rule(\Gamma) = \{Rule(c) : c \in \Gamma\}$.

To illustrate, suppose the case base Γ contains the case c_1 , considered above. Then the set $Rule(\Gamma)$ of rules derived from this case base will contain the particular rule r_1 , since this rule is the value of $Rule(c_1)$ and c_1 belongs to Γ .

It is customary, of course, to suppose that the rules derived from a case base play an important role in precedential constraint; some writers argue that these precedent rules play the entire role. The next concept we introduce—a preference relation on reasons, or considerations—is much less common as a focus of attention.

In order to motivate this concept, it will be useful to consider our previous example, the case c_1 , in more detail. The key idea underlying precedential constraint is that later courts must respect the decisions of earlier courts. So what information is actually carried by the earlier court's decision in the case of c_1 ; what is the earlier court telling us with this decision? Well, two things, at least. First of all, by appealing to the rule r_1 , the court is telling us that the premise of this rule—that is, $Premise(r_1)$, or $\{f_1^\pi, f_2^\pi\}$ —provides a reason for reaching a decision in favor of the plaintiff. But second, with its decision for the plaintiff, the court is also telling us that this reason outweighs the various factors present in the case that favor the defendant.

Let us define a *consideration* as a set of factors favoring one side or another, so that the class of considerations includes the class of reasons—premises of precedent rules—but is broader; not every consideration need be the premise of some precedent rule. The case c_1 contains the set $X_1 = \{f_1^\pi, f_2^\pi, f_3^\pi, f_1^\delta, f_2^\delta, f_3^\delta, f_4^\delta\}$ as its fact situation, and so the strongest consideration presented by this case for the defendant is the subset $X_1^\delta = \{f_1^\delta, f_2^\delta, f_3^\delta, f_4^\delta\}$, containing all those factors from the original fact situation that favor the defendant. Since the earlier court has decided for the plaintiff on the grounds of the reason provided by $Premise(r_1)$, even in the face of the consideration provided by X_1^δ for the defendant, it seems to follow as a consequence of the court's decision that the reason, or consideration, $Premise(r_1)$ for the plaintiff is stronger than the consideration X_1^δ for the defendant—that is, that the consideration $\{f_1^\pi, f_2^\pi\}$ outweighs the consideration $\{f_1^\delta, f_2^\delta, f_3^\delta, f_4^\delta\}$. If we introduce the symbol $<_{c_1}$ to represent the preference relation on considerations derived from the particular case c_1 , then this consequence of the court's decision can be put more formally as the claim that $X_1^\delta <_{c_1} Premise(r_1)$, or equivalently, that $\{f_1^\delta, f_2^\delta, f_3^\delta, f_4^\delta\} <_{c_1} \{f_1^\pi, f_2^\pi\}$.

As far as the preference ordering goes, then, the earlier court is telling us at least that $X_1^\delta <_{c_1} \text{Premise}(r_1)$, but is it telling us anything else? Perhaps not explicitly, but implicitly, yes. For if the consideration provided by $\text{Premise}(r_1)$ for the plaintiff outweighs the consideration provided by X_1^δ for the defendant, then surely any consideration for the plaintiff that is at least as strong as $\text{Premise}(r_1)$ must likewise outweigh X_1^δ , and just as surely, any consideration for the defendant that is at least as weak as X_1^δ must be outweighed by $\text{Premise}(r_1)$. It is natural to suppose that a consideration Z for the plaintiff is at least as strong $\text{Premise}(r_1)$ if it contains all the factors from $\text{Premise}(r_1)$ and possibly others, also favoring the plaintiff—that is, if $\text{Premise}(r_1) \subseteq Z$. And it is natural to suppose that a consideration W for the defendant is at least as weak as X_1^δ if it contains no more factors than X_1^δ itself—that is, if $W \subseteq X_1^\delta$. It therefore follows from the earlier court’s decision in c_1 , not only that $X_1^\delta <_{c_1} \text{Premise}(r_1)$, but that $W <_{c_1} Z$ whenever $W \subseteq X_1^\delta$ and $\text{Premise}(r_1) \subseteq Z$. To illustrate: it follows from the court’s decision, not only that $\{f_1^\delta, f_2^\delta, f_3^\delta, f_4^\delta\} <_{c_1} \{f_1^\pi, f_2^\pi\}$, but also that $\{f_3^\delta, f_4^\delta\} <_{c_1} \{f_1^\pi, f_2^\pi, f_5^\pi\}$, for example.

This line of reasoning leads to the following definition of the preference, or strength, relation among considerations that can be derived from a single case.

Definition 2 (Preference relation derived from a case) Let $c = \langle X, r, s \rangle$ be a case, and suppose W and Z are considerations. Then the relation $<_c$ representing the preferences on considerations derived from the case c is defined by stipulating that $W <_c Z$ if and only if $W \subseteq X^s$ and $\text{Premise}(r) \subseteq Z$.¹⁰

¹⁰From our coherence constraints on rules and cases we have $\text{Premise}(r) \subseteq F^s$. It therefore follows from the stipulation that $\text{Premise}(r) \subseteq Z$ together with the requirement that Z is a consideration—that is, a set of factors all favoring a single side—that $\text{Premise}(r) \subseteq Z \subseteq F^s$. Thus, W is a consideration lying between

And once we have defined the preference relation derived from a single case, we can then introduce the preference relation $<_{\Gamma}$ derived from an entire case base Γ in the natural way, by stipulating that one consideration is stronger than another according to the entire case base if that strength relation is supported by some particular case in the case base.

Definition 3 (Preference relation derived from a case base) Let Γ be a case base, and suppose W and Z are considerations. Then the relation $<_{\Gamma}$ representing the preferences on considerations 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 .

It is worth noting explicitly that the derived preference relation $<_{\Gamma}$ is very weak, both formally and conceptually. Formally, it is not even transitive: $X <_{\Gamma} Y$ and $Y <_{\Gamma} Z$ does not entail $X <_{\Gamma} Z$. Indeed, quite the opposite. For it is easy to see from our various definitions that, whenever $X <_{\Gamma} Y$, the two considerations X and Y must lie on opposite sides of some dispute, one favoring the plaintiff while the other favors the defendant. Hence, given $X <_{\Gamma} Y$ and $Y <_{\Gamma} Z$, we can conclude that X and Z , both lying opposed to Y , must themselves favor the same side, from which it follows that $X <_{\Gamma} Z$ fails.

And the preference relation is conceptually weak as well, because it might reflect very few of our ordinary judgments about strength relations among considerations. Consider, for example, a situation in which the issue at hand is the question whether an individual's residence in a foreign country qualifies as a change of fiscal domicile with respect to income tax.¹¹ The plaintiff is the individual's home country, which would like to collect tax on his the empty set and the entire subset $X^{\bar{s}}$ of factors from X that favor \bar{s} , while Z is a consideration lying between $Premise(r)$ and the entire universe F^s of factors that favor s .

¹¹This example is derived from some hypothetical cases considered by Prakken and Sartor (1998).

income; the defendant is the individual, who would prefer to pay his income taxes to the foreign country, where we can assume the rates are lower. Imagine that the fact situation contains the following factors, all favoring the defendant: the individual resigned from his old job and is now employed by a company in the foreign country; he has sold his old home and purchased a new home in the foreign country; he has sold his old car and both purchased and registered a new car in the foreign country. Suppose these three factors are represented as f_1^δ , f_2^δ , and f_3^δ . In favor of the plaintiff is the single factor that the individual has maintained a registered bicycle in his home country, which he uses while visiting his parents; this factor is f_1^π .

Can we now assume that the consideration $\{f_1^\delta, f_2^\delta, f_3^\delta\}$ favoring the defendant outweighs the consideration $\{f_1^\pi\}$ favoring the plaintiff? Intuitively, it would seem so—surely information about employment, residence, and automobile registration should outweigh bicycle registration. But as a matter of law, not necessarily. Unless the case base Γ contains a previous case in which bicycle registration was actually weighed against at least one of the factors supporting the defendant, we will not have $\{f_1^\pi\} <_\Gamma \{f_1^\delta, f_2^\delta, f_3^\delta\}$. The present approach thus reflects a positivist view according to which the legally sanctioned preference relations among considerations, as much a part of the law as precedent rules themselves, must have a basis in the acts of some appropriate legal authority—here, a court’s decision in some precedent case. Of course, it is likely in the present situation that the court confronting this case would be guided by our intuitive assessment concerning strength of the conflicting considerations and therefore decide for the defendant. As a result of this decision, the intuitive assessment would be given legal standing, and it would then hold, once the case is decided, that $\{f_1^\pi\} <_\Gamma \{f_1^\delta, f_2^\delta, f_3^\delta\}$. This is the genius of the common law—that it pro-

vides a mechanism through which our ordinary intuitions about the relative importance of various considerations are gradually filtered into the legal doctrine, on an incremental basis, in reaction to particular circumstances.

We now turn to the task of defining the class of precedent rules that should be considered as *binding* in a particular fact situation, those with the greatest bearing on that fact situation, which any court confronting that fact situation must take into account. The definition is simple, and proceeds in three steps.

First of all, a rule is said to be *applicable* in a fact situation whenever that situation contains all the factors required by the premise of the rule.

Definition 4 (Applicable rules) Let Γ be a case base, with $Rule(\Gamma)$ the derived set of rules, and suppose X is a fact situation. Then a rule r from $Rule(\Gamma)$ is applicable in the fact situation X if and only if $Premise(r) \subseteq X$.

Since our precedent rules are taken as defeasible, however, not every applicable rule can be classified as binding. Some will be overridden—or *trumped*—by stronger rules supporting the opposite side.

What makes one precedent rule stronger than another? This, again, is a matter of law, not intuition. The force of a precedent rule, we recall, is that the premise of that rule provides the court with a reason for deciding in favor of the side specified in its conclusion. Precedent rules can therefore be ranked in strength along with the reasons they provide, so that a rule r' is taken to be stronger than the rule r in the context of a case base Γ whenever $Premise(r')$ is stronger than $Premise(r)$ according to the preference relation $<_{\Gamma}$ derived from that case base—whenever, that is, $Premise(r) <_{\Gamma} Premise(r')$.

Given this notion of strength, we can now define an applicable rule as trumped whenever there is another rule, also applicable, that is stronger than the original and supports the opposite side.

Definition 5 (Trumped rules) Let Γ be a case base, with $Rule(\Gamma)$ the derived set of rules and $<_{\Gamma}$ the derived preference relation, and suppose X is a fact situation. Then a rule r from $Rule(\Gamma)$ that is applicable in X is trumped in the context of the case base Γ just in case there is another rule r' from $Rule(\Gamma)$ that is also applicable in X , but which is such that (1) $Premise(r) <_{\Gamma} Premise(r')$ and (2) $Conclusion(r') = \overline{Conclusion(r)}$.

And once we have defined both the applicable and the trumped rules, we can introduce the idea of a binding rule as one that is applicable but not trumped.

Definition 6 (Binding rules) Let Γ be a case base, with $Rule(\Gamma)$ the derived set of rules and $<_{\Gamma}$ the derived preference relation, and suppose X is a fact situation. A rule r from $Rule(\Gamma)$ is binding in X just in case it is triggered in the fact situation X and not trumped in the context of Γ .

These concepts can be illustrated by considering the simple case base $\Gamma_1 = \{c_1, c_2\}$ containing only these two precedent cases, with c_1 as before and with $c_2 = \langle X_2, r_2, s_2 \rangle$, where the fact situation $X_2 = \{f_1^{\pi}, f_2^{\pi}, f_4^{\pi}, f_5^{\delta}, f_6^{\delta}\}$, where the rule r_2 is $\{f_5^{\delta}\} \rightarrow \delta$, and where the outcome s_2 is δ , a decision for the defendant. Now suppose that, against the background of this case base, a new fact situation $X_3 = \{f_1^{\pi}, f_2^{\pi}, f_5^{\pi}, f_5^{\delta}, f_7^{\delta}\}$ comes before the court. The set $Rule(\Gamma_1)$ of derived precedent rules contains two rules, r_1 and r_2 ; and evidently, both of these two rules are applicable in the current situation, since we have both $Premise(r_1) \subseteq X_3$ and $Premise(r_2) \subseteq X_3$. The first of these rules, however, is trumped by the second. The

two rules favor different sides, of course, with r_1 favoring the plaintiff and r_2 the defendant. And it is easy to see also that, in the context of Γ_1 as a background case base, the reason provided by the second rule is stronger than that provided by the first: it follows from our definitions that $\{f_1^\pi, f_2^\pi\} <_{c_2} \{f_5^\delta\}$, from which we have $\{f_1^\pi, f_2^\pi\} <_{\Gamma_1} \{f_5^\delta\}$ as well—that is $Premise(r_1) <_{\Gamma_1} Premise(r_2)$. Because both r_1 and r_2 are applicable, but r_1 is trumped, only the rule r_2 is binding in this fact situation.

3 The theory

3.1 Constraint by precedent

The account of precedential constraint set out here can be seen as a version of the reason model, advocated by Lamond, according to which a later court is constrained to reach a decision that is consistent—not necessarily with the rules set out in earlier cases, or even with their results—but instead, with the assessments reached in these cases concerning the proper balance of reasons. In order to develop this idea, I first introduce a reason-centered notion of consistency for case bases. A later decision can then be defined as consistent with the precedents contained in a case base if it does not introduce an inconsistency into that case base.

As we have seen, a case base Γ leads to a derived preference relation $<_{\Gamma}$, where the statement $X <_{\Gamma} Y$ means that the consideration Y carries more weight than X , according to Γ . Such a statement is supported, of course, by some particular precedent case from Γ in which it was decided either explicitly that the consideration Y itself carries more weight than X , or else that some consideration at least as weak as Y carries more weight than some

consideration at least as strong as X , from which it follows implicitly that $X <_{\Gamma} Y$. We therefore define the case base Γ as *inconsistent* whenever there are two considerations X and Y for which both $X <_{\Gamma} Y$ and $Y <_{\Gamma} X$ —whenever, that is, Γ tells us both that Y is stronger than X and that X is stronger than Y —and *consistent* otherwise.

Definition 7 (Consistent and inconsistent case bases) Let Γ be a case base with $<_{\Gamma}$ the derived preference relation. Then Γ is inconsistent just in case there are considerations X and Y such that $X <_{\Gamma} Y$ and $Y <_{\Gamma} X$. Γ is consistent just in case it is not inconsistent.

Is this a good definition of case base inconsistency, and so consistency, from an intuitive point of view? I think so. The condition isolated by the definition is almost certainly sufficient with respect to our intuitive notion of inconsistency—surely any case base from which it can be derived that, of two considerations, each is more important than the other would have to be classified as inconsistent from an intuitive standpoint. But is the suggested condition also necessary? Perhaps a case base might exhibit some other anomaly that would lead us to classify it, from an intuitive standpoint, as inconsistent. Suppose, for example, that the case base contains two precedent cases of the form $\langle X, r, s \rangle$ and $\langle X, r', \bar{s} \rangle$ in which the very same fact situation leads to decisions for opposing sides; surely there is some kind of intuitive inconsistency in a case base like this. True enough, but as it turns out, this particular anomaly entails the formal condition set out in our definition to capture the intuitive notion of inconsistency, so that it cannot be used to challenge the claim that the formal condition is necessary.

Observation 1 Let Γ be a case base containing two precedent cases of the form $\langle X, r, s \rangle$ and $\langle X, r', \bar{s} \rangle$. Then Γ is inconsistent.

A different, slightly weaker anomaly will be shown later in the paper to entail our formal condition as well. I have not been able to find any others that do not, and will therefore take our formal definitions of consistency and inconsistency for a case base as intuitively acceptable.

Given this notion of consistency, then, we can now turn to the concept of precedential constraint itself. The guiding intuition is that, in deciding a case, a constrained court is required to preserve the consistency of the background case base. More exactly, where Γ is a consistent case base, suppose a court that is constrained by 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 into the case base.

Definition 8 (Precedential constraint) Let Γ be a consistent case base and X a new fact situation confronting the court. Then 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 itself consistent.

This notion of precedential constraint can be illustrated by returning to our previous example, in which $\Gamma_1 = \{c_1, c_2\}$ is the background case base, with c_1 and c_2 as before, and the court is confronted with the new fact situation $X_3 = \{f_1^\pi, f_2^\pi, f_5^\pi, f_5^\delta, f_7^\delta\}$. As we have seen, the rule r_2 , or $\{f_5^\delta\} \rightarrow \delta$, is the unique binding rule in this fact situation, so that, as far as precedent rules are concerned, the background case base unambiguously favors a decision for the defendant. And in many situations, there may indeed be a presumption that favors following a binding rule.¹² Still, on the view developed here, precedential constraint does

¹²See, for example, the discussion of “presumptive positivism” in Schauer (1989) and (1991).

not depend on binding rules, but instead, on consistency with the background case base.

The court, in this situation, would be free to follow the binding rule r_2 , leading to a decision for the defendant, and so augmenting the background case base with the new case $c_3 = \langle X_3, r_3, s_3 \rangle$, where the rule r_3 is simply r_2 and the outcome s_3 is δ , which would, in fact, preserve consistency. But the court is also free to decide, for example, that the new factor f_5^π , which favors the plaintiff, outweighs the considerations presented by this fact situation for the defendant. The court might then formulate a new rule $\{f_5^\pi\} \rightarrow \pi$, and on the basis of this rule, decide for the plaintiff. As a result, the background case base would be augmented with the new case $c_4 = \langle X_4, r_4, s_4 \rangle$, where X_4 is identical with X_3 , where r_4 represents this new rule, and s_4 is π . And again, the augmented case base resulting from this decision would remain consistent. The new case c_4 would allow us to derive, for example, the preference relation $\{f_5^\delta, f_7^\delta\} <_{c_4} \{f_5^\pi\}$, telling us that the consideration $\{f_5^\pi\}$ for the plaintiff carries more weight than the consideration $\{f_5^\delta, f_7^\delta\}$ for the defendant. But this new preference relation is consistent with the others already derivable from the background case base.

What the court cannot do, however—what precedential constraint rules out—is to find for the plaintiff on the basis of some consideration that does not include the new factor f_5^π . Suppose, for example, that the court formulates the rule $\{f_1^\pi\} \rightarrow \pi$, and wishes to decide for the plaintiff on the basis of this rule. Such a decision would result in an augmentation of the background case base with the new case $c_5 = \langle X_5, r_5, s_5 \rangle$, where X_5 is again identical with X_3 , where r_5 represents the new rule, and s_5 is π . But this augmented case base would now be inconsistent. The new case c_5 would support the preference relation $\{f_5^\delta\} <_{c_5} \{f_1^\pi\}$, telling us that the consideration $\{f_1^\pi\}$ for the plaintiff outweighs the consideration $\{f_5^\delta\}$ for the defendant. But the background case base already contains the case c_2 , from which we

can derive the preference relation $\{f_1^\pi\} <_{c_2} \{f_5^\delta\}$, telling us exactly the opposite. Since a decision for the plaintiff on these grounds would therefore lead to an inconsistent case base, it is ruled out by the present account of precedential constraint.¹³

It is worth acknowledging explicitly that the present theory of precedential constraint relies on the assumption that the background case base is itself consistent to begin with. This is, of course, an unrealistic assumption. Given the vagaries of judicial decision, with a body of case law developed by a number of different courts, at different places and different times, it would be surprising if any nontrivial case base were actually consistent. But in fact, this assumption is not essential. The notion of case base inconsistency at work here is not like logical inconsistency—it is local, not pervasive. A case base might be inconsistent in certain areas, providing conflicting information about the relative weight of particular considerations,

¹³As I have emphasized, the account proposed here can be seen as a version of the reason model advocated by Lamond in (2005) and elsewhere; however, there are at least two significant differences between Lamond’s discussion and the present account. First at all, at a conceptual level, Lamond sets out his reason model as an alternative to the conventional view of precedential constraint, with its emphasis on rules, while, on the present account, the roles of reasons and rules are integrated: reasons are identified with the antecedents of precedent rules; the preference ordering on rules is derived from the preference ordering on reasons. Second, Lamond is more concerned with advocating, rather than developing, the reason model—emphasizing the importance of reasons in precedential constraint. The present account, by contrast, focuses on the development of a particular, concrete version of the theory, rather than general advocacy; and the particular version of the theory developed here, furthermore, seems to differ from than envisioned by Lamond. For example, Lamond suggests (2005, pp. 18–19) that it might be best to understand the meaning precedents in terms of “protected reasons” and other forms of exclusionary reasons. Although initially sympathetic to this idea, I did not, in the end, find it necessary to appeal to exclusionary reasons in developing the present account.

while remaining consistent elsewhere. It would therefore be possible to extend the present theory of precedential constraint to apply also to inconsistent case bases, by requiring of a court, not necessarily that it should preserve the consistency of a consistent case base, but only that it should refrain from introducing any new inconsistencies, which were not present before, into a case base that may already be inconsistent. Nevertheless, simply in order to avoid additional complexity, this paper will be concerned only with situations in which the background case base is consistent.

3.2 Case base dynamics

I now want to consider how some of the most important concepts from the traditional theory of precedent—following a rule, distinguishing a rule, overturning a precedent—can be accommodated within the present framework. The goal is to understand the options available to the court under various circumstances, and also the way in which the case base and its associated constraints evolve when one of these options is selected—that is, the dynamics of case base update.

Imagine, then, that a court constrained by the precedents from a consistent case base Γ is confronted with a new problem situation X , and suppose to begin with that none of the rules belonging to $Rule(\Gamma)$ is even applicable to X . A situation like this is our theoretical analogue to the legal notion of a case of first impressions, presenting issues that have not previously been addressed within the law, so that the established rules of precedent have no bearing. In such a situation, the court is utterly unconstrained by precedent.¹⁴ The court is free to assess the issues in whatever way it thinks best, to formulate a rule r supporting an

¹⁴See, for example, the discussion in Cross (1968, pp. 190–195).

outcome s , and to reach a decision of the form $\langle X, r, s \rangle$. This decision, of course, leads to $\Gamma \cup \{\langle X, r, s \rangle\}$ as the updated case base, and fortunately—since there is no real precedential constraint—any such decision is guaranteed to preserve consistency.

Observation 2 Let Γ be a consistent case base, and suppose X is a fact situation in which none of the rules from $Rule(\Gamma)$ are applicable. Then $\Gamma \cup \{\langle X, r, s \rangle\}$ is also consistent, where r is any newly formulated rule applicable in X and supporting s as an outcome.

Next, still imagining that a court constrained by Γ is confronted with a new situation X , we now suppose that some precedent rule from $Rule(\Gamma)$ is, in fact, binding in this situation. Then, according to the standard theory, any court has at least two options, and some courts have a third: any court has the option of either following or distinguishing the binding precedent rule, and some courts, depending on their standing in the judicial hierarchy, have the further option of overruling the precedent.

Where r is such a binding rule from $Rule(\Gamma)$, supporting the outcome s , we can say that the court *follows* this rule whenever it reaches a decision of the form $\langle X, r, s \rangle$, generating $\Gamma \cup \{\langle X, r, s \rangle\}$ as the updated case base. As it turns out, the court is guaranteed to satisfy precedential constraint whenever it follows a binding rule, in the sense that the resulting decision preserves consistency.

Observation 3 Let Γ be a consistent case base with $Rule(\Gamma)$ the derived set of rules, and suppose X is a fact situation in which some rule r from $Rule(\Gamma)$, supporting the outcome s , is binding. Then the case base $\Gamma \cup \{\langle X, r, s \rangle\}$ is also consistent.

We can say that the court *distinguishes* the binding rule r , by contrast, whenever it either follows or formulates another rule r' supporting the opposite outcome \bar{s} , leading to a decision

of the form $\langle X, r', \bar{s} \rangle$, and so generating $\Gamma \cup \{\langle X, r', \bar{s} \rangle\}$ as the updated case base; this notion of distinguishing a rule differs from the traditional idea, to which it will be compared in the following section. And of course, while, as we have seen, a court is guaranteed to preserve consistency whenever it follows a binding rule, there is no such guaranteed with distinguishing; the operation must be performed with some care, guided by the constraints of precedent.

Both the ideas of following and distinguishing a binding rule were illustrated by our previous example, in which a court constrained by the case base $\Gamma_1 = \{c_1, c_2\}$ confronts the new fact situation $X_3 = \{f_1^\pi, f_2^\pi, f_5^\pi, f_5^\delta, f_7^\delta\}$. Here, as we noted, the rule r_2 , or $\{f_5^\delta\} \rightarrow \delta$, is the unique binding rule, so that the court follows this rule if it reaches the decision $c_3 = \langle X_3, r_3, s_3 \rangle$, where r_3 is r_2 and s_3 is δ , leading to $\Gamma_1 \cup \{\langle X_3, r_3, s_3 \rangle\}$ as the updated case base. But as we also noted, the court is likewise free to formulate and apply the new rule r_4 , or $\{f_5^\pi\} \rightarrow \pi$, leading to the decision $c_4 = \langle X_4, r_4, s_4 \rangle$, where X_4 is X_3 and s_4 is π , and so to $\Gamma_1 \cup \{\langle X_4, r_4, s_4 \rangle\}$ as the updated case base. In that circumstance, we would now say that the court distinguishes the binding rule r_2 .

Overruling a case is a more radical operation, generally available only to courts either above or, sometimes, at the same level as that which decided the case to be overruled. Even then, this option is avoided whenever possible, since the changes wrought in the underlying case base are so extreme: when a precedent case is overruled, it is as if the case were completely “wiped off the slate.”¹⁵ How can this operation be modeled in the present framework? One option might be to imagine that the logical effect of overruling a case

¹⁵Cross (1968, p 119) attributes the phrase to Lord Dunedin; Raz (1979, p. 189) speaks of overruling a case as removing it “root and branch.”

is simply that the overruled case should then be removed from the case base entirely. On this view, the updated case base that results from overruling a case c would be nothing more than $\Gamma - \{c\}$.¹⁶ But the precedent rule set out in the case c might have been employed in other cases as well. And it is also possible to imagine that, in overruling a case, the court wishes to invalidate this precedent rule itself, so that the logical effect of overruling the case c is to remove from the case base, not just the single case c , but every case c' sharing the same rule. On this view, the updated case base that results from overruling a case c would then be $\Gamma - \{c' : Rule(c') = Rule(c)\}$.

There are other options as well. Perhaps in overruling the case c , the court means to invalidate, not every other case that shares its rule with c , but only those that actually follow this rule, so that a temporal dimension would have to be introduced. Overruling raises a number of complicated issues, which I do not intend to explore in this paper, except to note this: on nearly any sensible view, the operation of overruling a case, unlike the other case base update operations considered here, must be treated as a destructive operation, in the sense that it modifies or removes information already present in a case base, rather than simply augmenting that information.

¹⁶This formula indicates the effect simply of overruling the previous case c , but of course, the matter would arise only through the consideration of a new fact situation, for which the court would then be likely to formulate an entirely new rule, supporting a different outcome from that contained in the overruled case, so that the actual result would be more complicated. Suppose, for example, that the case c is $\langle X, r, s \rangle$, and that the court is confronted with the new fact situation X' , in which the rule r is binding, but that the court wishes to avoid the outcome s by overruling the previous case. The court would then be likely to formulate another rule r' , also applicable to X' , but yielding the opposite outcome \bar{s} , so that the updated case base would then be, not just $\Gamma - \{c\}$, but $(\Gamma - \{c\}) \cup \{\langle X', r', \bar{s} \rangle\}$.

Returning, then, to the more straightforward case base update operations, I want to close with three points.

First of all, it follows at once from Observations 2 and 3 that a court, presented with a new fact situation, will always be in a position to comply with the constraints of precedent—the court will never be forced to introduce an inconsistency into the case base. For either no rule from the background case base is applicable in the current situation, or some binding rule is applicable.¹⁷ If no rule is applicable, then it follows from Observation 2 that any decision reached by the court will preserve consistency. If some binding rule is applicable, then it follows from Observation 3 that the court can preserve consistency by following that rule.

Second, although the traditional concepts of following and distinguishing a precedent rule can be accommodated within the framework set out here, the present framework is more general, allowing us to understand options available to a court that do not seem to fall naturally within the traditional classification. To illustrate, we return to the example of a court confronting the situation $X_3 = \{f_1^\pi, f_2^\pi, f_5^\pi, f_5^\delta, f_7^\delta\}$ against the background of the case base $\Gamma_1 = \{c_1, c_2\}$, with r_2 , or $\{f_5^\delta\} \rightarrow \delta$, as a binding rule. The rule r_2 , of course, was formulated in the context of the precedent case c_2 , with $X_2 = \{f_1^\pi, f_2^\pi, f_4^\pi, f_5^\delta, f_6^\delta\}$ as its fact situation. What the c_2 court is telling us with its decision, then, is that the reason provided by $\{f_5^\delta\}$ for the defendant outweighs the consideration $\{f_1^\pi, f_2^\pi, f_4^\pi\}$ present in X_2 for the plaintiff.

¹⁷We can assume that the set of precedent rules is finite, so that not every applicable rule can be trumped by another rule; as a result, if any rule is applicable to a fact situation, some rule must be binding—applicable and untrumped.

Now, we have previously imagined that the current court, confronted with X_3 , might conclude that the consideration $\{f_5^\pi\}$, which was not present in X_2 , outweighs the consideration $\{f_5^\delta\}$, and so decide X_3 for the plaintiff on the basis of the new rule r_4 , or $\{f_5^\pi\} \rightarrow \pi$, thus reaching the decision $c_4 = \langle X_4, r_4, s_4 \rangle$. But we could also suppose that the current court, while still recognizing that the consideration $\{f_5^\pi\}$ for the plaintiff outweighs $\{f_5^\delta\}$ for the defendant, also feels that the new consideration $\{f_7^\delta\}$ for the defendant, present here for the first time, itself outweighs $\{f_5^\pi\}$. The current court, then, might reach the decision $c_6 = \langle X_6, r_6, s_6 \rangle$, where X_6 is X_3 , where s_6 is δ , and where r_6 represents the new rule $\{f_7^\delta\} \rightarrow \delta$. In this case, then, although both the decision and the train of thought leading to it seem to make perfect sense, it would be hard to describe the court's decision in terms of the traditional vocabulary of following or distinguishing a rule. The court is not following the binding rule r_2 , or $\{f_5^\delta\} \rightarrow \delta$, since its decision is not based on this rule, and indeed, it recognizes that the fact situation contains the consideration $\{f_5^\pi\}$ for the plaintiff which outweighs the consideration $\{f_5^\delta\}$ for the defendant, on which the rule r_2 is based. But the court is not distinguishing the rule r_2 either, since it reaches a decision for the side that is favored by that rule.

Third, and finally, the present framework allows us to understand a feature of the common law that can seem very puzzling from a more traditional perspective—the idea that simply following a precedent rule can lead to a change in the law. This idea is often alluded to in the legal literature. Edward Levi, for example, writes that “the rules change from case to case and are remade with each case,” and later, that “the rules change as the rules are

applied.”¹⁸ And, Simpson likewise, after discussing the operations of distinguishing and following a precedent rule, writes that “the development of the law is normally brought about by just these two activities,” which leads him to a sort of paradox: “The legal process is conceived of as conditioned by rules, yet in a sense the rules change from case to case; the very point in having a system of rules to ensure consistency in decision seems to be frustrated if the rules themselves lack fixity.”¹⁹

These remarks can be hard to understand. It is easy enough to see how distinguishing a precedent rule might introduce a change into the law. In the present view, as we have seen, distinguishing a rule often involves introducing a new rule into the case base; and the change is even more palpable on the standard treatment, which we will consider shortly, where distinguishing a precedent rule involves an actual modification of that rule. But if a court simply follows a precedent rule—if it does no more than draw a rule from some precedent case and apply that same rule to a new fact situation—how can we say that the rules are changed, or that the law is affected at all?

The present account allows us to answer this question, if we take phrases such as “the rules change as the rules are applied” as referring, not necessarily to the set of precedent rules themselves, but, in a more metaphorical way, to the precedential constraints generated by the underlying case base. In that case, although simply following a familiar precedent rule, applying it in a new situation, does not lead to any modifications in the set of rules derived from a case base, it does indeed affect the precedential constraints generated by that case base. For on the present view, these constraints depend, not on the set of precedent

¹⁸See Levi (1949, pp. 2–4).

¹⁹See Simpson (1961, p. 172), who himself cites Levi.

rules, but on consistency with the background case base. Consistency, in turn, is defined in terms of the preference ordering on reasons, or considerations, that is generated by the case base. And as it turns out, even following a precedent rule—simply applying a familiar rule in a new situation—can change the preference ordering on considerations in a way that affects precedential constraint. Each time a familiar rule is followed in a new case, the court makes the decision that certain considerations, which might have been judged as strong enough to override that rule, are in fact not strong enough. These decisions, encoded in new cases, then modify the preference relations derived from the case base, which affects the options open to the court when it encounters those same considerations again in the future.

The point can be illustrated by returning once again to our familiar example, in which a court constrained by the case base $\Gamma_1 = \{c_1, c_2\}$ confronts the new fact situation $X_3 = \{f_1^\pi, f_2^\pi, f_5^\pi, f_5^\delta, f_7^\delta\}$. Now, as we have seen, even though r_2 , or $\{f_5^\delta\} \rightarrow \delta$, is the unique binding rule in this new situation, precedential constraints allow the court to distinguish this rule, arriving instead at the decision $c_4 = \langle X_4, r_4, s_4 \rangle$, where r_4 is $\{f_5^\pi\} \rightarrow \pi$.

Suppose, however, that, prior to confronting the fact situation X_3 , the court is first faced with the fact situation $X_7 = \{f_5^\pi, f_6^\pi, f_5^\delta, f_7^\delta\}$, again with r_2 as the unique binding rule; but here, it decides simply to apply this rule, resulting in the decision $c_7 = \langle X_7, r_7, s_7 \rangle$, where r_7 is r_2 and s_7 is δ . In that case, the court would then be confronting the new situation X_3 against the background of the case base $\Gamma_2 = \{c_1, c_2, c_7\}$, rather than Γ_1 . And even though Γ_2 results from Γ_1 simply through the application to a new situation of a rule that is already present in Γ_1 , the mere application of this rule changes the precedential constraints derivable from the background case base to such an extent that the court is no longer free, when faced with the situation X_3 , to reach the decision $c_4 = \langle X_4, r_4, s_4 \rangle$. For it is easy to

see that $\{f_5^\pi\} <_{c_7} \{f_5^\delta\}$ —it follows from the court’s decision in c_7 that the consideration $\{f_5^\delta\}$ for the defendant is stronger than the consideration $\{f_5^\pi\}$ for the plaintiff. But we also have $\{f_5^\delta\} <_{c_4} \{f_5^\pi\}$, so that the decision c_4 would tell us exactly the opposite—that the consideration $\{f_5^\pi\}$ for the plaintiff is stronger than the consideration $\{f_5^\delta\}$ for the defendant. The decision c_4 , therefore, is inconsistent with Γ_2 or any other case base containing c_7 , and so ruled out by precedential constraint.²⁰

²⁰Lamond (2005) makes exactly this point—that following, as well as distinguishing, a previous precedent rule can change the law. However, the way in which he develops this idea indicates, I believe, a problem with his picture of case base dynamics and the evolution of legal doctrine. On Lamond’s view, the doctrine provided by a precedent rule, or *ratio*, consists, not only in the particular reason carried by that rule, but in the set of facts that have been judged as insufficient in strength to defeat that reason. Following a rule thus leads to a change in legal doctrine because: “every time a precedent is followed, further facts are added to the list of those regarded as insufficient to defeat the reason provided by the *ratio* (2005, p. 17; see also p. 20).” To illustrate with our current example, the doctrine provided by the case c_2 would, on this view, consist in the rule r_2 , or $\{f_5^\delta\} \rightarrow \delta$, along with the list $\{f_1^\pi, f_2^\pi, f_4^\pi\}$ of factors judged by the c_2 court as insufficient in strength to defeat this rule; the application of the same rule in the case of c_7 would change the doctrine by expanding this list to $\{f_1^\pi, f_2^\pi, f_4^\pi, f_5^\pi, f_6^\pi\}$, which includes also the new factors judged by the c_7 court as insufficient to defeat r_2 . Although I agree with this general idea behind this approach, of course, I feel that the use of a simple list to represent insufficient defeaters is not adequate. For this reason, I rely here on a more general preference relation to represent strength comparisons among conflicting considerations—sets of factors, rather than individual factors—and it is this relation that is then updated as rules are applied. In the present paper, then, the doctrine of c_2 consists in the rule the rule r_2 together with the information that $\{f_5^\delta\}$ is stronger than the consideration $\{f_1^\pi, f_2^\pi, f_4^\pi\}$ or any of its subsets, so that none of these considerations can defeat that rule; the case c_7 then adds the further information that $\{f_5^\delta\}$ is likewise stronger than $\{f_5^\pi, f_6^\pi\}$ or its subsets. But all of this is consistent with the idea that a court might eventually decide that $\{f_5^\delta\}$ is not stronger than, say, the consideration $\{f_4^\pi, f_5^\pi\}$ —which contains a pair of potential defeaters, one from each of

4 Comparisons

This section compares aspects of the current theory of precedential constraint to two previous accounts. First, the overall theory of constraint presented here is compared to the result model. Second, the present notion of distinguishing a precedent rule is compared to the conventional view.

4.1 The result model

The result model of precedent holds that a body of precedent cases constraints later courts only when these courts are presented with a fortiori fact situations—situations that are at least as strong for the winning side of some precedent case as the precedent case itself. Obviously, this model presupposes some ordering through which different fact situations can be compared in strength for one side or another. As mentioned earlier, Alexander has objected to the result model on the grounds that any such ordering would be unattractive and perhaps incoherent, but I have elsewhere proposed a strength ordering that I believe avoids these objections and allow us to provide a sensible definition of the result model.

I do not intend to review these various arguments here; the reader is referred to my earlier paper, as well as to Alexander’s work.²¹ My present goal is simply to compare the notion of constraint derived from the result model—which I refer to here as *a fortiori constraint*—with the concept of precedential constraint set out in the current paper.

the two cases, which had not yet been considered in combination. It is hard to see how this latter possibility could be captured in Lamond’s list notation, which fails to represent the information that individual factors such as f_4^π and f_5^π originated in different cases, and so had not yet been considered together.

²¹See Horty (2004) and Alexander (1989).

We begin with the proposed strength ordering on fact situations, which is motivated in more detail in my earlier paper. The idea is that a fact situation Y presents a case for the side s which is at least as strong as that presented by the fact situation X whenever Y contains all the factors from X that support s , and X contains all the factors from Y that support \bar{s} , the opposite side. If we let \leq^s represent the strength ordering for the side s , this idea can then be defined formally as follows.

Definition 9 (Strength for a side) Let X and Y be fact situations. Then Y is at least as strong as X for the side s —written, $X \leq^s Y$ —just in case $X^s \subseteq Y^s$ and $Y^{\bar{s}} \subseteq X^{\bar{s}}$.

To illustrate, consider the fact situations $X_8 = \{f_1^\pi, f_1^\delta, f_2^\delta\}$ and $X_9 = \{f_1^\pi, f_2^\pi, f_1^\delta\}$. We then have $X_8 \leq^\pi X_9$, since X_9 contains all the factors from X_8 that support π , and X_8 contains all the factors from X_9 that support δ ; and we can see, likewise, that $X_9 \leq^\delta X_8$. This definition, I have argued, conforms to our intuitions, and it exhibits a number of plausible formal properties as well.²²

Given this notion of strength, we can now note, as an aside, that it allows us to generalize the earlier Observation 1, according to which a certain anomaly in a case base entails inconsistency. What this observation shows is that any case base containing two cases in which the same fact situation X is decided for two different sides, both s and \bar{s} , must be

²²It is easy to verify that \leq^s is a partial ordering: reflexive, transitive, and antisymmetric. That is, for any situations X , Y , and Z , we have: $X \leq^s X$; $X \leq^s Y$ and $Y \leq^s Z$ implies $X \leq^s Z$; and $X \leq^s Y$ and $Y \leq^s X$ implies $X = Y$. The ordering also exhibits a duality property, according to which $X \leq^s Y$ if and only if $Y \leq^{\bar{s}} X$. It is important to note, however, that this strength ordering is not linear, or connected. Given two situations X and Y , it may well be that we have neither $Y \leq^s X$ nor $X \leq^s Y$, in which case the two situations would be incomparable in strength.

inconsistent, in our formal sense. We can now see that a case base is likewise inconsistent if it contains cases in which one fact situation X is decided for the side s while another fact situation Y , which is at least as strong for s as X , happens to be decided for \bar{s} .

Observation 4 Let Γ be a case base containing two precedent cases of the form $\langle X, r, s \rangle$ and $\langle Y, r', \bar{s} \rangle$ where $X \leq^s Y$. Then Γ is inconsistent.

Since any fact situation is at least as strong as itself for either side, this result generalizes our earlier observation, and so provides further support that our formal notion of inconsistency is intuitively correct.

Once this strength ordering \leq^s is in place, it is then a straightforward matter to define the concept of a fortiori constraint, the notion of constraint at work in the result model of precedent. The idea, once again, is that a court faced with a fact situation X should reach a decision for the side s whenever X is at least as strong for s as some precedent case that was itself decided for that side—whenever, that is, X is at least as strong for s as the fact situation of some precedent case whose outcome was a decision for s .

Definition 10 (A fortiori constraint) Let Γ be a case base and X a new fact situation confronting the court. Then a fortiori constraint requires the court to reach a decision in X for the side s just in case there is some precedent case c from Γ such that $Outcome(c) = s$ and $Facts(c) \leq^s X$.

To illustrate, suppose the background case base Γ contains the familiar case $c_1 = \langle X_1, r_1, s_1 \rangle$, where $X_1 = \{f_1^\pi, f_2^\pi, f_3^\pi, f_1^\delta, f_2^\delta, f_3^\delta, f_4^\delta\}$, where the rule r_1 is $\{f_1^\pi, f_2^\pi\} \rightarrow \pi$, and where the outcome s_1 is π ; and imagine that the court is confronting the new fact situation $X_{10} = \{f_1^\pi, f_2^\pi, f_3^\pi, f_1^\delta\}$. Of course, $Outcome(c_1) = \pi$ and $Facts(c_1) = X_1$, and it is easy to see that

$X_1 \leq^\pi X_{10}$, since X_{10} contains all the factors from X_1 that favor π and fewer that favor δ . The notion of a fortiori constraint therefore requires that X_{10} should be decided for the plaintiff, since Γ contains the case c_1 with $Outcome(c_1) = \pi$ and $Facts(c_1) \leq^\pi X_{10}$.

Now, what is the relation between this notion of a fortiori constraint, derived from the result model, and the concept of precedential constraint advanced in this paper, which involves reasons and rules? There are two initial points to make. The first is that a fortiori constraint is at least as strong as precedential constraint, in the sense that, in any situation in which a fortiori constraint requires an outcome, a decision for some particular side, precedential constraint requires that same outcome.

Observation 5 Let Γ be a consistent case base and X a new fact situation confronting the court, and suppose a fortiori constraint requires the court to reach a decision for the side s in the situation X . Then precedential constraint also requires the court to reach a decision for the side s .

The second point is that there are some situations in which precedential constraint requires an outcome that is not required by a fortiori constraint. Suppose, for example, that the background case base contains the single precedent case c_1 , as above, and that the court is confronted with the new fact situation $X_{11} = \{f_1^\pi, f_2^\pi, f_1^\delta\}$. It is then easy to see that the relation $Facts(c_1) \leq^\pi X_{11}$ fails—the fact situation X_{11} is not stronger for the plaintiff than X_1 , since, although X_{11} contains fewer factors than X_1 that favor δ , it contains fewer that favor π as well. The court, therefore, is not required by a fortiori constraint to reach a decision for the plaintiff in this situation.

On the other hand, the court would be required by precedential constraint to decide for

the plaintiff in the fact situation X_{11} . After all, the rule of the precedent case c_1 is r_1 , or $\{f_1^\pi, f_2^\pi\} \rightarrow \pi$. What the c_1 court is telling us with its decision for the plaintiff, then, is that the consideration $Premise(r_1) = \{f_1^\pi, f_2^\pi\}$ outweighs the consideration $X_1^\delta = \{f_1^\delta, f_2^\delta, f_3^\delta, f_4^\delta\}$, or any of its subsets. In particular, we have $\{f_1^\delta\} <_{c_1} \{f_1^\pi, f_2^\pi\}$ —that is, the consideration $\{f_1^\pi, f_2^\pi\}$ for the plaintiff outweighs the consideration $\{f_1^\delta\}$ for the defendant, according to the c_1 court. In the new situation X_{11} , then, a decision in for the defendant would carry exactly the opposite information, that the consideration $\{f_1^\delta\}$ for the defendant outweighs the consideration $\{f_1^\pi, f_2^\pi\}$ for the plaintiff, thus introducing an inconsistency into the case base. Therefore, since a decision for the defendant in this situation leads to inconsistency, precedential constraint requires a decision for the plaintiff.

Putting these two points together, we can conclude that the concept of a fortiori constraint at work in the result model is stronger, strictly speaking, than the notion of precedential constraint advanced here: in any situation in which a fortiori constraint requires an outcome, precedential constraint requires that same outcome, but there are some situations in which precedential constraint requires an outcome that a fortiori constraint does not. Furthermore, we can begin to see how it is, exactly, that the explicit appeal to rules, or reasons, can enhance the notion of constraint that is derived purely from results.

Put abstractly, the force of a court's decision for s in the situation X , according to the result model, is simply that $X^{\bar{s}}$ is outweighed by X^s —that the strongest consideration favoring \bar{s} is outweighed by the strongest consideration favoring s . This decision thus constrains only those future cases in which the considerations favoring \bar{s} are no stronger than $X^{\bar{s}}$ and the considerations favoring s are at least as strong as X^s . On the present theory, by contrast, the force of a court's decision for s in the situation X is that $X^{\bar{s}}$ is outweighed, not simply by

X^s , but by $Premise(r)$, where r is the rule appealed to in the decision—that the strongest consideration favoring \bar{s} is outweighed by the premise of the rule supporting a decision for s . Since the premise of some rule favoring s can be weaker than the strongest consideration supporting s —that is, since $Premise(r) \subseteq X^s$ —the precedent has broader reach: it constrains cases in which the considerations favoring \bar{s} are no stronger than $X^{\bar{s}}$ and the considerations favoring s are at least as strong as $Premise(r)$, even though they may not be as strong as X^s . The additional expressive resources provided by explicitly formulated rules, or reasons, thus allow courts to weaken the conditions necessary for constraint, and so formulate precedents with greater generality.

The same point can be seen from the other side as well. Rather than considering the way in which rules add expressive resources to the result model, we can instead view the result model itself as a special case of the present theory, in which rules are restricted to a particular form. Consider again a fact situation X that is decided for s . As we have seen, the present theory allows the court, by formulating a rule r , to broaden the scope of its decision by citing as its reason, not necessarily X^s , the strongest possible consideration favoring s , containing every factor from X that favors this result, but only $Premise(r)$, a subset of this consideration. But suppose the court does, in fact, cite as the reason for its decision the very strong consideration X^s , by formulating a rule of the form $X^s \rightarrow s$. In that case, as it turns out, the decision constrains exactly the same future cases under both the present theory and the result model. If we now imagine that courts are limited to rules of this special form, the present theory of precedential constraint collapses into the result model.

Observation 6 Consider only cases $\langle X, r, s \rangle$ in which the rule r has the form $X^s \rightarrow s$. Let

Γ be a consistent set of such cases, and suppose that Y is a new fact situation confronting the court. Then a fortiori constraint requires the court to reach a decision for some particular side in this situation if and only if precedential constraint requires a decision for the same side.

This result suggests two lines of interpretation, which I will simply mention here. It provides us, first of all, with a way of interpreting within the present framework the occasional case in which a court reaches some decision on a the basis of a fact situation but no rule is explicitly supplied: if the fact situation is X and the court reaches a decision for the side s , we can thus imagine the court as working with an implicit rule of the form $X^s \rightarrow s$. And second, the result also suggests a charitable interpretation of Arthur Goodhart’s famous “material facts” version of the rule of a case.²³ On Goodhart’s view, legal decisions are often influenced by principles of which the court is unaware, or which the court misunderstands; accordingly, he places less emphasis on the rule actually formulated by the court to justify its decision and more emphasis on an implicit rule which has as its premise the material facts of the case and as its conclusion the outcome arrived at by the court. But surely not all the material facts of a given case can be taken to support its outcome. If the fact situation is X and the outcome arrived at by the court is s , only those facts belonging to X^s actually support s as an outcome; the others—those facts belonging to $X^{\bar{s}}$ —instead support \bar{s} , the opposite outcome. Goodhart can thus be interpreted as holding that, regardless of the rule explicitly formulated by the court, the principle actually guiding the court’s decision for s in the situation S is simply the rule $X^s \rightarrow s$, according to which the material facts from

²³See Goodhart (1930).

X that favor s provide a reason for s , and a stronger reason for than those that favor the opposite side.

4.2 Distinguishing

The power of courts to distinguish previous rulings is one of the most striking features of the common law, lying precisely at the fulcrum between constraint and freedom; it is, as Raz writes, the “power to develop the law even when deciding regulated cases and even by courts which have no power to overrule.”²⁴ On the present account, as we have seen, distinguishing is a simple process: a court confronting a new fact situation is said to distinguish a binding rule whenever—while still conforming to precedential constraint—the court formulates or follows a different rule in that situation, leading to the opposite outcome.

This is not the standard view of distinguishing, however, which differs in two ways. First of all, the present account is cast against the general background of precedent rules as defeasible, telling us that the premise of a rule entails its conclusion by default—or equivalently, that the premise of the rule provides a reason for reaching its conclusion. On the standard view, by contrast, precedent rules tend to be interpreted as universal generalizations of the kind familiar from ordinary classical logic, telling us that the conclusion holds in any situation whatsoever in which the premise is satisfied. Some legal theorists argue, in fact, that this very strong interpretation of rules as universal generalizations is actually required if precedent rules are to play the appropriate role in justifying a court’s decisions. Neil MacCormick, for example, writes that “justification requires universalization,”—that a precedent rule should have the form of a “universal, ‘whenever f_1, f_2, \dots, f_n , then j ’,” and that “the

²⁴Raz (1979, p. 185).

‘whenever’ clause is essential to justify the actual decision.”²⁵

We can mark the difference between these two kinds of rules with a shift of notation. Our present, defeasible rules are written with a single arrow notation, so that $Y \rightarrow s$ means that, in any situation containing all the factors from Y , the court ought, by default, to hold for the side s , or again, that the presence of these factors provides a reason to hold for this side. The interpretation of precedent rules as universal generalizations will now be indicated with a stronger double arrow, so that $Y \Rightarrow s$ is taken to mean that the court must find for s in any situation containing the factors from Y .

The second difference between the present account and the standard view is this. While, on the present account, a court can distinguish an earlier precedent rule in a particular fact situation simply by applying another rule that leads to a different result, the standard view requires that the later court must actually modify the original precedent rule. As Raz describes it: “one distinguishes a rule by changing it so that a rule which did apply to the present case no longer applies to it in its modified form.”²⁶

Of course, later courts do not have complete freedom to modify earlier precedent rules; the process is generally taken to be subject to two conditions. First, an earlier precedent rule can only be narrowed, and narrowed, furthermore, only by the addition to its premise of some minimal requirement necessary to distinguish the current fact situation from the previous cases in which the rule has been applied. And second, even once it has been modified, the rule must remain applicable to the previous cases in which it has been applied, supporting the same result in each of those cases. The first of these conditions derives from the idea

²⁵MacCormick (1987, p. 162).

²⁶Raz (1979, p. 185).

that courts have no general power to draft legislation, but can only respond to the particular fact situations before them. The second condition indicates the way in which the standard view attempts to capture the fundamental idea that the decisions of later courts must be consistent with those of earlier courts. Because later courts can modify the rules set out by earlier courts only to the extent that they continue to support earlier decisions, these modified rules can be taken, not as real changes, but only as refinements of the earlier rules. They are often interpreted as the rules that the earlier courts “actually had in mind,” or as the rules these earlier courts “would have formulated” had they envisioned the current fact situation.²⁷

The differences between these two approaches to distinguishing—the present account and the standard view—can be highlighted by returning once again to our familiar example. Consider once again, then, a court confronted with the new fact situation $X_3 = \{f_1^\pi, f_2^\pi, f_5^\pi, f_5^\delta, f_7^\delta\}$ against the background of the case base $\Gamma_1 = \{c_1, c_2\}$, so that the only rule binding in this situation is r_2 , or $\{f_5^\delta\} \rightarrow \delta$, derived from the precedent case $c_2 = \langle X_2, r_2, s_2 \rangle$, where $X_2 = \{f_1^\pi, f_2^\pi, f_4^\pi, f_5^\delta, f_6^\delta\}$ and where s_2 is δ . And let us again imagine that, even though the rule r_2 supports a decision for the defendant, the court is sufficiently impressed with the weight of the new factor f_5^π in the situation X_3 that it wishes to distinguish the precedent rule in order to reach a decision, instead, for the plaintiff.

On the present account, as we have seen, the court can distinguish the previous rule r_2 simply by formulating and appealing to the new rule r_4 , or $\{f_5^\pi\} \rightarrow \pi$, thus introducing into the case base the new case $c_4 = \langle X_4, r_4, s_4 \rangle$, where X_4 is X_3 and s_4 is π . This new decision

²⁷See Raz (1979, p. 188) for a discussion of this interpretation.

carries with it the preference information that $\{f_5^\delta\} <_{c_4} \{f_5^\pi\}$ —that the consideration $\{f_5^\pi\}$ for the plaintiff is stronger than the consideration $\{f_5^\delta\}$ for the defendant, according to the c_4 court—which must then be respected by later courts as well. As a result, although the previous rule r_2 may continue to serve as a binding rule in cases containing f_5^δ without f_5^π , in any future cases that contain both f_5^δ and f_5^π , where both the previous r_2 and the new r_4 are applicable, the previous rule will be trumped by the new rule, so that only the later can be binding.

On the standard view of distinguishing, things are more complicated. To begin with, the rule set out in the precedent case c_2 must now be interpreted as an ordinary universal generalization, here represented as $\{f_5^\delta\} \Rightarrow \delta$, and so taken to mean that the court must find for the defendant in any situation whatsoever containing the factor f_5^δ . Because of the meaning attached to this previous rule, the current court—still impressed with the weight of the new factor f_5^π —must now explicitly modify the rule so that it does not yield the undesired result in the case of X_3 , which does, after all, contain f_5^δ . The process might lead, on the standard view, to the replacement of the original rule $\{f_5^\delta\} \Rightarrow \delta$ by the modified version $\{f_5^\delta, \neg f_5^\pi\} \Rightarrow \delta$, where $\neg f_5^\pi$ indicates the absence of the feature f_5^π , so that the modified rule now means that a court must decide for the defendant in any situation containing the factor f_5^δ but lacking f_5^π . This modification satisfies the two conditions associated with the traditional view: it narrows the previous rule with some minimal requirement distinguishing the current situation from the earlier case in which the rule was applied, yet it continues to support the decision in that earlier case. And the modification blocks the application of the rule in X_3 , since this new situation does contain f_5^π .

Once the previous rule has been modified in order to avoid generating the undesired result

in the case of X_3 , the current court can then formulate the fresh rule $\{f_5^\pi\} \Rightarrow \pi$, leading to the result it desires, a decision for the plaintiff. And just as before, the modified rule will now continue to apply in situations containing f_5^δ alone, without f_5^π , but only the new rule will apply in situations containing f_5^π as well, since the application of the modified rule has now been explicitly blocked in these circumstances.

In a sense, then, these two approaches to distinguishing—the present account and the standard view—accomplish much the same thing. On both approaches, the result of distinguishing the previous rule, set out in c_2 , allows the current court to decide the fact situation X_3 for the plaintiff, as well as any others containing both the factors f_5^δ and f_5^π , while still supporting the earlier court’s decision in c_2 . Nevertheless, in spite of their similarities, I feel that the present account of distinguishing allows us to avoid two real problems confronting the standard view.²⁸

The first problem is broadly conceptual, and flows from the standard view of precedent rules as universal generalizations. The problem is this: if precedent rules are actually taken as ordinary universal generalizations, then these rules are so strong that it is hard to see how they can allow for the practice of distinguishing at all—it is hard to see how a later court can, in fact, consistently arrive at a conclusion contrary to that directed by an earlier precedent rule. Consider again our example in which the current court, confronted with the situation X_3 , is guided only the by the rule $\{f_5^\delta\} \Rightarrow \delta$, set out earlier in response to the situation X_2 . Precedent requires the decisions of later courts to be consistent with those of earlier courts and it is true, as we have seen, that, in distinguishing the earlier rule, the

²⁸Additional arguments against the standard view of distinguishing are provided by Lamond (2005).

current court continues to support the earlier court's decision in the case of X_2 . But it is also reasonable to ask: how can the current court's own decision in X_3 , the current situation, be construed as consistent with the rule set out by the earlier court? If rules are to be interpreted as universal generalizations, then what this earlier rule means is that a situation should be decided for the defendant whenever it contains the factor f_5^δ . The current court is facing just such a situation in X_3 , yet it decides for the plaintiff all the same, apparently in direct violation of the earlier rule.

Of course, the standard view is that, by distinguishing, later courts have the power to refine the rules set out by earlier courts; the idea, then, is that their decisions must be consistent only with these later refinements, not with the rules as originally stated. But putting the point this way merely throws a different light on the problem: if later courts are free to modify the rules set out by earlier courts, in order to avoid their consequences, then they are free to reach decisions inconsistent with the rules as originally formulated.²⁹

²⁹This argument, I believe, may help us understand the very strict version of the rule model formulated by Alexander in (1989) and elsewhere, according to which a court constrained by a precedent rule is allowed only two real options: following that rule, or overruling the precedent. Like many writers, Alexander (p. 19) interprets precedent rules as ordinary universal generalizations. And if the concern raised in the text is legitimate, he can therefore be viewed as recognizing, correctly, a tension between this interpretation of precedent rules and the idea that they can be narrowed, or distinguished, by later courts. The present theory, of course, also recognizes the tension between the interpretation of precedent rules as ordinary universal generalizations and the idea that they can be distinguished, but I take it in the opposite direction: rather than holding on to the interpretation of precedent rules as universal generalizations and rejecting the possibility of distinguishing these rules, I allow for distinguishing while interpreting precedent rules, instead, only as defaults.

This same problem—the excessive strength of rules as universal generalizations—can be seen from the other side as well: not from the standpoint of a later court that wishes to distinguish an earlier rule, but from standpoint of an earlier court as it formulates that rule. Consider a variant of our example, in which the c_2 court, facing the situation $X_2 = \{f_1^\pi, f_2^\pi, f_4^\pi, f_5^\delta, f_6^\delta\}$, still wishes to decide for the defendant on the basis of the factor f_5^δ , but also feels, even at the time of its decision, that a variety of competing factors, had they been present, would have outweighed f_5^δ and led instead to a decision for the plaintiff. Of course, the court might attempt to enumerate some of these factors. Perhaps it envisions f_5^π itself along with f_8^π , say, as factors for the plaintiff that might outweigh f_5^δ , and so formulates the rule $\{f_5^\delta, \neg f_5^\pi, \neg f_8^\pi\} \Rightarrow \delta$ in an effort to justify its decision; interpreted as a universal generalization, what this rule means is that a case should be decided for the defendant whenever f_5^δ is present and both f_5^π and f_8^π are absent.

But we can easily imagine that the court has no faith in its own ability to consider every possible circumstance, and is fully aware that other factors, in addition to f_5^π and f_8^π , might also outweigh f_5^δ . And in that case, the court must be aware, also, that the universal generalization contained in its own rule is actually incorrect—that there are some situations in which f_5^δ is present and both f_5^π and f_8^π are absent that ought to be decided for the plaintiff all the same. As this example suggests, the very strong view of precedent rules as universal generalizations seems to force courts into the odd position of justifying their decisions on the basis of rules that these courts themselves regard as mistaken. Indeed, exactly this point—that the standard view forces courts to distance themselves from their own rules—is recognized, and endorsed, by MacCormick, who emphasizes a distinction between the “logical properties” of the rules set out by a court and the court’s own “quality of commitment”

to those rules: “One’s rulings,” he writes, “ should be universalized or at the very least universalizable; but one’s commitment to them can and should be relativistic, open-ended, and therefore revisable.”³⁰

The present account avoids this problem all together. To begin with, there is no conceptual problem involved in the process of distinguishing a rule. In the case of c_2 , for example, the force of the rule $\{f_5^\delta\} \rightarrow \delta$, interpreted defeasibly, is simply that $\{f_5^\delta\}$ provides a reason favoring a decision for the defendant; this is entirely consistent with the c_4 court’s later decision that the consideration $\{f_5^\pi\}$, which was not present in the earlier case, provides a stronger reason for the opposite outcome. And there is no need either to try to modulate, or qualify, a court’s commitment to its own rules. The c_2 court, for example, is now able to take its own stated rule, $\{f_5^\delta\} \rightarrow \delta$, as thoroughly correct—representing, now, the idea that $\{f_5^\delta\}$ always counts as a reason for the defendant, though a reason that might elsewhere be defeated by stronger reasons to the contrary. The present account thus allows courts to take their own rules seriously: rather than placing these courts in the somewhat disingenuous position of having to qualify their commitment to excessively strong rules, the present account allows them to make a wholehearted commitment to rules that are themselves weaker, or defeasible.

The second problem with the standard view of distinguishing can best be described as computational, or perhaps information-theoretic. A case base is, among other things, a repository in which, ideally, a body of case law can be represented and updated efficiently. We have already studied the dynamics of case base update from the standpoint of the

³⁰See MacCormick (1987, p. 164–165).

present account. There we noted that the operation of distinguishing a previous rule can be accommodated very simply, as the mere addition of a new case to the previous case base. As we have seen, for example, beginning with the case base $\Gamma_1 = \{c_1, c_2\}$, the binding rule r_2 , or $\{f_5^\delta\} \rightarrow \delta$, can be distinguished in the context of the new fact situation X_3 simply by the addition to this case base of $\langle X_4, r_4, s_4 \rangle$, where X_4 is X_3 , r_4 is the rule $\{f_5^\pi\} \rightarrow \pi$, and s_4 is π . This operation of distinguishing can thus be accomplished simply by moving from the original case base to the new case base $\Gamma_1 \cup \{\langle X_4, r_4, s_4 \rangle\}$.

On the standard view, however, the operation of updating a case base by distinguishing a precedent rule is much more complex. Like overruling a precedent, distinguishing must be modeled as a destructive operation, involving the modification or removal of existing elements of the case base, not simply the addition of new information. In the case of our example, as we have seen, the process of distinguishing the c_2 court's original rule in the context of X_3 would lead to the replacement of $\{f_5^\delta\} \Rightarrow \delta$ with the modified rule $\{f_5^\delta, \neg f_5^\pi\} \Rightarrow \delta$, not just the addition of the new rule $\{f_5^\pi\} \Rightarrow \pi$.

Now, it might be argued that, although the replacement of the existing rule $\{f_5^\delta\} \Rightarrow \delta$ with the new rule $\{f_5^\delta, \neg f_5^\pi\} \Rightarrow \delta$ is, in fact, a destructive operation, the operation is sufficiently routine, or mechanical, as to cause no real difficulty. Indeed, much of the standard literature on distinguishing, which generally concentrates only on a unique case, suggests that the process is always this straightforward, involving nothing more than the addition to the premise of the previous rule of a single atomic formula or its negation. But that is not so, as we can see by considering even a very simple sequence of cases.

Suppose the court is confronted, first of all, with the situations $X_{12} = \{f_1^\pi, f_1^\delta\}$ and $X_{13} = \{f_1^\pi, f_2^\delta\}$, both of which it decides for the plaintiff on the basis of the rule $\{f_1^\pi\} \Rightarrow \pi$.

Next, imagine that the court is confronted with the situation $X_{14} = \{f_1^\pi, f_1^\delta, f_2^\delta\}$ and wishes to decide for the defendant; even though it had previously concluded that the factor f_1^π outweighs either f_1^δ or f_2^δ alone, the court might reasonably feel that f_1^π is outweighed by the combination of f_1^δ and f_2^δ taken together. In that case, in order to avoid being driven to an undesirable conclusion, the court would need to distinguish the previous rule, $\{f_1^\pi\} \Rightarrow \pi$, narrowing it with some minimal requirement necessary to separate the current situation from the earlier cases in which the rule was applied. What might this requirement be? Well, the court cannot require $\neg f_1^\delta$, that f_1^δ cannot be present, as an additional premise component, since the modified rule would then fail to support the earlier decision in the case of X_{12} ; nor can it require $\neg f_2^\delta$, that f_2^δ cannot be present, since the modified rule would then fail to support the earlier decision in the case of X_{13} . Instead, it seems, the court must require that f_1^δ and f_2^δ cannot both be present. This requirement could be represented only through a more complex formula along the lines of $\neg(f_1^\delta \wedge f_2^\delta)$, where the symbol \wedge stands for conjunction, leading to $\{f_1^\pi, \neg(f_1^\delta \wedge f_2^\delta)\} \Rightarrow \pi$ as the revised rule.

To take the example one step further, suppose the court is next confronted with the situation $X_{15} = \{f_1^\pi, f_1^\delta, f_3^\delta\}$, which it again wishes to decide for the defendant, feeling that the combination of f_1^δ and f_3^δ together again outweighs f_1^π . In order to avoid an undesirable conclusion for the plaintiff, once again, while still supporting the desired conclusions in each of the previous situations, the court must now supplement the original rule with the requirement, not only that f_1^δ and f_2^δ cannot both be present, but that f_1^δ and f_3^δ cannot both be present either. This new requirement would have to be represented with a still more complex formula along the lines of $\neg((f_1^\delta \wedge f_2^\delta) \vee (f_1^\delta \wedge f_3^\delta))$, where \vee stands for disjunction, now leading to $\{f_1^\pi, \neg((f_1^\delta \wedge f_2^\delta) \vee (f_1^\delta \wedge f_3^\delta))\} \Rightarrow \pi$ as a revised rule.

Even this simple sequence of cases is enough to show, then, that distinguishing a precedent rule, on the standard view, is not necessarily a routine, mechanical matter; the process may require some logical sophistication, and may lead to new rules of some complexity. The idea that distinguishing involves the explicit modification of a universal generalization in order to avoid unwanted conclusions seems especially odd in the area of case law, where, as is frequently noted, there are no canonical rule formulations, so that the increasingly complex precedent rules are not explicitly recorded. And the entire process compares unfavorably, I think, to the present account, where rules are defeasible, there is no explicit rule modification or destructive operations of any sort, and the priority relations among precedent rules is derived implicitly from the preference ordering on reasons.

5 Conclusion

My goal in this paper has been to provide a new account of precedential constraint, which unifies certain previous accounts from the literature—the rule, result, and reasons models. The account presented here is precise, and it is set out within a representational framework, due to Ashley and his colleagues, that has already shown itself to be useful in both analytic and empirical studies of legal precedent. I hope that the present account can be developed into a full theory of precedential reasoning. But the goal of precision has forced a number of limitations in the present treatment, which would need to be generalized in several directions before any such ambitions could be explored. I close simply by mentioning three of these.

To begin with, then, the fact situations at work here are specified only in terms of the factors they contain, not the factors they fail to contain; but often, the explicit absence of a

factor is as important to the meaning of a case as its presence. The current representation of a fact situation simply as a set of positive factors leaves us with no middle ground: if a factor is not listed as present in a fact situation, it must be regarded as absent. Indeed this perspective is often adopted by legal theorists: Raz, for example, suggests that it is reasonable to suppose that a case can be characterized as one in which *not-F* whenever “there is no record whether it was a case of *F* or of *not-F*.”³¹ This style of reasoning—from the absence of positive information to the presence of negative information—is known in the computer science literature as “closed world reasoning.” It is certainly appropriate in some situations: if an airline schedule does not explicitly list a flight from Pittsburgh to London on Saturday evening, we can conclude from this that there is no such flight; if I do not know that I have a brother, it is reasonable to conclude from this that I have no brother. But there are other situations in which closed world reasoning is much less appropriate: to take an extreme example, if I do not happen to know that the President of the United States is wearing a blue suit today, it would be unreasonable to conclude from this that he is, in fact, not wearing a blue suit.

Rather than assuming that closed world reasoning applies uniformly for legal factors, or that it uniformly fails, Ashley and Vincent Aleven suppose, more sensibly, that this form of reasoning applies for some factors but fails for others.³² In the domain of trade secrets law, for example, they argue that closed world reasoning can be applied to the factor representing the existence of an explicit confidentiality agreement between plaintiff and

³¹Raz (1979, p. 187).

³²See Aleven (1997, pp. 239–247) for a list of various legal factors to which closed world reasoning does or does not apply.

defendant—lacking positive information indicating such an agreement, that is, the court can legitimately conclude that there is none. But they deny that this form of reasoning can be applied to the factor representing the existence simply of a confidential relationship between plaintiff and defendant—even lacking explicit information indicating such a relationship, the court cannot conclude that no such relationship exists; the presumption of confidentiality is too strong.

If this analysis is correct, then the presence of a negative factor cannot be defined as nothing but the absence of the corresponding positive factor. A fact situation cannot therefore be represented simply as a list of positive factors, with all negative factors calculated, implicitly, through closed world reasoning. Instead, both positive and negative factors must, at times, be listed as basic components of the fact situation; fact situations must be allowed to contain both factors and their negations explicitly. And in that case, many of our fundamental ideas, including the crucial notions of a preference ordering on considerations and a strength ordering on cases, which currently apply to fact situations containing only positive factors, will have to be generalized to apply to these richer fact situations as well.

The remaining generalizations I would like to mention involve weakening two of simplifying assumptions underlying the present account.

First of all, as noted earlier, this account relies on the idea that every legal factor favors one side or the other, either the plaintiff or the defendant. This is, in many ways, a plausible assumption, especially for relatively mature areas of the law. It is hard to think of a factor that, while legally relevant, does not favor one side or another; and certainly, the analysis mentioned earlier of actual legal cases by Ashley and his colleagues involves only factors favoring some particular side. Still, there are arguments within moral philosophy suggesting

that the polarity of certain factors, the side they favor, might vary depending on the context in which they appear—that a particular factor might favor one side when taken together with one group of factors, and a different side when taken together with a different group.³³ The basic idea can be illustrated with an example entirely outside the moral or legal domain, by considering a situation in which an individual is trying to decide whether conditions are favorable for an afternoon run. It is easy to imagine that both extreme heat and rain might count as unfavorable factors, tending to rule out a run, but that a combination of heat and rain together is acceptable, perhaps even refreshing. On one natural interpretation, what this example suggests is that neither heat nor rain itself has any independent polarity with respect to the classification of a situation as favorable for running, since each of these features tends to make the situation less favorable in one context, when present alone, but more favorable in another, when both features are present together.

The issues surrounding examples like this are complicated, and of course, other interpretations are possible as well; perhaps what this particular example shows is that the basic factors involved in the domain are actually heat-without-rain and rain-without-heat, both of which would have negative polarity, and heat-and-rain-together, which would have positive polarity. Nevertheless, such examples, as well as other considerations from the literature, give life to the possibility that certain factors might have variable polarity, favoring different sides of an issue depending on the context in which they occur. If this turns out to be true, then several of our current ideas—including, once again, the crucial notions of a preference ordering on considerations and a strength ordering on cases, which currently apply

³³See Dancy (1993) and (2004) for arguments in favor of this position, and Hooker and Little (2000) for a collection of essays on the topic; my own views on Dancy's argument are developed in Horty (2007).

to fact situations containing only positive factors—will have to be generalized to allow for this possibility.³⁴

And second, as we have seen, the current account is based on the idea that precedential reasoning involves only a single step, proceeding from the factors present in a case directly to a decision, in favor of the plaintiff or the defendant, without moving through a series of intermediate legal concepts. In fact, this hard-headed perspective, approaching that of legal realism, still has some adherents. However, it has often been argued in the literature on artificial intelligence and law, especially by Karl Branting and by Henry Prakken and Giovanni Sartor, that an accurate model of legal reasoning must accord at least some meaning to the intermediate steps through which the outcome in a precedent case is determined.³⁵

In order to accommodate this perspective, the current framework would have to be generalized considerably. Rather than representing precedents as items of the form $\langle X, r, s \rangle$, where X is a set of factors, r is a rule, and s represent an immediate decision for the plaintiff or defendant, we would have to focus instead on structures of the form $\langle X, r, g \rangle$, with X and r as above, but where g represents yet another factor. Adapting Branting's phrase, we might refer to such an item as a *precedent constituent*, noting that, if the outcomes π and δ are considered as special factors, then the class of precedent constituents include our previous precedent cases. Such a precedent constituent could be thought of as representing a decision, on the basis of some particular set of factors, that some further factor should also be included in a description of the fact situation. An entire precedent case could then be represented as

³⁴Prakken and Sartor (1998) develop a model of precedential reasoning with polarity-free factors; the idea is explored in Horty (1999) as well.

³⁵See Branting (1991) and (1994), and Prakken and Sartor (1998).

a linked set of precedent constituents, beginning with a characterization of the fact situation through a set of base-level factors, proceeding through a series of higher-level legal concepts, and eventually arriving at a decision, based in part on these higher-level concepts, for the plaintiff or defendant. Incorporating this more detailed picture of legal reasoning into the present framework would be a complicated project, of course, but some ideas along these lines have already been developed in the work of Branting and Prakken and Sartor, noted above, as well as in my own previous work.³⁶

³⁶See Horty (1999).

A Cases and fact situations

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 is $\{f_1^\pi, f_2^\pi\} \rightarrow \pi$,
and where s_1 is π .
2. $c_2 = \langle X_2, r_2, s_2 \rangle$, where $X_2 = \{f_1^\pi, f_2^\pi, f_4^\pi, f_5^\delta, f_6^\delta\}$, where r_2 is $\{f_5^\delta\} \rightarrow \delta$, and
where s_2 is δ .
3. $c_3 = \langle X_3, r_3, s_3 \rangle$, where $X_3 = \{f_1^\pi, f_2^\pi, f_5^\pi, f_5^\delta, f_7^\delta\}$, where $r_3 (= r_2)$ is $\{f_5^\delta\} \rightarrow \delta$,
and where s_3 is δ .
4. $c_4 = \langle X_4, r_4, s_4 \rangle$, where $X_4 = X_3 = \{f_1^\pi, f_2^\pi, f_5^\pi, f_5^\delta, f_7^\delta\}$, where r_4 is $\{f_5^\pi\} \rightarrow \pi$,
and where s_4 is π .
5. $c_5 = \langle X_5, r_5, s_5 \rangle$, where where $X_5 = X_4 = X_3 = \{f_1^\pi, f_2^\pi, f_5^\pi, f_5^\delta, f_7^\delta\}$, where
 r_5 is $\{f_1^\pi\} \rightarrow \pi$, and where s_5 is π .
6. $c_6 = \langle X_6, r_6, s_6 \rangle$, where $X_6 = X_5 = X_4 = X_3 = \{f_1^\pi, f_2^\pi, f_5^\pi, f_5^\delta, f_7^\delta\}$, where
 r_6 is $\{f_7^\delta\} \rightarrow \delta$, and where s_6 is δ .
7. $c_7 = \langle X_7, r_7, s_7 \rangle$, where $X_7 = \{f_5^\pi, f_6^\pi, f_5^\delta, f_7^\delta\}$, where $r_7 (= r_2)$ is $\{f_5^\delta\} \rightarrow \delta$,
and where s_7 is δ .
8. $X_8 = \{f_1^\pi, f_1^\delta, f_2^\delta\}$.
9. $X_9 = \{f_1^\pi, f_2^\pi, f_1^\delta\}$.
10. $X_{10} = \{f_1^\pi, f_2^\pi, f_3^\pi, f_1^\delta\}$.
11. $X_{11} = \{f_1^\pi, f_2^\pi, f_1^\delta\}$.
12. $X_{12} = \{f_1^\pi, f_1^\delta\}$.
13. $X_{13} = \{f_1^\pi, f_2^\delta\}$.

14. $X_{14} = \{f_1^\pi, f_1^\delta, f_2^\delta\}$.

15. $X_{15} = \{f_1^\pi, f_1^\delta, f_3^\delta\}$.

B Observations and proofs

Observation 1 Let Γ be a case base containing two precedent cases of the form $\langle X, r, s \rangle$ and $\langle X, r', \bar{s} \rangle$. Then Γ is inconsistent.

Proof Suppose the case base Γ contains the cases $c = \langle X, r, s \rangle$ and $c' = \langle X, r', \bar{s} \rangle$. By our coherence constraints on rules and cases—particularly the requirements that the premise of a rule can contain only factors supporting its conclusion along with the requirement that the rule of a case must be applicable to its fact situation—we have (1) $Premise(r) \subseteq X^s$ and (2) $Premise(r') \subseteq X^{\bar{s}}$, and obviously (3) $Premise(r) \subseteq Premise(r)$ and (4) $Premise(r') \subseteq Premise(r')$. From (2) and (3), we have $Premise(r') <_c Premise(r)$ by Definition 2, and likewise $Premise(r) <_{c'} Premise(r')$ by (1) and (4). We therefore have both $Premise(r') <_{\Gamma} Premise(r)$ and $Premise(r) <_{\Gamma} Premise(r')$ by Definition 3, and so Definition 7 tells us that Γ is inconsistent. ■

Observation 2 Let Γ be a consistent case base, and suppose X is a fact situation in which none of the rules from $Rule(\Gamma)$ are applicable. Then $\Gamma \cup \{\langle X, r, s \rangle\}$ is also consistent, where r is any newly formulated rule applicable in X and supporting s as an outcome.

Proof Assume that Γ is consistent and that X is a fact situation in which none of the rules from $Rule(\Gamma)$ are applicable. Now suppose for contradiction that $\Gamma \cup \{c\}$ is not consistent, where $c = \langle X, r, s \rangle$ with r a newly formulated rule supporting s . Since Γ is consistent but $\Gamma \cup \{c\}$ is not, the preference relation derived from c must conflict with the preference relation derived from some other case already belonging to Γ —that is, there must be considerations A and B such that (1) $A <_c B$ and (2) $B <_{c'} A$, where $c' = \langle Y, r', \bar{s} \rangle$ is a

case from Γ . From (1), we have (3) $A \subseteq X^{\bar{s}}$ and (4) $Premise(r) \subseteq B$ by Definition 2, and likewise, from (2), we have (5) $B \subseteq Y^s$ and (6) $Premise(r') \subseteq A$. From (3) and (6), together with the fact that $X^{\bar{s}} \subseteq X$ we have $Premise(r') \subseteq X$, and of course r' belongs to $Rule(\Gamma)$, so that some rule from $Rule(\Gamma)$ is applicable to X , contrary to assumption. ■

Observation 3 Let Γ be a consistent case base with $Rule(\Gamma)$ the derived set of rules, and suppose X is a fact situation in which some rule r from $Rule(\Gamma)$, supporting the outcome s , is binding. Then the case base $\Gamma \cup \{\langle X, r, s \rangle\}$ is also consistent.

Proof Assume that Γ is consistent with X is a fact situation in which the rule r from $Rule(\Gamma)$, supporting the outcome s , is binding. Now suppose for contradiction that $\Gamma \cup \{c\}$ is not consistent, where $c = \langle X, r, s \rangle$. Since Γ is consistent but $\Gamma \cup \{c\}$ is not, the preference relation derived from c must conflict with the preference relation derived from some other case already belonging to Γ —that is, there must be considerations A and B such that (1) $A <_c B$ and (2) $B <_{c'} A$, where $c' = \langle Y, r', \bar{s} \rangle$ is a case from Γ . From (1), we have (3) $A \subseteq X^{\bar{s}}$ and (4) $Premise(r) \subseteq B$ by Definition 2, and likewise, from (2), we have (5) $B \subseteq Y^s$ and (6) $Premise(r') \subseteq A$. From (3) and (6), together with the fact that $X^{\bar{s}} \subseteq X$, we have $Premise(r') \subseteq X$, so that the rule r' is applicable in the situation X . From (4) and (5), we have (6) $Premise(r) \subseteq Y^s$, and obviously (7) $Premise(r') \subseteq Premise(r')$. By Definition 2, then, (6) and (7) tell us that $Premise(r) <_{c'} Premise(r')$, from which it follows that $Premise(r) <_{\Gamma} Premise(r')$. Since the rule r' from $Rule(\Gamma)$, which supports a different conclusion, is both applicable in X and stronger than r , it follows from Definition 5 that r is trumped. By Definition 6, therefore, the rule r is not binding, contrary to assumption. ■

Observation 4 Let Γ be a case base containing two precedent cases of the form $\langle X, r, s \rangle$

and $\langle Y, r', \bar{s} \rangle$ where $X \leq^s Y$. Then Γ is inconsistent.

Proof Suppose the case base Γ contains the cases $c = \langle X, r, s \rangle$ and $c' = \langle Y, r', \bar{s} \rangle$ where $X \leq^s Y$. By our coherence constraints on rules and cases—particularly the requirements that the premise of a rule can contain only factors supporting its conclusion along with the requirement that the rule of a case must be applicable to its fact situation—we have (1) $Premise(r) \subseteq X^s$ and (2) $Premise(r') \subseteq Y^{\bar{s}}$, and obviously (3) $X^{\bar{s}} \subseteq X^{\bar{s}}$ and (4) $Y^s \subseteq Y^s$. Because $X \leq^s Y$, it follows from Definition 9 that $X^s \subseteq Y^s$, and so from (1) that (5) $Premise(r) \subseteq Y^s$. From (3) and (5), we therefore have (6) $X^{\bar{s}} <_c Y^s$ by Definition 2. Because $X \leq^s Y$, again, we also have $Y \leq^{\bar{s}} X$ by the duality property of the \leq^s relation. It therefore follows from Definition 9 that $Y^{\bar{s}} \subseteq X^{\bar{s}}$, and so from (2) that (7) $Premise(r') \subseteq X^{\bar{s}}$. From (4) and (7), we have (8) $Y^s <_{c'} X^{\bar{s}}$ by Definition 2. From (6) and (8) we therefore have both $X^{\bar{s}} <_{\Gamma} Y^s$ and $Y^s <_{\Gamma} X^{\bar{s}}$ by Definition 3, and so Definition 7 tells us that Γ is inconsistent. ■

Observation 5 Let Γ be a consistent case base and X a new fact situation confronting the court, and suppose a fortiori constraint requires the court to reach a decision for the side s in the situation X . Then precedential constraint also requires the court to reach a decision for the side s .

Proof Suppose Γ is a consistent case base, with X a fact situation in which the court is required by a fortiori constraint to reach a decision for the side s . By Definition 10, then, there must be some case of the form $\langle Y, r, s \rangle$ in Γ for which $Y \leq^s X$. In order to show that precedential constraint also requires a decision for s , we must show, by Definition 8, that the result of augmenting Γ with a contrary decision—that is, with a case of the form $\langle X, r', \bar{s} \rangle$,

where r' is some rule supporting the outcome \bar{s} —would be inconsistent. But this follows at once from Observation 4. ■

Observation 6 Consider only cases $\langle X, r, s \rangle$ in which the rule r has the form $X^s \rightarrow s$. Let Γ be a consistent set of such cases, and suppose that Y is a new fact situation confronting the court. Then a fortiori constraint requires the court to reach a decision for some particular side in this situation if and only if precedential constraint requires a decision for the same side.

Proof We have already seen, in Observation 5, that following a fortiori constraint implies following precedential constraint in general, and so it remains only to show that following precedential constraint implies following a fortiori constraint under the assumption that any rule r in a case of the form $\langle X, r, s \rangle$ itself has the form $X^s \rightarrow s$. Under this assumption, then, suppose that Γ is a consistent case base against the background of which precedential constraint requires a decision for s in the new fact situation Y . What this means, according to Definition 8, is that the result of augmenting Γ with a contrary decision, of the form $c = \langle Y, r, \bar{s} \rangle$, would then be inconsistent—that is, there must be considerations A and B such that (1) $A <_c B$ and (2) $B <_{c'} A$, where $c' = \langle X, r', s \rangle$ is a case already belonging to Γ . From (1), we have by Definition 2 both (3) $A \subseteq Y^s$ and also (4) $Y^{\bar{s}} \subseteq B$, since $Premise(r)$ has the form $Y^{\bar{s}}$ by our assumption; and likewise, from (2) we have both (5) $B \subseteq X^{\bar{s}}$ and also (6) $X^s \subseteq A$, since $Premise(r')$ has the form X^s . Evidently, (3) and (6) tell us that $X^s \subseteq Y^s$, and (4) and (5) that $Y^{\bar{s}} \subseteq X^{\bar{s}}$. We therefore have $X \leq^s Y$ by Definition 9, and so, since c' belongs to Γ , it follows from Definition 10 that a fortiori constraint requires a decision for s in the case of Y . ■

References

- [Aleven and Ashley, 1997] Vincent Aleven and Kevin Ashley. 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, 1997.
- [Aleven, 1997] Vincent Aleven. *Teaching Case-Based Argumentation Through a Model and Examples*. PhD thesis, Intelligent Systems Program, University of Pittsburgh, 1997.
- [Alexander and Sherwin, 2001] Larry Alexander and Emily Sherwin. *The Rule of Rules: Morality, Rules, and the Dilemmas of Law*. Duke University Press, 2001.
- [Alexander and Sherwin, 2007] Larry Alexander and Emily Sherwin. Judges as rule makers. In Douglas Eldin, editor, *Common Law Theory*, pages 27–50. Cambridge University Press, 2007.
- [Alexander, 1989] Larry Alexander. Constrained by precedent. *Southern California Law Review*, 63:1–64, 1989.
- [Alexander, 1996] Larry Alexander. Precedent. In Dennis Patterson, editor, *A Companion to Philosophy of Law and Legal Theory*, pages 503–513. Blackwell Publishers, 1996.
- [Ashley, 1989] Kevin Ashley. Toward a computational theory of arguing with precedents: accomodating multiple interpretations of cases. In *Proceedings of the Second International Conference on Artificial Intelligence and Law (ICAIL-89)*, pages 93–110. The Association for Computing Machinery Press, 1989.

- [Ashley, 1990] Kevin Ashley. *Modeling Legal Argument: Reasoning with Cases and Hypotheticals*. The MIT Press, 1990.
- [Branting, 1991] L. Karl Branting. Reasoning with portions of precedents. In *Proceedings of the Third International Conference on Artificial Intelligence and Law (ICAIL-91)*, pages 145–154. ACM Press, 1991.
- [Branting, 1994] L. Karl Branting. A computational model of *ratio decidendi*. *Artificial Intelligence and Law*, 2:1–31, 1994.
- [Cross, 1968] Rupert Cross. *Precedent in English Law*. Oxford University Press, 1968. Second edition.
- [Dancy, 1993] Jonathan Dancy. *Moral Reasons*. Basil Blackwell Publisher, 1993.
- [Dancy, 2004] Jonathan Dancy. *Ethics Without Principles*. Oxford University Press, 2004.
- [Goodhart, 1930] Arthur Goodhart. Determining the *ratio decidendi* of a case. *Yale Law Journal*, 40:161–183, 1930.
- [Hooker and Little, 2000] Brad Hooker and Margaret Little. *Moral Particularism*. Oxford University Press, 2000.
- [Horty, 1999] John Horty. Precedent, deontic logic, and inheritance. In *Proceedings of the Seventh International Conference on Artificial Intelligence and Law (ICAIL-99)*, pages 63–72. The Association for Computing Machinery Press, 1999.
- [Horty, 2004] John Horty. The result model of precedent. *Legal Theory*, 10:19–31, 2004.

- [Horty, 2007] John Horty. Reasons as defaults. *Philosopher's Imprint*, 7(3), 2007.
- [Lamond, 2005] Grant Lamond. Do precedents create rules? *Legal Theory*, 11:1–26, 2005.
- [Levi, 1949] Edward Levi. *An Introduction to Legal Reasoning*. The University of Chicago Press, 1949.
- [MacCormick, 1987] Neil MacCormick. Why cases have *rationes* and what these are. In Laurence Goldstein, editor, *Precedent in Law*, pages 155–182. Oxford University Press, 1987.
- [Prakken and Sartor, 1998] Henry Prakken and Giovanni Sartor. Modelling reasoning with precedents in a formal dialogue game. *Artificial Intelligence and Law*, 6:231–287, 1998.
- [Raz, 1975] Joseph Raz. *Practical Reasoning and Norms*. Hutchinson and Company, 1975. Reprinted by Oxford University Press, 2002.
- [Raz, 1979] Joseph Raz. *The Authority of Law*. Oxford University Press, 1979.
- [Rissland, 1990] Edwina Rissland. Artificial intelligence and law: stepping stones to a model of legal reasoning. *Yale Law Review*, 99:1957–1981, 1990.
- [Schauer, 1989] Frederick Schauer. Is the common law law? *California Law Review*, 77:455–471, 1989.
- [Schauer, 1991] Frederick Schauer. *Playing by the Rules: A Philosophical Examination of Rule-Based Decision-Making in Law and Life*. Oxford University Press, 1991.

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