

A factor-based definition of precedential constraint

John F. Horty · Trevor J. M. Bench-Capon

Published online: 14 June 2012
© Springer Science+Business Media B.V. 2012

Abstract This paper describes one way in which a precise *reason model* of precedent could be developed, based on the general idea that courts are constrained to reach a decision that is consistent with the assessment of the balance of reasons made in relevant earlier decisions. The account provided here has the additional advantage of showing how this reason model can be reconciled with the traditional idea that precedential constraint involves rules, as long as these rules are taken to be defeasible. The account presented is firmly based on a body of work that has emerged in AI and Law. This work is discussed, and there is a particular discussion of approaches based on theory construction, and how that work relates to the model described in this paper.

Keywords Case-based reasoning · Precedent · Theory construction · Factor-based reasoning · Rationales

1 Introduction

The doctrine of precedent, as it has evolved within the common law, has at its heart a form of reasoning—broadly speaking, a *logic*—according to which the decisions that relevant courts have previously taken in particular cases somehow generalize to constrain subsequent decisions, while still allowing courts a degree of freedom in responding to fresh circumstances when making these decisions. The techniques for arguing on the basis of precedent are taught early on in law schools, mastered with

J. F. Horty
Philosophy Department and Institute for Advanced Computer Studies, University of Maryland,
College Park, MD 20742, USA

T. J. M. Bench-Capon (✉)
Department of Computer Science, University of Liverpool, Liverpool L69 3BX, UK
e-mail: tbc@csc.liv.ac.uk

relative ease, and applied on a daily basis by legal practitioners; these techniques have also been the subject of investigation in AI and Law through three decades, resulting in several interesting implemented systems. It has, however, proved to be difficult to arrive at a precise theoretical understanding of the doctrine itself, a clear articulation of the underlying logic.

Our purpose in this paper¹ is to provide a formal framework within which we can begin to address this problem. It will build on previous work in AI and Law, so as to consolidate and bring together various strands of this previous work. We 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 courts for developing the law?

The view we present will be contrasted with two other views, or models, of precedential constraint appearing in literature from legal theory. 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 justification of the decision which can serve as an explanation of how that decision was reached. This explanation will often refer to a rule derived from a past case, or make a pronouncement which may be construed as a rule which is applicable to future cases. 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 a subsequent decision when the rule contained in that precedent applies to the fact situation confronted in making the subsequent decision.

A particularly strict version of the rule model has been advanced by Alexander (1989), and developed by Alexander and Sherwin (2001).² According to this strict rule model, the constraints carried by precedent rules are very nearly absolute, with the result that the freedoms given to 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 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 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 Simpson (1961) and Raz (1979).

¹ This paper is in part based on a revised version of work originally presented (Horty 2011a, b). This paper extends those papers by firmly locating the work developed from them in the context of AI and Law, and providing extensive discussion of points of similarity and difference with previous approaches to draw out its contribution to AI and Law.

² 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; We use the term more broadly to apply to any model in which precedential constraint is thought to be carried by rules.

The second view of precedential constraint to be considered here, recently suggested by Lamond (2005), may be termed the *reason model*. According to Lamond, what is most important about a decision in a precedent case is, not the particular rule it sets out, but instead the court's identification of reasons for and against a particular decision and its assessment of the balance of reasons presented by the facts of the case under consideration. The requirement of precedential constraint can then be defined as follows: a court is constrained to reach a decision that is consistent with the earlier assessments of the balance of reasons. A principal goal of this paper is to provide a way of making this idea precise.

The account presented here has the advantage of showing how these two theoretical models of precedent—the rule model and the reason model—can be unified; it helps us see what is correct in each of these views, and how they are related. The unification is achieved by interpreting the rules contained in precedent cases, not as strict rules, but as default, or defeasible, rules, with reasons serving as the premises of these default rules. This analysis of reasons as the premises of default rules 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 has been explored extensively in Horty (2007, 2012).³

A 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 underlying judicial decisions. Like the set of rules contained in precedent cases, this priority ordering 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. Once this priority ordering on reasons has been made explicit, the notion of consistency with past decisions, and so precedential constraint, can then be defined.

The paper is organized as follows. After beginning, in the second section, with a brief summary of relevant developments in AI and Law, we move on, in the third section, to lay out basic ideas and notation. The fourth section defines our central ordering relation on reasons, which is then related to rules in the fifth section. In the sixth section, we appeal to this ordering to present our core theory of precedential constraint, and also to explore some of the issues surrounding a transitive extension of the core theory. The seventh section focuses on case base dynamics—the effects of following or distinguishing a precedent. We conclude with a discussion and comparison of the approach presented here with other approaches from AI and Law. An appendix contains verifications of some of the observations mentioned in the paper.

³ The idea is also, of course, closely related to the work of Jaap Hage on reason based logic. In Hage (1993) and Hager (1997, 2005), Hage provided a logical account of case-based reasoning on the basis of a reason—rather than a rule-based model. In Hage (1993), conflicts were resolved either through explicitly stated general principles, such as *rule1 is more specific than rule2*, or using numeric weights. The later books identify precedent cases as one sources of reasons why a set of reasons may be considered stronger than another, but, like Prakken and Sartor (1998) and Roth and Verheij (2004), these are always used to justify a form of a fortiori reasoning. The need to be able to go beyond a fortiori reasoning and explain how conclusions can be drawn when the precedents are not unequivocal is one of the motivations of this paper.

2 Reasoning with cases in AI and Law

Reasoning with cases has long been a major concern of AI and Law. We begin by reviewing some key pieces of work in this developing story, and in particular those that are adopted and adapted in the account we present here.

In the late nineteen eighties two central approaches to the topic emerged. One, corresponding to the rule model from legal theory, involved the representation of decisions in cases as supplying sufficient conditions to allow terms used in legislation to be applied on the basis of a set of case facts.⁴ Such conditions were typically used in conjunction with a formalisation of the legislative provisions, and the whole executed using Prolog, or a similar reasoning engine. Where there were conflicts in these conditions, the priority between them was resolved by the programmer, by, for example, ordering the clauses to give the desired behaviour.⁵ This approach is still in use today, and underlines many systems that use rule based representations of knowledge, and a knowledge base intended to summarise expertise in the area. Such systems, however, have no explicit representation of cases, and do not reason with cases, but merely apply the results of previous analysis of cases. Additionally they impose a fixed interpretation, both in the rule summarising the decision, and in the priority given to the various rules in cases of conflict.

For systems with an explicit representation of cases, and reasoning with cases as part of the system itself, we will take as a starting point the HYPO system introduced by Rissland and Ashley in (1987) and most fully described in Ashley (1990). In HYPO there was a representation of facts, and of *dimensions*. Dimensions applied or did not apply to a case on the basis of the facts, and if applicable, represented an issue which required consideration; the facts of the case would then indicate that the case was at some particular point on the dimension. For example one dimension in HYPO, which was developed in domain of US Trade Secret Law, concerned the security measures taken by the plaintiff. At one end of the dimension, the plaintiff will have taken no security measures, and the dimension will favour the defendant, while at the other extreme, where the plaintiff has taken every conceivable measure, the dimension will favour the plaintiff. Most cases will fall at a point between these extremes, and one may look at precedents to see which party is favoured at that particular point on the dimension. Dimensions were particularly useful for hypothetical reasoning, since one could posit cases with the facts adjusted along one or more dimensions to be more in favour of one side or the other.

The notion of dimensions was adapted in the successor to HYPO, the CATO system, developed by Ashley along with Vincent Aleven in Aleven and Ashley (1997) and Aleven (1997), which relied instead on the idea of *factors*.⁶ Factors correspond to particular points on dimensions. In consequence, they can be either present in or absent from a case, but always favour one side or the other when

⁴ This approach can be found, for example in, Bench-Capon (1991).

⁵ See Schild and Herzog (1993) for an approach along these lines.

⁶ CATO was designed to support the teaching of reasoning with cases to law students; the two cited papers also report the results of using the system with students.

present. For example, we may recognise four points on the dimension of security measures: the two extremes, *took no measures* and *took strong measures*; and two intermediate points, *took minimal measures* and *took substantial measures*. We might then use our knowledge of decided cases to recognise that taking any measures at all appears to favour the plaintiff. Now one of these points on the dimension, *took no measures*, can be recognized as a pro-defendant factor, and the other three can be grouped together into a single pro-plaintiff factor, *took some measures*. It is this notion of a factor that has been used in much subsequent work in AI and Law, and it is this notion that we will use in this paper.

Two points should be noted. First, factors are not just facts, but are ascribed to cases on the basis of legally significant facts or patterns of facts. The second point is that some cases actually turn on whether a particular point on a dimension favours the plaintiff or the defendant. Such cases cannot be considered using factors, since factors require a commitment to one side or the other in the representation itself.⁷ An additional feature of CATO was that factors were placed in a factor hierarchy, with each factor being related, positively or negatively, to one or more abstract factors: for example taking security measures would be positively related to the more abstract factor *information trade secret*, since taking security measures would be an indication that the plaintiff at least believed the information to be worth protecting.

At the time at which CATO appeared, the rule based approaches and the case based approaches were considered rather separate, even opposed, and little comparison could be made between them. Soon, however, people started to consider whether formal work on defeasible reasoning could be used to explore legal case based argumentation in the style of HYPO and CATO. One early attempt was made in Loui et al. (1993). In that work, however, the authors were intent on using a purely syntactic criterion to decide between competing arguments. The criterion they used was to prefer the more specific rule, and they identified the most-on-point case of HYPO and CATO as the source of the most specific applicable rule. Thus precedents were seen only as the source of rules, not of preferences and priorities between rules. The notion of *rationales*, also introduced in Loui et al. (1993) and further developed in Loui and Norman (1995) will be discussed in Sect. 8 below. Prakken and Sartor (1998) went beyond this purely syntactic criterion by using CATO style precedents to generate both rules and priorities between them. They proposed representing a decision in a case represented in terms of factors as a set of three defeasible rules. The idea was that a case is to be mapped into two conflicting defeasible rules and an expression of priority between these rules. To produce the rules, the factors present in a case are split into the pro-plaintiff factors and the pro-defendant factors. The first rule then says that the set of pro-plaintiff factors, taken together, favors the plaintiff; the second rule says that the set of pro-defendant factors favors the defendant. The third rule expresses priority, and can vary depending on the decision reached in the case: if the case was found for the plaintiff, then the third rule states that the first rule, favoring the plaintiff, has greater priority than the second, favoring the defendant; if the case was found for the

⁷ See Bench-Capon and Rissland (2001) for a discussion of this issue.

defendant, then the third rule states that the second rule has greater priority than the first.

It is important to note that Prakken and Sartor's proposal presents a very safe approach: the rules for both sides are the strongest that can be formed for the sides from the available factors. Yet, while it is reasonable that the losing side should be represented as having advanced its strongest position, it may be (and often is) the case that a weaker rule would have been enough for the winning side. Suppose, for example, that a case involves four pro-plaintiff factors and one pro-defendant factor. Then the import of the case would be represented by three rules: one stating that the four pro-plaintiff factors favor the plaintiff, another stating that the single pro-defendant factor favors the defendant, and a third stating that the first rule has higher priority than the second. But, while it is certain, given this decision, that the rule containing the four pro-plaintiff factors has higher priority than the rule containing the single pro-defendant factor, it may well be that a weaker rule would have done just as well. Perhaps a rule containing only one particular pro-plaintiff factor would better reflect the court's decision; such a rule could be then be applied in a broader range of future cases. By limiting their attention only to the strongest rule possible for the winning side of any case, Prakken and Sartor essentially explain only a fortiori reasoning.⁸ The current proposal is motivated, in part, by the desire to move beyond simple a fortiori reasoning.

3 Factors, rules, and cases

We follow the line of work described in the preceding section 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 stands, once again, for a legally significant fact or pattern of facts. Factors are always taken as 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. Of course, it must be noted also that the mere ability to understand and analyse a case in terms of the factors it contains itself requires a significant degree of legal expertise, which is presupposed here. As with much work in the tradition of HYPO and its successors, we take the description of cases in terms of factors as our input: we are concerned here only with the subsequent reasoning.⁹

⁸ The approach was further developed in Bench-Capon (1999), which represented the power sets of pro-plaintiff and pro-defendant factors as separate lattices, and then joined them into a partial order using the precedent relations derived from the priority rules; in this way *all* factor combinations could be systematically represented. A different approach to a fortiori reasoning, developed in the setting of legal theory, can be found in Horty (2004).

⁹ This is not to understate the challenges involved in producing an analysis of cases in terms of factors. Certainly judges rarely if ever explicitly refer to factors, in the sense used here, in their decisions. Sometimes they do use some standard phrases (e.g. 'imminent danger' in the series of cases in Levi 1949) which is a good indicator of a factor. For the domain of CATO, Trade Secrets law, treatises such as the Restatement of Torts were used to supplement the decisions and were found to give a good guide to the

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. Since 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.¹⁰ 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 $Outcome(c) = s$.

Because reasoning with factors 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 $\bar{\delta} = \pi$.¹¹ 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.¹² 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, in the absence of any conflicting

Footnote 9 continued

considerations relevant to the decisions, and hence the factors that needed to be looked for. Automatic analysis of decisions in terms of factors has also been explored by Bruninghaus and Ashley in (2005) and Ashley and Brüninghaus (2009). Although there is no large scale validation of the approach with respect to legal practice, empirical studies in Alevén (1997) showed that the approach was effective in teaching law students how to reason with cases.

¹⁰ For the purpose of this paper, we take it that the rule underlying a court's decision is plain, ignoring the extensive literature on methods for determining the rule, or *ratio decidendi*, of a case, which are rarely if ever explicitly stated by judges in their decisions. Validation of such rules would pose problems which are beyond the scope of this paper. We will also assume that a case always contains a single rule, ignoring situations in which a judge might offer several rules for a decision, or in which a court reaches a decision by majority, with different judges offering different rules, or in which a judge might simply render a decision in a case without setting out any general rule at all.

¹¹ This view of reasoning with factors is challenged by work such as Roth and Verheij (2004), which uses a chain of reasoning through a tree of successively more abstract factors. We will discuss this work further in Sect. 8.

¹² Note that, as in Prakken and Sartor (1998), the factors in the antecedent of the rule all favour the same side.

considerations, 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 , or that the rule gives an argument for deciding for s .

This connection between rules from precedent cases 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; for example, I might be called upon to save a life, or, more prosaically, attend an important meeting. 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 or a reason with special moral force, since it is based on a promise, but still a reason that might be defeated by stronger reasons, or perhaps excluded from consideration entirely.¹³

The idea behind the current account is that precedent rules work in exactly the same way, identifying legal reasons that support particular decisions. 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—or trumped, as we will say—by a stronger legal reason favoring the opposite side, in a way that will be explained shortly.

Let us return, now, to our concept of a precedent case $c = \langle X, r, s \rangle$, which we defined earlier as containing a fact situation X along with a rule r leading to the outcome s . In order for this concept to make sense, we impose 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: $Conclusion(r) = Outcome(c)$.¹⁴

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

¹³ This general picture is described in much more detail in Horty (2007), which appeals to techniques from nonmonotonic logic to develop a detailed theory of reasons as the premises 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).

¹⁴ Note that this is a departure from Prakken and Sartor (1998), in which precedents also gave rise to a rule for the losing side, albeit one with lower priority than the rule for the winning side.

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 by applying or introducing 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.

4 Ordered reasons

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 will be taken to represent the common law in some area, and to constrain the decisions of future courts.¹⁵ 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 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, of course, customary 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 second concept we introduce—a preference relation on reasons—is much less common as a focus of attention.¹⁶

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 courts must respect the decisions previously taken by relevant courts. So what information is actually carried by the earlier decision in the case of c_1 ; what is the court telling us with this decision? Well, two things, at least. First of all, by

¹⁵ Although we focus in this paper on which might be called “pure” common law, we believe the model would also apply to cases in which precedential reasoning is used to interpret statutory or constitutional language, or even in civil law jurisdictions. For the relevance of precedents to such reasoning see, for example, the papers in MacCormick and Summers (1997).

¹⁶ Although preference relations are of central importance in AI and Law treatments, such as those of Prakken and Sartor (1998) and Bench-Capon and Sartor (2003), the preferences defined there typically relate rules (and values), rather than reasons. Expressing preferences on reasons has similar advantages to the use of preferences between reasons rather than rules in the treatment of accrual of arguments; see Modgil and Bench-Capon (2010).

appealing to the rule r_1 , the court is telling us that the premise of this rule—that is, $Premise(r_1)$, or $\{f_1^\pi, f_2^\pi\}$ —is a sufficient 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 is preferred to whatever other reasons the case might present that favor the defendant.

To make this point precisely, let us now define a *legal reason* as a set of factors uniformly favoring one side or the other. For example, $\{f_1^\pi, f_2^\pi\}$ is a reason favoring the side π , a decision for the plaintiff, while $\{f_1^\delta, f_2^\delta\}$ is a reason favoring the side δ , a decision for the defendant; but according to this definition $\{f_1^\pi, f_1^\delta\}$ is not a reason at all, since the factors contained in this set do not uniformly favor either side. We can say that a fact situation *presents* a reason if all the factors from that reason are contained in that fact situation, and that a case presents a reason if its fact situation does so. Finally, if X and Y are reasons favoring the same side, we can say that Y is *at least as strong as* X whenever Y contains all the factors contained by X —whenever, that is, $X \subseteq Y$.¹⁷

Returning to our example, the case c_1 , as we have seen, 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 reason 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 a relevant court has previously decided for the plaintiff on the grounds of the reason provided by $Premise(r_1)$, even in the face of the reason provided by X_1^δ for the defendant, it seems to follow as a consequence of the court's decision that the reason $Premise(r_1)$ for the plaintiff is preferred to the reason X_1^δ for the defendant—that is, that the reason $\{f_1^\pi, f_2^\pi\}$ is preferred to the reason $\{f_1^\delta, f_2^\delta, f_3^\delta, f_4^\delta\}$. If we introduce the symbol $<_{c_1}$, to represent the preference relation on reasons that is derived from the particular case c_1 , then this consequence of the court's decision can be put more formally as the claim that $\{f_1^\delta, f_2^\delta, f_3^\delta, f_4^\delta\} <_{c_1} \{f_1^\pi, f_2^\pi\}$, or equivalently, that $X_1^\delta <_{c_1} Premise(r_1)$.¹⁸

As far as the preference ordering goes, then, the previous decision tells us at least that $X_1^\delta <_{c_1} Premise(r_1)$, but is it telling us anything else? Perhaps not explicitly, but implicitly, yes. For if the reason $Premise(r_1)$ for the plaintiff is preferred to the reason X_1^δ for the defendant, then surely any reason for the plaintiff that is at least as strong as $Premise(r_1)$ must likewise be preferred to X_1^δ , and just as surely, $Premise(r_1)$ must be preferred to any reason for the defendant that is at least as weak as X_1^δ . As we have seen, a reason Z for the plaintiff is at least as strong as $Premise(r_1)$ if it contains all the factors contained by $Premise(r_1)$ —that is, if $Premise(r_1) \subseteq Z$. And it is natural to conclude, likewise, that a reason W for the

¹⁷ Prakken (2005) gives an example in which an additional *fact* may weaken a case: such a fact would mean that different *factors* are applicable to the cases and so this principle is not threatened. Thus, in his example, if hot weather favours jogging and wet weather favours not jogging but hot and wet together favour jogging, we have three factors, *hot*, *wet* and *hot-and-wet*, based on the two facts *hot* and *wet* rather than two factors and their combination.

¹⁸ Note that the rejected reason is the same as the antecedent in the losing rule formed using the method of Prakken and Sartor (1998).

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 previous decision in c_1 , not only that $X_1^\delta <_{c_1} \text{Premise}(r_1)$, but that $W <_{c_1} Z$ whenever W is at least as weak a reason for the defendant as X_1^δ and Z is at least as strong a reason for the plaintiff as $\text{Premise}(r_1)$ —whenever, that is, $W \subseteq X_1^\delta$ and $\text{Premise}(r_1) \subseteq Z$. To illustrate: from the court's explicit decision that $\{f_1^\delta, f_2^\delta, f_3^\delta, f_4^\delta\} <_{c_1} \{f_1^\pi, f_2^\pi\}$, we can conclude also that $\{f_3^\delta, f_4^\delta\} <_{c_1} \{f_1^\pi, f_2^\pi, f_5^\pi\}$, for example.¹⁹

This line of argument leads to the following definition of the preference relation among reasons 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 reasons. Then the relation $<_c$ representing the preferences on reasons derived from the case c is defined by stipulating that $W <_c Z$ if and only if $W \subseteq X^s$ and $\text{Premise}(r) \subseteq Z$.

It is easy—indeed, trivial—to verify that the preference relation derived from any particular case c is transitive: whenever $X <_c Y$ and $Y <_c Z$, it follows that $X <_c Z$. It is not, however, a connected relation: we do not invariably have either $X <_c Y$ or $Y <_c X$ —the case c may tell us nothing at all about the relative strength of X and Y . To illustrate by returning to c_1 , we do not have either $\{f_1^\delta\} <_{c_1} \{f_1^\pi\}$ or $\{f_1^\pi\} <_{c_1} \{f_1^\delta\}$, for example.

Once we have defined the preference relation derived from a single case, we can then introduce a preference relation $<_\Gamma$ derived from an entire case base Γ in the natural way, by stipulating that one reason 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 reasons. Then the relation $<_\Gamma$ representing the preferences on reasons derived from the case base Γ is defined by stipulating that $W <_\Gamma Z$ if and only if $W <_c Z$ for some case c from Γ .

It is worth emphasizing that the derived preference relation $<_\Gamma$ is very weak, both formally and conceptually. From a formal standpoint, we should note that the preference relation derived from an entire case base, like that derived from a single case, is not connected: again we may have neither $X <_\Gamma Y$ nor $Y <_\Gamma X$. More surprisingly, this new relation is not transitive either: $X <_\Gamma Y$ and $Y <_\Gamma Z$ does not entail $X <_\Gamma Z$. We will return in the next section to consider the issues surrounding transitivity in more detail.

The preference relation derived from a case base is conceptually weak as well, in the sense that it might reflect very few of our ordinary judgments about strength relations among reasons. 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.²⁰ The plaintiff is the

¹⁹ Capturing this style of reasoning was the motivation of Bench-Capon (1999).

²⁰ This example is modeled on some hypothetical cases considered by Prakken and Sartor (1998).

individual's home country, which would like to collect tax on her income; the defendant is the individual, who would prefer to pay her 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 her old job and is now employed by a company in the foreign country; she has sold her old home and purchased a new home in the foreign country; she has sold her old car and both purchased and registered a new car in the foreign country. Suppose these three factors are represented as f_1^π , 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 her home country, which she uses while visiting her parents; this factor is f_1^π .²¹

Can we now assume that the reason $\{f_1^\delta, f_2^\delta, f_3^\delta\}$ favoring the defendant should be preferred to the reason $\{f_1^\pi\}$ favoring the plaintiff? Intuitively, it would seem so: surely, the mass of information about employment, residence, and automobile registration should outweigh some stray fact about bicycle registration. But as a matter of precedential constraint, not necessarily. Unless the case base Γ contains a previous case in which bicycle registration was actually compared to at least one of the factors supporting the defendant, and found to be less weighty, we will not have $\{f_1^\pi\} <_\Gamma \{f_1^\delta, f_2^\delta, f_3^\delta\}$. The present approach thus reflects a broadly positivist view of precedential constraint, according to which the legally sanctioned preference relations among reasons must have a basis, not simply in our everyday intuitions about which reasons are stronger than which, but in the acts of an appropriate legal authority—here, a court's decision in some precedent case.²² Of course, it is likely in the present situation that the court confronting this case would be guided by similar intuitions concerning weight of the conflicting reasons, and therefore decide for the defendant, perhaps using only the factor concerning automobile registration, f_2^δ . As a result of this decision, the intuitive assessment would be given legal standing, and it would then hold, once the case is decided and the case base is updated accordingly, that $\{f_1^\pi\} <_\Gamma \{f_2^\delta\}$, and in consequence $\{f_1^\pi\} <_\Gamma \{f_1^\delta, f_2^\delta, f_3^\delta\}$. This is the genius of the common law—that it provides a mechanism through which our ordinary intuitions about the relative importance of various reasons are gradually filtered into legal doctrine, on an incremental basis, in reaction to particular circumstances.

²¹ Note that the numbers of factors in a set should not be important. One factor may easily outweigh several factors. It is important that this is so, since the number of factors is always, to a greater or lesser extent, an artifact of the analysis: whether the facts of a case are represented as one factor or several will depend on the factors chosen by the analyst. We are, of course, relying on the factors capturing all significant nuances: for example whether the individual left her job may be detail enough, but we may need to distinguish between resignation and other reasons for termination of contract. This will depend on the analysis of the decisions.

²² Note that this positivist treatment of precedential constraint in particular does not entail any view of whether the law more generally should, or should not, be understood in positivist terms.

5 Binding rules

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. 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, or more preferred, rules supporting the opposite side.

When is one of two conflicting precedent rules preferable to the other? 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 themselves can therefore be placed in a preference ranking exactly in accord with the reasons they provide, so that a rule r' is ranked as preferable to the rule r in the context of a case base Γ whenever the reason $Premise(r')$ is ranked as preferable to the reason $Premise(r)$, according to the preference relation $<_{\Gamma}$ derived from that case base—whenever, that is, $Premise(r) <_{\Gamma} Premise(r')$.

Given this preference ranking among rules, we can now characterize an applicable rule as trumped whenever there is another rule, also applicable, that is preferred to it 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 if and only if 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 quite simply, 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 if and only if it is applicable in the fact situation X and not trumped in the context of Γ .

These concepts can be illustrated by considering the very simple case base $\Gamma_1 = \{c_1, c_2\}$, containing the familiar case $c_1 = \langle X_1, r_1, s_1 \rangle$ —where, once again, $X_1 = \{f_1^{\pi}, f_2^{\pi}, f_3^{\pi}, f_1^{\delta}, f_2^{\delta}, f_3^{\delta}, f_4^{\delta}\}$, where r_1 is $\{f_1^{\pi}, f_2^{\pi}\} \rightarrow \pi$, and where s_1 is π —as well as the new case $c_2 = \langle X_2, r_2, s_2 \rangle$, where $X_2 = \{f_1^{\pi}, f_2^{\pi}, f_4^{\pi}, f_5^{\delta}, f_6^{\delta}\}$, where r_2 is $\{f_5^{\delta}\} \rightarrow \delta$, and where 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 precedent rules derived from Γ_1 contains r_1 and r_2 ; and evidently, both rules are applicable in the new 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 preferable to 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, and the defendant wins, with r_2 as the rule of this case. Note then, that this case adds nothing new to the theory, although it could be cited along with c_2 in future cases where this rule was applicable to provide additional support for the rule.

6 Constraint by precedent

The account of precedential constraint set out here is a version of the reason model, according to which courts are constrained to reach a decision that is consistent, not necessarily with the rules set out in earlier cases, but with the assessments reached in those cases concerning the proper balance of reasons. In order to develop this idea, we 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 reason Y is preferred to the reason 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 reason Y itself is preferred to X , or else that some reason at least as weak as Y is preferred to some reason 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 reasons X and Y for which both $X <_{\Gamma} Y$ and $Y <_{\Gamma} X$ —whenever, that is, Γ tells us both that Y is preferred to X and that X is preferred to 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 if and only if there are reasons X and Y such that $X <_{\Gamma} Y$ and $Y <_{\Gamma} X$. Γ is consistent if and only if it is not inconsistent.

Is this a good definition of case base inconsistency, and so consistency, from an intuitive point of view? We 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 reasons, each is preferred to 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 that the formal condition set out in our definition of inconsistency has been met, 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.

Other anomalies can likewise be shown to entail our formal condition as well. Indeed, we 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 $X_1 = \{f_1^\pi, f_2^\pi, f_5^\pi, f_5^\delta, f_7^\delta\}$. 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.²³ Still, on the view developed here, precedential constraint does not depend on binding rules, but instead, on consistency with the background case base.

The court, in this situation, would of course 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, of course, preserve consistency. But the court is also free to decide, for example, that the new reason $\{f_5^\pi\}$, which favors the plaintiff and did not appear in c_1 , is itself preferable to the various reasons 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

²³ See, for example, the discussion of “presumptive positivism” in Schauer (1989, 1991).

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 where s_4 is π . 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 $\{f_5^\delta, f_7^\delta\} <_{c_4} \{f_5^\pi\}$. But this new preference relation is consistent with the others already derivable from the background case base.²⁴

What the court cannot do, however—what precedential constraint rules out—is find for the plaintiff on the basis of some reason that does not include the new factor f_5^π , since it is a consequence of the c_2 court's decision that any such reason is less preferable than the reason $\{f_5^\delta\}$ for the defendant. 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 reason $\{f_1^\pi\}$ for the plaintiff outweighs the reason $\{f_5^\delta\}$ 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.²⁵

Having defined the core concept of precedential constraint, we should now to discuss two related issues, one briefly and one at more length.

First of all, it is worth noting that our core account of 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.²⁶ 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

²⁴ Had the court not considered f_5^π to be by itself stronger than f_5^δ , then other pro-plaintiff factors could have been included in r_4 to allow a decision for the plaintiff: for example the antecedent could have been $\{f_1^\pi, f_2^\pi, f_5^\pi\}$, which would be a way of showing explicitly that the plaintiff's position was stronger than in c_1 , and would present weaker constraints on future decisions.

²⁵ Although this account of precedential constraint is related to the reason model originally proposed by Lamond (2005), there are at least two significant differences between Lamond's proposal and the present account. First, 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: it is reasons that function as the premises of precedent rules, and the preference ordering on rules is derived from the preference ordering on reasons. And second, although Lamond himself does not develop a precise implementation of his reason model, the account set out here does seem to differ from what he has in mind. In particular, Lamond suggests (pp. 18–19) that the meaning of precedents should be understood in terms of “protected reasons” and other forms of exclusionary reasons. Although there may be merit in this idea, we did not, in the end, find it necessary to appeal to exclusionary reasons in developing the present account.

²⁶ Chorley and Bench-Capon (2005), for example, demonstrates the rarity of finding a theory capable of explaining all the cases in a substantial case base. The difficulty may be further compounded by the analysis requiring some degree of subjectivity in the assignment of factors.

pervasive. A case base might be inconsistent in certain areas, providing conflicting information about the relative weight of particular reasons, 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.

Second, we must now return to the vexed issues surrounding transitivity of the preference relation derived from an entire case base. As noted earlier, the relation $<_{\Gamma}$, introduced to represent the preferences among reasons derived from the case base Γ , is not transitive: $X <_{\Gamma} Y$ and $Y <_{\Gamma} Z$ do 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 reasons 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.

What blocks transitivity, then, is the assumption—built into our definition—that two reasons can be related by the $<_{\Gamma}$ relation only if they favor opposite sides. In fact, this assumption is not unnatural. The $<_{\Gamma}$ relation is built on top of the $<_c$ relation, representing the preferences among reasons derived from the single case c , and what the court decides in any single case is whether, subject to the constraints of precedent, the reasons presented for one side are or are not stronger than the reasons presented for another; any observation that a reason for one side happens to be stronger than another reason for that same side would likely be taken as mere *dicta*, and not authoritative in future decisions.

But even if a strength comparison between reasons favoring the same side cannot be derived from a single case, perhaps such a comparison can be derived by combining information from several cases within a case base. Suppose, for example, that our background case base contains the case $c_6 = \langle X_6, r_6, s_6 \rangle$, where $X_6 = \{f_1^{\pi}, f_1^{\delta}\}$, where r_6 is the rule $\{f_1^{\pi}\} \rightarrow \pi$, and where s_6 is π , as well as the case $c_7 = \langle X_7, r_7, s_7 \rangle$, where $X_7 = \{f_1^{\pi}, f_2^{\delta}\}$, where r_7 is the rule $\{f_2^{\delta}\} \rightarrow \delta$, and where s_7 is δ . From these two cases, we have $\{f_1^{\delta}\} <_{c_6} \{f_1^{\pi}\}$ and $\{f_1^{\pi}\} <_{c_7} \{f_2^{\delta}\}$ —that is, $\{f_1^{\pi}\}$ is a stronger reason for the plaintiff than $\{f_1^{\delta}\}$ is for the defendant, and $\{f_2^{\delta}\}$ is a stronger reason for that defendant than $\{f_1^{\pi}\}$ is for the plaintiff. It is therefore tempting to conclude, through a form of transitivity, that $\{f_2^{\delta}\}$ is itself a stronger reason for the defendant than $\{f_1^{\pi}\}$ is—otherwise, why would $\{f_2^{\delta}\}$ but not $\{f_1^{\pi}\}$ be preferred to $\{f_1^{\pi}\}$?

Now, if we were to embrace this temptation, a new, stronger form of precedential constraint would then be available. Imagine that the case base also contains the case $c_8 = \langle X_8, r_8, s_8 \rangle$, where $X_8 = \{f_2^{\delta}, f_1^{\pi}\}$, where r_8 is the rule $\{f_1^{\pi}\} \rightarrow \delta$, and where s_8 is δ ; and suppose the court is currently confronting the fact situation $X_9 = \{f_2^{\delta}, f_2^{\pi}\}$. An advocate for the defendant could then argue as follows:

It was concluded by the c_6 , court that $\{f_1^{\pi}\}$ is stronger for the plaintiff than $\{f_1^{\delta}\}$ is for the defendant, and by the c_7 , court that $\{f_2^{\delta}\}$ is stronger for the defendant than $\{f_1^{\pi}\}$ is for the plaintiff. Therefore, combining these results by transitivity, it follows that $\{f_2^{\delta}\}$ must stronger for the defendant than $\{f_1^{\pi}\}$. Now, it was also

concluded by the c_8 , court that $\{f_1^\delta\}$ itself is sufficient to guarantee a decision for the defendant over $\{f_2^\delta\}$ for the plaintiff. Therefore, since we have already seen that $\{f_2^\delta\}$ is stronger for the defendant than $\{f_1^\delta\}$, it follows that X_9 should be decided for the defendant as well.

This argument is tempting, and there is no technical difficulty in extending our definitions to support the notion of constraint it suggests. To do so, we need only move from the familiar relation $<_\Gamma$, introduced in Definition 3 to represent the intransitive preferences derived from the case base Γ , to a stronger relation—say, \prec_Γ —defined simply as the transitive closure of the previous intransitive relation.²⁷ More exactly, the relation \prec_Γ , representing the *transitive preferences* derived from Γ , can be defined by stipulating that, where W and Z are reasons, then $W \prec_\Gamma Z$ if and only if there is a sequence of reasons X_1, X_2, \dots, X_n such that (i) $X_1 = W$ and $X_n = Z$ and (ii) $X_i <_\Gamma X_{i+1}$ for i from 1 through $n - 1$. Using this new idea of transitive preference, we could then mirror our Definition 7 and 8 ideas of consistency and constraint to arrive at their transitive analogues by stipulating: first, that the case base Γ possesses the property of *transitive consistency* if and only if there are no reasons W and Z such that $W \prec_\Gamma Z$ and $Z \prec_\Gamma W$; and second, that *transitive precedential constraint* requires a court confronting a new fact situation X against the background of a transitive consistent case base Γ to reach a decision based on a rule r leading to an outcome s such that $\Gamma \cup \{ \langle X, r, s \rangle \}$ preserves transitive consistency.

The resulting transitive theory of precedential constraint would then allow us to validate the advocate's argument for the defendant in the situation $X_9 = \{f_2^\delta, f_2^\delta\}$. The case c_8 tells us that $\{f_2^\pi\} <_{c_8} \{f_1^\delta\}$, and, as we have already seen, c_6 and c_7 establish that $\{f_1^\delta\} <_{c_6} \{f_1^\pi\}$ and $\{f_1^\pi\} <_{c_7} \{f_2^\delta\}$. If we take Γ as the background case base containing each of these cases, we therefore know that $\{f_2^\pi\} <_\Gamma \{f_1^\delta\}$, that $\{f_1^\delta\} <_\Gamma \{f_1^\pi\}$, and that $\{f_1^\pi\} <_\Gamma \{f_2^\delta\}$. From this, our new definition of transitive preference allows us to conclude that $\{f_2^\pi\} \prec_\Gamma \{f_2^\delta\}$. Transitive precedential constraint thus forces a decision in X_9 for the defendant, since a decision for the plaintiff, would then establish that $\{f_2^\delta\} \prec_\Gamma \{f_2^\pi\}$ as well, and so lead to a transitive inconsistency in the case base.

Still, even though this kind of argument is tempting, and even though the present account can be extended in a straightforward way to support the transitive reasoning necessary to validate the argument, we are not entirely convinced that we should allow this extension. The concerns have to do with transitivity itself, and in particular, with the way in which transitivity allows the preference relations among different reasons established by different courts to be amalgamated into a sort of group preference, even though the various reasons involved may never have been considered together by any single court—as in our example, where the separate judgments of the c_6 , c_7 , and c_8 courts are combined to support the overall judgment that $\{f_2^\delta\}$ outweighs $\{f_2^\pi\}$, even though no case presenting both of these reasons together has yet been considered.

²⁷ It is this stronger preference relation that is given by the reasoning in Bench-Capon (1999). The use of values, as in Bench-Capon and Sartor (2003), provides a relation of intermediate strength, where transitive reasoning is permitted for factors that promote the same value, but not for factors promoting different values.

The appeal to transitive reasoning introduces a number of complex issues concerning the amalgamation of judgments and preferences from different sources.²⁸ In order to avoid these additional complexities, except in so far as the precedent mechanism can itself be seen as expressing group preferences, we concentrate in this paper only on the core account of precedential constraint set out in Definitions 3, 7, and 8, leaving the promises and problems associated with any possible transitive extension of this core account for another time.

7 Case base dynamics

We now want to consider how two of the most important concepts from the traditional theory of precedent, the concepts of following and of distinguishing a precedent, can be accommodated within the present framework. The more general 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 impression, presenting—in the paradigm case—issues that have not previously been addressed within the law, so that the established rules of precedent have no bearing.²⁹ The court is therefore free to assess the issues in whatever way it thinks best, to formulate a rule r supporting an 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 any such decision is guaranteed to preserve consistency. In such a case of first impression, there is no precedential constraint at all.

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 , let us suppose that some precedent rule from $Rule(\Gamma)$ is, in fact, binding in this situation. Then, according to the traditional theory, any court has the option of either following or distinguishing the binding precedent rule.³⁰

²⁸ Any reader who is unfamiliar with these issues is invited to consult Kornhauser and Sager (1986) and the extensive literature on judgment aggregation spawned by this paper, in addition to the standard work in economics on the derivation of group preferences from individual preferences.

²⁹ See, for example, the discussion in Cross (1968, pp. 190–195); at times, other, more complicated situations are also referred to as cases of first impression, such as situations in which multiple precedent rules apply yet none clearly trumps the other, or situations in which entirely novel factors are present.

³⁰ Of course, depending on their standing in the judicial hierarchy, a court may have the further option of overruling the precedent. We will not consider this further option in this paper, but see, for example, Wyner and Bench-Capon (2009).

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 we have seen, precedential constraint is not, on the present theory, carried through precedent rules, so one might wonder what role these rules do play. Here is one answer. Although satisfying precedential constraint does not require following a binding rule, it turns out that following a binding rule is sufficient to guarantee the satisfaction of precedential constraint; any decision that follows a binding rule 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. 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 guarantee with distinguishing; the operation must be performed with some care, guided by the constraints of precedent.

Both the ideas of following and of distinguishing a binding rule are 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 .

With these notions of following and distinguishing a rule before us, three points can be made.

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 will be applicable in the new situation, or some binding rule will be applicable.³¹ 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 simply by following that rule.

³¹ We 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.

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 $X_3 = \{f_1^\sigma, f_2^\sigma, f_5^\sigma, f_5^\delta, f_7^\delta\}$ 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 $X_2 = \{f_1^\sigma, f_2^\sigma, f_4^\sigma, f_5^\delta, f_6^\delta\}$ 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 reason $\{f_1^\sigma, f_2^\sigma, f_4^\sigma\}$ present in X_2 for the plaintiff.

Now, we have previously imagined that the current court, confronted with X_3 , might conclude that the new reason $\{f_5^\delta\}$, which was not present in X_2 , is preferred to the previous reason $\{f_5^\delta\}$ for the defendant, and so decide X_3 for the plaintiff. But we could likewise suppose that the current court, while still recognizing that the reason $\{f_5^\delta\}$ for the plaintiff is preferred to $\{f_6^\delta\}$ for the defendant, also feels that the new reason $\{f_7^\delta\}$ for the defendant, also present here for the first time, itself outweighs $\{f_5^\delta\}$. The current court might then reach the decision $c_{10} = \langle X_{10}, r_{10}, s_{10} \rangle$, where X_{10} is X_3 , where s_{10} is δ , and where r_{10} represents the new rule $\{f_7^\delta\} \rightarrow \delta$. In this case, 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 situation presents a reason $\{f_5^\delta\}$ for the plaintiff that outweighs the reason $\{f_5^\delta\}$ 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.”³² 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.”³³ These remarks can be hard to understand, particularly when we see judges themselves frequently claiming to be following rules and tests established in previous cases. It is easy enough to see how distinguishing a precedent rule might introduce a change into the law; as we have seen, distinguishing a rule often involves introducing a new rule into the case base.

³² See Levi (1949, pp. 2–4). This is followed literally by Bench-Capon and Sartor in (2003) and Chorley and Bench-Capon (2005), where the theory construction is started afresh for each new case.

³³ See Simpson (1961, p. 172), who himself cites Levi.

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 current account gives us the resources to answer this question, as long as we take phrases such as “the rules change as the rules are applied” to refer, 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 these constraints depend, not on the set of precedent rules, but on consistency with the background case base. Consistency, in turn, is defined in terms of the preference ordering on reasons 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 reasons in a way that affects precedential constraint. For each time a familiar rule is followed in a new case, the court makes the decision that certain reasons, which might, conceivably, 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 reasons 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 $X_{11} = \{f_5^\pi, f_6^\pi, f_5^\delta, f_7^\delta\}$, again with r_2 as the unique binding rule; but here, it decides simply to apply this rule, resulting in the decision $c_{11} = \langle X_{11}, r_{11}, s_{11} \rangle$, where r_{11} is r_2 and s_{11} is δ . The court would then be confronting the new situation X_3 against the background of the case base $\Gamma_2 = \{c_1, c_2, c_{11}\}$, rather than the case base Γ_1 . Of course, since c_{11} is decided by the application of a rule already present in c_2 , the set of rules derived from these two case bases is identical: $Rule(\Gamma_2)$ is simply $Rule(\Gamma_1)$. Nevertheless, although the rules remain the same, the addition of c_{11} does affect the derived priority ordering: $<_{\Gamma_2}$ is not identical with $<_{\Gamma_1}$. And this new priority ordering then changes precedential constraint in such a way that the court, when faced with the situation X_3 , would no longer be free to reach the decision $c_4 = \langle X_4, r_4, s_4 \rangle$. Why? Well, we can conclude from the court's decision in c_{11} that the reason $\{f_5^\delta\}$ for the defendant is preferred to the reason $\{f_5^\pi\}$ for the plaintiff—that is, $\{f_5^\pi\} <_{c_{11}} \{f_5^\delta\}$, from which it follows that $\{f_5^\pi\} <_{\Gamma_2} \{f_5^\delta\}$. But the decision c_4 would tell us exactly the opposite, that the reason $\{f_5^\pi\}$ for the plaintiff is preferred to the reason $\{f_5^\delta\}$ for the defendant—that is, $\{f_5^\delta\} <_{c_4} \{f_5^\pi\}$, from which it would follow that $\{f_5^\delta\} <_{\Gamma_2} \{f_5^\pi\}$.

The decision c_4 is therefore inconsistent with Γ_2 , or with any other case base containing c_{11} , and so would be ruled out by precedential constraint.³⁴

8 Discussion

The approach that we have described above broadly represents a development from the position of Prakken and Sartor (1998). Like them, we represent cases as sets of factors, and precedent decisions in terms of rules and priorities. Key differences are that our priorities are between sets of factors, or reasons, rather than rules, and that the reason for the winning side that we derive from a precedent need not be the strongest available to that side in that case, enabling us to go beyond a fortiori reasoning. Within this framework, we have been able to give a precise account of what it is to follow, distinguish and overrule precedents, and how case law can develop over time. In particular we have presented a precise model in which adjudicators have freedom to decide on the particular facts of a case while retaining consistency with previous decisions.

In this section, we will first compare the account developed in this paper with approaches based on theory construction, which also can be seen as starting from the work of Prakken and Sartor, but then moving in a rather different direction; the particular version of theory construction we focus on here is that proposed by Bench-Capon and Sartor in (2001), refined in Bench-Capon and Sartor (2003), and then empirically investigated by Alison Chorley in Chorley and Bench-Capon (2005a, b). We will then consider approaches which relax our condition that all rules involve a single step from factors to the side that should win. In that subsection we will consider two approaches: the *entangled factor hierarchy* of Roth and Verheij (2004) and the *rationales* of Loui and Norman (1995).

³⁴ Lamond (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, we 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* (p. 17; see also p. 20).” To illustrate with our 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 $\{f_1^c, f_2^c, f_4^c\}$ 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_{11} would change the doctrine by expanding this list to $\{f_1^c, f_2^c, f_4^c, f_5^c, f_6^c\}$, which includes also the new factors judged by the c_{11} court as insufficient to defeat r_2 . We feel, however, that the use of a simple list of individual factors that are viewed as insufficient to override a rule is not adequate, and instead rely here on a more general preference relation to represent strength comparisons among conflicting reasons—sets of factors, rather than individual factors—and it is this relation that is then updated as rules are applied. On our view, then, the doctrine of c_2 consists in the rule r_2 together with the information that $\{f_5^\delta\}$ is preferred to the reason $\{f_1^c, f_2^c, f_4^c\}$ or any of its subsets, so that none of these reasons can defeat that rule; the case c_{11} then adds the further information that $\{f_5^\delta\}$ is likewise preferred to $\{f_5^c, f_6^c\}$ or its subsets. But all of this is consistent with the idea that a court might eventually decide that $\{f_5^\delta\}$ is not preferred to, say, the reason $\{f_4^c, f_5^c\}$ —which contains a pair of potential defeaters, one from each of 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 refers only to individual factors, rather than reasons, or sets of factors.

8.1 Theory construction

The theory construction approach was also motivated by the desire to go beyond a fortiori reasoning, and so extend the applicability of precedents. Steps to achieve this had already been a key motivation of the introduction of a factor hierarchy in CATO as discussed above. The idea of the factor hierarchy was that if two factors both contributed to the same abstract factor, they could be substituted for one another. Thus a plaintiff who did not take security measures might establish that the information was a trade secret in some other way, and then use a rule derived from a case in which this issue had been decided by the existence of security measures. Similar considerations motivated the use of teleological reasoning, first studied by Berman and Hafner (1993), who argued that when a case presents strengths and weaknesses for both sides, the conflict is often resolved by focusing on the social purposes promoted by findings for different parties; in the much-discussed property law case of *Pierson v Post*, for example, the social purpose of clarity of the law was preferred to a certain economic benefit. The decision can thus be seen as reflecting a preference for one social purpose over another, and this preference can then be projected forward to constrain reasoning in future cases.³⁵ The theory construction approach was an attempt to incorporate teleological considerations of this kind into legal case based reasoning in a systematic fashion.

The key idea of Bench-Capon and Sartor (2003) was that a factor favours a party in virtue of the purpose served by finding for that party when that factor is present, and that the decisions in precedent cases can be explained in terms of a preference ordering on the purposes, or *values*, relevant to the body of case law.³⁶ To illustrate, in CATO there is a pro-plaintiff factor *bribed-employee* and a pro-plaintiff factor *used-deception*. Both these factors favour the plaintiff, because deciding for the plaintiff in such cases discourages improper or illegal behaviour from being used to obtain trade secrets, and so promotes the social value of acting within the law. Now if, for example we have a precedent in which a case with *bribed-employee* is decided for the plaintiff despite the presence of the pro-defendant factor *reverse-engineerable*, indicating that the purported secret was reverse engineerable and promoting the value of healthy competition, we can see this decision as a preference for the value of respect for law over the value of healthy competition. Once this value preference has been established, it can be applied to comparison between different factors which involve the same value conflict, so that *used-deception* will also be preferred to *reverse-engineerable*. This can achieve the same results as the factor hierarchy, allowing factors with the same value to be substituted for one another, and permits additional inferences based on the transitivity of the value preference relation.

Case based reasoning is then conceived in Bench-Capon and Sartor (2003) as a process of *theory construction*. The process starts from a background comprising a set of factors, each associated with a value, along with a set of cases, each

³⁵ Note that it is possible for difference preferences to be expressed in different jurisdictions or at different times.

³⁶ The terminology of *values*, in this context, was derived from Bench-Capon (2003).

represented only as a set of factors together with an outcome. Note that these cases are like those considered in this paper, except that the cases from Bench-Capon and Sartor (2003) do not also contain rules; instead, rules are developed as part of the process of theory construction and associated with the theory as a whole, rather than with a particular case. The idea is to use the cases to construct a *theory*—a set of rules, priorities between rules and a set of preferences between values—that will explain the outcomes in a selected subset of the cases. This process begins with an empty theory, and is conceived of as adversarial; the two sides take turns to modify the theory using one of a set of theory construction operators. The main operators are: include factor, include case, merge factors, preference from case, preference from value, add rule preference and add value preference.³⁷ The idea is that each party should use their turn to refine the theory so that it gives the outcome they want for the current case, and improves on the previous theory. When a side is unable to produce an improved theory which changes the result, the current theory is adopted and applied to the current case.

Theories are assessed so that the more cases explained, the fewer the rules, and the fewer the unjustified preferences required, the better the theory. Very simple rules, according to which a factor favoring a side actually favors that side, are included automatically. More complex rules are then constructed from these simple rules by combining antecedents, using the merge factors constructor, so as to produce a set of rules capable of producing the decisions in the entire body of cases. Thus while Prakken and Sartor (1998) work with the strongest rule available from each case, and the approach described in this paper uses a specific rule associated with a case, typically containing a subset of the available factors, Bench-Capon and Sartor (2003) constructs the weakest rules capable of explaining the set of cases included in the case base. New cases are decided by constructing the best possible theory that explains a case base including that new case as well, using the adversarial process described above, and seeing what outcome for the new case follows from that theory. As noted above, each new case is considered from first principles, starting from an empty theory, and so there is no notion of building on or modifying a set of rules expressing the current understanding of a law relating to some domain. While this is consequently well able to deal with cases which require a radical reinterpretation, it is not a good reflection of actual legal practice.

In the original work of Bench-Capon and Sartor, theories were constructed and evaluated by hand. Later work by Chorley, in Chorley and Bench-Capon (2005a, b), explored the possibilities for automating the theory construction process, which permitted substantially larger collections of cases to be considered. The idea was to use search to select from the theory construction operators mentioned above, in the context of the case to be decided, so as to produce the best available theory according to a particular set of evaluation heuristics, and again, to use that theory to determine the outcome of the case in question. Searches were implemented using exhaustive search, heuristic search (A^*), and alpha-beta pruning on an adversarial dialogue game between plaintiff and defendant. Experiments using a range of

³⁷ The last two need to be used with caution since the preferences they add cannot be justified from the existing case base, and so tend to weaken the theory.

treatments of factors and values (uniform weights, weights according to value, differential factor weights within values) were conducted. Results were encouraging, attaining accuracy of between 85 and 92 %, which is as good as was obtained in the best competitor system for classifying legal cases based on precedents.³⁸

We see the approach presented in the current paper and the theory construction approach developed by Bench-Capon, Sartor, and Chorley as essentially complementary, but directed at rather different aims. The approach presented here begins with a set F of factors, and a set Γ of cases; from these we can derive the set $Rule(\Gamma)$ of rules contained within these cases, as well as the relation $<_{\Gamma}$ expressing preferences among the reasons, or sets of factors sets of factors, that compose the antecedents of the rules belonging to $Rule(\Gamma)$. The theory construction approach also starts with the set F of factors along with a set Γ of cases, like those studied here but from which rules have been omitted. The theory construction process then produces a theory T comprising: a set $F' \subseteq F$ of factors, a set $\Gamma' \subseteq \Gamma$ of cases, a set R of rules, a relation $<_T$ of preferences on these rules, a a further relation $<_V$ of preferences among the underlying values. Importantly, the rules studied here and those produced by the theory construction approach are identical in format: both use a set of factors in their antecedents and an outcome in their consequents. Thus in both the present approach and the theory construction account we move from factors to outcome in a single step, without any chains of inference, since we have no rules capable of being chained, the antecedents and consequents being from different domains. This makes the two approaches compatible in this respect.³⁹

Now suppose that we first construct a theory, in which we use the include factor constructor to ensure that $F' = F$ and the include case constructor to ensure that $\Gamma' = \Gamma$. The best theory T will have a set of rules R and preferences between them $<_T$. We can then associate the rules of R with the cases from Γ so that these cases, like those considered in this paper, will also contain rules; it is also possible to assign these rules to cases in such a way that the preference ordering $<_{\Gamma}$ defined in this paper is such that the preference relation $<_T$ on rules can be derived from it. At this point, we will be able to apply the model of following, distinguishing and overruling precedents developed in this paper to a theory constructed using the original techniques of Bench-Capon and Sartor (2003), or even one of the automatically generated theories of Chorley and Bench-Capon (2005). What this buys us, in addition, is the value preference $<_V$ to explain the rules and the preferences between them, whereas in the account developed here, there was no specification of how the association of particular rules and cases was to be justified.

³⁸ See Ashley and Brüninghaus's IBP (2003). IBP was developed from CATO and represents a third direction of development: it broke down the decision in trade secret cases into a number of issues based on the Restatement of Torts, used the factors relevant to those issues to decide the particular issue at hand, and then combined these results.

³⁹ Compare this with the model in Lindhal and Odelstat (2006), where the authors, following Ross (1957), view legal reasoning as a two step process: first from facts to intermediate predicates, and then from intermediate predicates to legal conclusions. Effectively, the approach studied in this paper, along with the theory construction approach, performs only the second of these steps; as noted by Ashley and Brüninghaus (2003), factors function as intermediate legal predicates, so that the first step is already carried out once a case has been analyzed in terms of factors. We will look further at work containing chains of reasoning below.

Thus, seen in this way, theory construction justifies the theory of case law, and our present account justifies the way in which that theory is applied as new cases present themselves.

In the well known theory of Levi (1949), legal development follows a three stage process. An initial period of fluctuation is followed by a period of stability; during such a period—which typically lasts for a reasonable period of time—there is normally a consensus as to what the rules associated with particular cases are, and to work from a *tabula rasa* on each occasion would not be efficient, and would not promote consistency and predictability. Eventually, though, tensions will develop and this stability will break down, typically through a landmark case, whereupon a new interpretation is developed and the existing cases are subject to reinterpretation. Usually, breakdown occurs in response to significant social or technological shifts, such as changes in the role of women in society or the invention of motor cars, which make traditional understanding a fact situation no longer appropriate.⁴⁰

In relation to Levi's model, the theory construction approach is appropriate in the first and third of these stages, involving fluctuation and reinterpretation, while the account developed in the current paper is applicable in the second stage, the period of relative calm. This second stage will typically last for some time, but after a breakdown the analysis will need revision, and precedents will become unreliable. Thus, the approach developed here remedies some deficiencies in the theory construction approach—most notably, that the idea of developing a fresh interpretation with each new case, reflecting the absence of any fixed understanding of the law, does not seem descriptively plausible; the dialogues generated by Chorley and Bench-Capon (2005) are quite unlike those we find in actual legal discourse. In turn, however, the theory construction approach can more closely model the concerns at work in periods of broad legal reinterpretation, and can likewise provide a justification for the set of rules that is assigned to cases, which is lacking on the present account.⁴¹

8.2 Chains of reasoning

Throughout this paper our rules have had a form in which the antecedents are a set of factors and the conclusion is a legal consequence of the presence of these factors in a case, namely *decide for the plaintiff* or *decide for the defendant*. Thus all inference comprises a single step: rules cannot be chained since the conclusions are of a different type from the premises. Earlier we noted the approach of Lindahl and Odelstad (2006) which, following Ross (1957), sees legal inference as a two step process: from facts to intermediate predicates and from intermediate predicates to legal consequences. Since factors are ascribed to cases on the basis of patterns of facts, they correspond closely to the intermediate predicates of Lindahl and Odelstad (2006), and so our account is consistent with this: the first step is done by

⁴⁰ The appearance of “red flag” cases, as discussed by Berman and Hafner (1995), or “black swans” and “gray cygnets,” as discussed by Rissland and Xu (2011), indicates that we are approaching a period of breakdown: hence the interest in spotting such cases.

⁴¹ There are, of course, other ways of justifying these rules, such as an analysis of opinions in the cases.

the analyst when cases are represented as sets of factors, and in this paper we are concerned only with the second step. The correspondence of factors and intermediate predicates is also noted by Ashley and Brüninghaus (2003).

There are, however, other approaches using factors in which chains of inference are developed. We will first discuss the account of Roth and Verheij (2004) which was based on the doctoral thesis of Bram Roth. In that work, which addresses domain of unfair dismissal law in the Netherlands, we find dialectical support represented as trees of factors with issues at the root. In an example taken from Roth and Verheij (2004): the issue is *Dismissal-Can-Be-Voided* which has a supporting child, *Always-Behaved-Good-Employee*, and an attacking child, *Serious-Act-Of-Violence*. The outcome of a precedent may show that here *Always-Behaved-Good-Employee* is preferred to *Serious-Act-Of-Violence*. If we stopped here, the approach would correspond to that of CATO and this paper. Roth and Verheij, however, continue down the tree: thus the presence of *Always-Behaved-Good-Employee* is seen as the result of weighing *Always-Arrived-On-Time* against *Once-Insulted-Superior*. If we stop here, we have the picture of Lindahl and Odelstad (2006) with *Always-Behaved-Good-Employee* an intermediate factor opposed by facts such as *Once-Insulted-Superior*. But Roth and Verheij (2004) does not need to stop here either: in principle they could produce further factors pro and con the presence of *Always-Arrived-On-Time*. The number of levels is given by what are stipulated to be the *base level* factors in a *particular* comparison. These are the factors which are taken as given and not further analysed. Thus in our approach all factors are considered base level factors. Note that the factor hierarchy is used quite differently from the way it is used in CATO: in CATO it is used to determine which factors may substitute for or cancel one another (an idea introduced in Alven and Ashley 1997 and recently discussed in Wyner et al. 2011): in Roth and Verheij (2004) they are a series of intermediate steps in deciding an issue. These intermediate steps are much more akin to the use of factors in IBP (Brüninghaus and Ashley 2003). In IBP there is a top level and-or tree of issues and sub-issues. The leaves of IBP's issue tree are resolved using CATO's factors, and these results then provide the basis to resolve the tree using standard propositional logic. Two points should be made about Roth and Verheij (2004): first that the reasoning is purely a fortiori: the issue can only be concluded if the precedent is not distinguishable from the current case. Second that the base level factors do tend to shade into facts rather than into the stereotypical patterns of facts that factors are supposed to be. Consider *Once-Insulted-Superior* in their example: this looks very much more like a fact than a factor.

We may, however, ask why they want to consider a chain of reasoning rather than a single step: what is the gain in seeing the intermediate steps, since a rule moving directly from the base level factors to the issue would be equivalent from a logical point of view? The answer is that Roth and Verheij are not simply interested in finding the outcome, but are at least as interested in how that outcome was established, whether that outcome was justified and how that outcome might be disputed. This motivation is clear from the second interesting innovation of Roth and Verheij (2004): not only do they have a factor hierarchy, but it is an *entangled factor hierarchy*. By *entangled* is meant that not only do lower level factors support

and attack factors at the next higher level, but the attack and support relations may themselves be attacked and supported⁴². Thus that a *Serious-Act-Of-Violence* supports *Pressing-Ground-For-Dismissal* is itself supported by *Act-Directed-Against-Employer* and attacked by *Acted-In-Selfdefence*. In order for these arguments about whether support and attack relationships hold to be considered, the steps in deriving the issue from the base level factor need to be articulated and so exposed to attack. In contrast, our aim is to provide a logical mechanism rather than a detailed representation of the argumentation in particular cases. Since our interest is in summarising the conclusions and applying them to subsequent cases, we can afford to conceal the inner workings.

The point may also be made in considering the rationales of Loui et al. (1993) and Loui and Norman (1995). Those papers identify several kinds of rationale, but most relevant for the present purposes is their *compression* rationale, or *c-rationale*. A compression rationale justifies the adoption of a rule because it compresses a line of reasoning. That is, we may use compression to summarise the import of a perhaps lengthy dispute. One example would be to move directly from base level factors to issues in Roth and Verheij (2004), thus simplifying a perhaps highly entangled and complicated tree of supports and attacks by recording only the outcome. Another example would be the rules of Bench-Capon and Sartor in (2003). There *every* rule is associated with a value and so has its rationale in the following line of reasoning: *factor promotes value*, and *value promotion supports outcome* which is compressed to the rule *factor supports outcome*. The purpose of uncompressing the rule to recover the rationale is to provide additional ways to attack the rule. This is stated in Loui Ronald and Norman (1995):

To attack an argument using a rule with a c-rationale, restate the argument in an uncompressed form. The resulting uncompressed argument will be more susceptible to attack.

Thus in Bench-Capon and Sartor (2003) we have a rule in the US Trade Secrets domain that *Reverse-Engineered supports defendant*, with the rationale that *Reverse-Engineered promotes Enterprise*. The idea is that we want to encourage enterprise, and if some one uses a legitimate practice, such as reverse engineering, to come to understand an idea discovered by another person, they should be allowed to use that understanding in their own business. Without the rationale it is difficult to see how one can attack the rule: to argue for the plaintiff one would need to argue either that the factor was not present, or to provide stronger reasons for deciding for the plaintiff. With the rationale, however, one could also argue that reverse engineering did not promote enterprise, or that enterprise was not something the law should be encouraging. While in this particular example these arguments are unlikely to succeed, recovering the rationale allows them to made. In other, more controversial, cases such arguments might be persuasive. For example in *Pierson v Post* as analysed in Berman and Hafner (1993), whether encouraging fox hunting

⁴² This idea interestingly anticipates Extended Argumentation Frameworks (Modgil 2009), which was applied to legal case based reasoning in Bench-Capon and Modgil (2009).

does indeed promote economic benefits, and whether the promotion of economic benefits is the business of the court are both debatable.

If we wish to represent and critique arguments, allowing for the intermediate steps in the reasoning to be made explicit and hence open to criticism, is essential, and chaining is required. From a purely logical standpoint, however, where we are concerned solely to apply rather than justify the rules, the compressed version is enough. Since this is indeed our purpose in this paper, we are therefore able to use single step rules. But this is in no way to deny the importance of the dialectic aspects of case based reasoning, and for this rules capable of being chained will be needed. Note that this applies both to the ways we can move from facts to factors, so that the analysis can be disputed, and to the rationales that take us from factors to legal conclusions.

9 Conclusion

The goal of this paper has been to provide a precise reason model of precedent, based on the general idea that courts are constrained to respect the relative importance assigned to various reasons given in earlier decisions by relevant courts. The account provided here is set out within a representational framework, developed over two decades of AI and Law as described in Sect. 2, that has shown itself to be useful in both analytic and empirical studies of legal precedent. It has the additional advantage of allowing us to see how this reason model can be reconciled with the traditional idea that precedential constraint involves rules, as long as these rules are taken to be defeasible.

We would like to think that the present account could be developed into a more general theory of precedential reasoning, but that would require a good deal of further work along both defensive and constructive lines. Defensively, the present treatment of precedent rules—balanced, as it is, in the center—is vulnerable to attack from both the left and the right. On the left, there are those who feel that rules play no real role in precedential constraint, including writers in the tradition of legal realism, of course, but also those following the original spirit of the factor based approach, where very successful empirical and educational work on precedent involved no appeal to rules whatsoever; this can be seen most fully in the work of Aleven (1997). And on the right, there are those, such as Alexander and Sherwin, who recognize the importance of precedent rules, but argue that coherence requires these rules to be even stronger than those at work here—strict, rather than defeasible. None the less our intermediate position is one which has emerged from the work in AI and Law discussed above and is quite consistent with, indeed complementary to, other approaches in this field, such as that of Bench-Capon and Sartor (2003). From a constructive standpoint, the present account, like other factor based accounts, relies on a number of simplifications, which could fruitfully be relaxed in a more general theory. These include, most notably, the various issues involved in the initial choice to work with factors rather than with dimensions, as discussed earlier, as well as the idea that legal reasoning involves only a single step, connecting base-level factors directly to an outcome, rather than moving through a

series of intermediate concepts. Each of these limitations put some kinds of reasoning with cases outside the scope of the approach described here. The second limitation relates, most directly, to the initial ascription of factors to cases. While we start with a set of factors, in practice the fact patterns that are relevant, and that are thought to form these factors, must be identified, and then each case must be considered to see whether its particular facts satisfy the patterns. Either of these activities could form the basis of legal argument.

While there is some promising work on these issues in AI and Law, there is currently no firm and well understood basis comparable to the account of cases in terms of factors on which we can build. These issues, along with the application of the approach to substantial examples, must therefore be left for future work: in this paper, however, we have, we believe, given a precise account, for some clearly defined, and clearly important, aspects of reasoning with legal cases.

Appendix: 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. \square

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 reasons 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. \square

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 reasons 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. \square

References

- Aleven V (1997) Teaching case-based argumentation through a model and examples. PhD thesis, Intelligent Systems Program, University of Pittsburgh
- Aleven V, Ashley KD (1997) Evaluating a learning environment for case-based argumentation skills. In: Proceedings of the sixth international conference on artificial intelligence and law (ICAIL-97). The Association for Computing Machinery Press, pp 170–179
- Alexander L (1989) Constrained by precedent. *South Calif Law Rev* 63:1–64
- Alexander L, Sherwin E (2001) The rule of rules: morality, rules, and the dilemmas of law. Duke University Press, Durham
- Ashley KD (1990) Modeling legal argument: reasoning with cases and hypotheticals. The MIT Press, Cambridge
- Ashley KD, Brüninghaus S (2003) A predictive role for intermediate legal concepts. In: Bourcier D (ed) JURIX 2003: the sixteenth annual conference on legal knowledge and information systems. IOS Press, pp 153–62
- Ashley KD, Brüninghaus S (2009) Automatically classifying case texts and predicting outcomes. *Artif Intell Law* 17(2):125–165
- Bench-Capon TJM, Sartor G (2001) Theory based explanation of case law domains. In: Proceedings of the 8th international conference on artificial intelligence and law. ACM, pp 12–21
- Bench-Capon TJM (2003) Persuasion in practical argument using value-based argumentation frameworks. *J Logic Comput* 13(3):429–448
- Bench-Capon TJM (1991) Practical legal expert systems: the relation between a formalisation of law and expert knowledge. In: Bennun M, Narayanan A (eds) Computers, law and AI. Ablex Publishers, Norwood, pp 191–201
- Bench-Capon TJM (1999) Some observations on modelling case based reasoning with formal argument models. In: Proceedings of the 7th international conference on AI and Law, pp 36–42

- Bench-Capon TJM, Modgil S (2009) Case law in extended argumentation frameworks. In: Proceedings of the 12th international conference on AI and Law, pp 118–127
- Bench-Capon TJM, Rissland EL (2001) Back to the future: dimensions revisited. In: Verheij B, Lodder A, Loui RP, Muntjewerff A (eds) JURIX 2001: the fourteenth annual conference on legal knowledge and information systems. IOS Press, Amsterdam, pp 41–52
- Bench-Capon TJM, Sartor G (2003) A model of legal reasoning with cases incorporating theories and values. *Artif Intell* 150(1–2):97–143
- Berman DH, Hafner CL (1993) Representing teleological structure in case-based legal reasoning: the missing link. In: Proceedings of fourth international conference on AI and Law, pp 50–59
- Berman DH, Hafner CL (1995) Understanding precedents in a temporal context of evolving legal doctrine. In: Proceedings of fifth international conference on AI and Law, pp 42–51
- Brüninghaus S, Ashley KD (2003) Predicting outcomes of case-based legal arguments. In: Proceedings of the ninth international conference on AI and Law, pp 233–242
- Brüninghaus S, Ashley KD (2005) Reasoning with textual cases. In: Muñoz-Avila H, Ricci F (eds) ICCBR, Lecture notes in computer science, vol 3620. Springer, Berlin, pp 137–151
- Chorley A, Bench-Capon TJM (2005) Agatha: using heuristic search to automate the construction of case law theories. *Artif Intell Law* 13(1):9–51
- Chorley A, Bench-Capon TJM (2005) An empirical investigation of reasoning with legal cases through theory construction and application. *Artif Intell Law* 13(3–4):323–371
- Cross R (1968) *Precedent in English law*, 2nd edn. Oxford University Press, Oxford
- Hage J (1993) Monological reason-based logic: A low level integration of rule-based reasoning and case-based reasoning. In: Proceedings of the fourth international conference on AI and Law, pp 30–39
- Hage J (1997) *Reasoning with rules*. Kluwer, Dordrecht
- Hage J (2005) *Studies in legal logic*. Springer, Dordrecht
- Horty JF (2004) The result model of precedent. *Leg Theory* 10:19–31
- Horty JF (2007) Reasons as defaults. *Philosopher's Imprint* 7(3):1–28
- Horty JF (2012) *Reasons as defaults*. Oxford University Press, Oxford
- Horty JF (2011) Reasons and precedent. In: Proceedings of the 13th international conference on artificial intelligence and law, pp 41–50
- Horty JF (2011) Rules and reasons in the theory of precedent. *Leg Theory* 17:1–33
- Kornhauser L, Sager L (1986) Unpacking the court. *Yale Law J* 96:82–117
- Lamond G (2005) Do precedents create rules. *Leg Theory* 11:1–26
- Levi E (1949) *An introduction to legal reasoning*. The University of Chicago Press, Chicago
- Lindahl L, Odelstad J (2006) Intermediate concepts in normative systems. In: Goble L, Meyer J-J (eds) *Deontic logic and artificial normative systems*, Lecture notes in computer science, vol 4048. Springer, Berlin, pp 187–200
- Loui RP, Norman J (1995) Rationales and argument moves. *Artif Intell Law* 3(3):159–189
- Loui RP, Norman J, Olson J, Merrill A (1993) A design for reasoning with policies, precedents, and rationales. In: Proceedings of the 4th international conference on AI and Law, pp 202–211
- MacCormick DN, Summers R (1997) *Interpreting precedents: a comparative study*. Dartmouth Publishing, Hanover
- Modgil S (2009) Reasoning about preferences in argumentation frameworks. *Artif Intell* 173(9–10):901–934
- Modgil S, Bench-Capon TJM (2010) Integrating dialectical and accrual modes of argumentation. In: Baroni P, Cerutti F, Giacomini M, Simari GR (eds) *Computational models of argument*. Proceedings of COMMA 2010, *Frontiers in artificial intelligence and applications*, vol 216. IOS Press, Amsterdam, pp. 335–346
- Prakken H (2005) A study of accrual of arguments, with applications to evidential reasoning. In: Proceedings of the tenth international conference on artificial intelligence and law, pp 85–94
- Prakken H, Sartor G (1998) Modelling reasoning with precedents in a formal dialogue game. *Artif Intell Law* 6:231–287
- Raz J (1975) *Practical reasoning and norms*. Hutchinson and Company (second edition with new Postscript printed in 1990 by Princeton University Press, and reprinted by Oxford University Press in 2002; pagination refers to the Oxford edition)
- Raz J (1979) *The authority of law*. Oxford University Press, Oxford
- Rissland EL, Ashley KD (1987) A case-based system for trade secrets law. In: Proceedings of the first international conference on AI and Law, pp 60–66

- Rissland EL, Xu X (2011) Catching gray cygnets: an initial exploration. In: Proceedings of the 13th international conference on AI and Law, pp 151–60
- Ross A (1957) Tu-tu. *Harvard Law Rev* 70:812–825
- Roth B, Verheij B (2004) Cases and dialectical arguments—an approach to case-based reasoning. In: Meersman R, Tari Z, Corsaro A (eds) OTM workshops, Lecture notes in computer science, vol 3292. Springer, Berlin, pp 634–651
- Schauer F (1989) Is the common law law. *Calif Law Rev* 77:455–471
- Schauer F (1991) *Playing by the rules: a philosophical examination of rule-based decision-making in law and life*. Oxford University Press, Oxford
- Schild U, Herzog S (1993) The use of meta-rules in rule based legal computer systems. In: Proceedings of the 4th international conference on AI and Law, pp 100–109
- Simpson AWB (1961) The *ratio decidendi* of a case and the doctrine of binding precedent. In: Guest AG (ed) *Oxford essays in jurisprudence*. Oxford University Press, Oxford, pp 148–175
- Wyner AZ, Bench-Capon TJM (2009) Modelling judicial context in argumentation frameworks. *J Logic Comput* 19(6):941–968
- Wyner AZ, Bench-Capon TJM, Atkinson KM (2011) Towards formalising argumentation about legal cases. In: Proceedings of the 13th international conference on artificial intelligence and law, pp 1–10