

Grant Lamond's Account of Precedent:
A Personal Encounter

John Horty

Version of: June 8, 2021

1 Introduction

Ever since reading Aristotle as an undergraduate, I have been intrigued by the particularist idea—which I believe to be correct—that normative generalizations cannot be formulated carefully enough to apply properly in every situation: there will always be circumstances in which the straightforward application of a normative generalization, no matter how finely articulated, will yield the wrong result. How can we live in a world like this? Except by appealing to normative generalizations, how can we justify our own actions, or evaluate the actions of others? Questions like these have recently received sustained attention in moral philosophy due, especially, to the work of Jonathan Dancy, who argues for an extreme version of particularism according to which generalizations have no role at all to play in the moral justification or evaluation of actions.¹

Several years ago, I wrote a book that provided, as I thought, a sensible response to Dancy’s extreme version of particularism, arguing that moral justification and evaluation could be provided by defeasible generalizations, and sketching a logic for reasoning with these generalizations.² Even if this response was successful in deflecting Dancy’s broad-scale assault on normative generalizations, however, it did not address the more detailed problem of explaining how our moral justifications or evaluations could actually be guided, or constrained, by generalizations that might, at any moment, allow for exceptions.

It struck me—as it has many people—that we might begin to understand how normative constraints could be derived even from generalizations that allow for exceptions by looking carefully at the patterns of reasoning at work in systems of common law. Here we find a stable institution in which, on a daily basis, courts confront new situations against a background of constraints derived from precedent decisions. These constraints are sufficiently firm that decisions in most situations are a matter of routine, but also flexible enough that courts have the freedom to respond to fresh circumstances in creative ways. How, in the common law,

¹For a canonical statement of the argument, see Dancy (2004).

²See Horty (2012); previous appeals to defeasibility in response to Dancy’s argument can be found in Lance and Little (2004, 2007).

is this balance between constraint and freedom achieved?

With this question in mind, I embarked on a haphazard study of the literature on common law reasoning—which is, of course, vast and contentious, touching on central problems in legal theory. As I sifted through this literature, I found that it was helpful to bear in mind the perspective adopted by Larry Alexander and Emily Sherwin.³ Although Alexander and Sherwin set out a broader taxonomy of positions—or *models*, as they say—concerning precedential constraint, they are particularly focused on the opposition between two central positions.

Both of these positions are based on the observation that a precedent case normally contains, not only a factual description of some situation as well as a decision on the basis of that description, but some rule through which that decision is justified; and both suppose that the constraints derived from a precedent case depend, in some sense, on this rule. Where the two positions differ is in the exact way this rule leads to constraints. According to the first position—which Alexander and Sherwin advocate, and which they describe as the *rule model*—the rule formulated in a precedent case governs any later situation to which it applies, unless the court in the later case wishes to overrule that earlier decision and has the authority to do so. According to the second position—which Alexander and Sherwin do not explicitly label, but which I refer to as the *standard model*—courts have the power of *distinguishing* the present case from previous cases in which a rule was formulated or applied: the power, that is, of identifying important, or material, differences between the facts of these cases, and so modifying the rule set out earlier in order to avoid an inappropriate application of that rule to the present situation.⁴ The idea that later cases can be distinguishing, in this

³See Alexander (1989) for an initial statement, and then Alexander and Sherwin (2001) and, especially, (2008) for later and more substantial development.

⁴I refer to this model as “standard” because of its prevalence among legal theorists; versions have been advocated by Levi (1949), Simpson (1961), Raz (1979), Eisenberg (1988), and Schauer (1989, 1991), along with many others. Of course, even according to the standard model, rules cannot be modified entirely at will; acceptable modifications must be subject to conditions that I describe as the *Raz/Simpson* conditions, after two prominent advocates, and explore at length in Horty (2015).

sense, can be seen as the common law response to the particularist view that generalizations formulated in earlier cases may not apply properly to new situations.

The contrast between these two positions on precedential constraint, the rule model and the standard model, can be illustrated with a domestic scenario: Suppose that Jack and Jo are the parents of two children—Emma, who has just turned nine, and Max, age twelve—and that they have agreed to respect each other’s decisions concerning the children, treating these decisions, in effect, as binding precedents. And imagine that, one night, Emma, who has completed her homework and finished her chores, but did not finish dinner, asks Jo if she can stay up and watch TV. This is like a legal case: a situation is presented to an authority, Jo, who must make a decision and, ideally, provide a rationale for her decision. Suppose that Jo resolves the case by granting the request, stating that Emma can stay up to watch TV since she is now nine years old. This decision can be seen as introducing a household version of a common law rule—perhaps, “Children age nine or greater can stay up and watch TV”—fashioned in response to a particular set of circumstances, but applicable to future situations as well.

Now imagine that, the next day, Max, who has likewise finished chores, but neither finished dinner nor completed homework, asks Jack whether he can stay up to watch TV. And suppose that Jack views the fact that Max, unlike Emma, has not completed homework as a material difference between the two situations, providing strong support for the conclusion that he should not be allowed to stay up and watch TV. On the other hand, Max is also, like Emma, at least nine years old, so that the rule formulated by Jo in the previous case applies to Max as well. It therefore follows from the rule model that, in spite of his own preference, Jack is nevertheless required to apply the previous rule to the current situation, and conclude that Max can stay up and watch TV. According to the standard model, by contrast, Jack can arrive at his preferred decision by distinguishing the two cases, arguing that the previous rule should not apply to the new case of Max, since this new case, unlike the previous case of Emma, presents the additional feature that the child in question has not completed his homework. An effect of Jack’s decision would then be that the rule set out

by Jo in the previous case of Emma is now modified—perhaps reading, “Children age nine or greater can stay up and watch TV, unless they have failed to complete their homework.”

In their formulation of the rule model, Alexander and Sherwin suppose that a common law rule must behave like a rule from ordinary logic, entailing its consequent, without fail, whenever its premise is satisfied—so that, for example, Jo’s original rule entails that a child must be allowed to watch TV whenever that child is at least nine years old. They therefore miss the possibility that these rules might be thought of as defeasible—so that Jo’s rule could be taken as entailing that a child should, as a default, be allowed to watch TV whenever that child is nine years old, unless there is some particular reason for overriding that default. Logics governing defeasible rules like this had already been applied, very successfully, within the field of artificial intelligence and law, primarily in the analysis of common law argumentation.⁵ I was convinced that a useful treatment of common law constraint could likewise be developed within the framework of these logics, and was in the process of working out the details of a proposal.

It was at this point that I came across Grant Lamond’s “reasons account” of precedent.⁶ Lamond’s paper seemed to contain that rare thing—not a development, an application, a response to criticism, a move in a debate, but a genuinely new idea. The idea, as I see it, is this: common law constraint can be understood as depending, not on rules at all, whether ordinary or defeasible, but on relations among reasons. As soon as I read this paper, I realized that exactly the proposal I was trying to develop using a complex logic of defeasible rules could be stated much more simply without any appeal to rules at all, but directly in terms of a relation among reasons, just as Lamond recommends. The result was my first paper on the topic, setting out what I described as the *reason model* of constraint.⁷

⁵For the starting point in this very large body of research, see Prakken and Sartor (1998).

⁶See Lamond (2005), and then Lamond (2021) for more recent reflections.

⁷See Horty (2011). I use the phrase “reason model,” rather than Lamond’s “reasons account,” primarily to fit the view into Alexander and Sherwin’s taxonomy of models, but it is also useful to have two different phrases to distinguish my account, which I think of as a particular implementation of Lamond’s general idea, from that general idea itself.

Like many important philosophical proposals, Lamond’s idea allows for multiple interpretations. There is my own interpretation, which has been applied within artificial intelligence and law by Trevor Bench-Capon, refined, improved, and compared to arguments from analogy by Adam Rigoni, limited in scope by Gabriel Broughton, and explored from a formal perspective by Henry Prakken.⁸ There is also a new interpretation by Robert Mullins, which I feel is much closer to Lamond’s own views, and which also connects more closely with some of the traditional literature from legal theory.⁹

Rather than consider any of this more recent work, I want to return to my original version of the reason model and try to set it out as simply and clearly as I can. Why do this? Because I have noticed that, in several recent papers, the original account has been misunderstood. This is a matter for which I take responsibility. In my first presentation of the reason model, I had not, in my mind, fully separated the reason model itself from its connections to defeasible logic, with the result that some readers have concluded that rules play an essential role in my formulation of the reason model. In later presentations, I was concerned either with exploring connections between the reason model and other models, or else with developing the reason model in a richer environment, with the result that the presentation of the reason model itself was tangled up with other complexities.¹⁰ There has not yet been a simple, clear statement of the reason model itself, without unnecessary connections to defeasible logics, or to other theories. That is what I hope to provide.

I begin, in the next section, by reviewing the framework within which the reason model is developed. Section 3 then defines the reason model itself. Section 4 shows how the reason model can be seen as an interpretation of Lamond’s proposal, and discusses a few additional points.

⁸See Horty and Bench-Capon (2012), Rigoni (2014, 2015), Broughton (2019), and Prakken (20xx).

⁹See Mullins (2020); interestingly, Mullins also shows that his alternative interpretation is, in a precise sense, equivalent to my own.

¹⁰See Horty (2015, 2016, 2019) for connection between the reason model and, respectively, the standard model and a defeasible rule model, and for development of the model in a richer environment.

2 Basic concepts

2.1 Factors and fact situations

We will suppose that a situation presented to a court for decision can usefully be represented as a set of *factors*, where a factor is a legally significant fact or pattern of facts bearing on that decision. In our domestic scenario, the legal, or quasi-legal, issue at hand is whether a child can stay up and watch TV, and the factors involved might reasonably include those already considered—whether the child has reached the age of nine, completed chores, eaten dinner, finished homework—as well as countless others.

But the factor-based representation of legal situations is not restricted to everyday examples of this kind. In fact, this style of representation has been used to analyze case-based reasoning in a number of complex legal domains within the field of artificial intelligence and law, where it originated in the work of Edwina Rissland and Kevin Ashley.¹¹ Cases in different areas of the law will be characterized by different sets of factors, of course. In the domain of trade-secrets law, for example, where the factor-based analysis has been explored most extensively, a case typically concerns the issue of whether the defendant has gained an unfair competitive advantage over the plaintiff through the misappropriation of a trade secret; and here the factors involved might turn on, say, questions concerning whether the plaintiff took measures to protect the trade secret, whether a confidential relationship existed between the plaintiff and the defendant, whether the information acquired was reverse-engineerable or in some other way publicly available, and the extent to which this information did, in fact, lead to a real competitive advantage for the defendant.¹²

¹¹The analysis of legal cases in terms of factors was first introduced by Rissland and Ashley (1987); see Ashley (1990) for a canonical treatment, Rissland (1990) for an overview of research in artificial intelligence and law that places this work in a broader perspective, and Rissland and Ashley (2002) for later reflections on factor-based representation of legal information.

¹²The most detailed analysis in this domain is presented by Aleven (1997), who analyzed 147 cases from trade-secrets law in terms of a factor hierarchy that includes five high-level issues, eleven intermediate-level concerns, and twenty-six base-level factors. The resulting knowledge base is used in an intelligent tutoring system for teaching elementary skills in legal argumentation, which has achieved results comparable to

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

Formally, then, we start by postulating a set F of legal factors bearing on some particular issue. We will let $F^\pi = \{f_1^\pi, \dots, f_n^\pi\}$ represent the set of factors favoring the plaintiff and $F^\delta = \{f_1^\delta, \dots, f_m^\delta\}$ the set of factors favoring the defendant. Given our assumption that each factor favors one side or the other, the entire set of legal factors will be exhausted by those favoring the plaintiff together with those favoring the defendant: $F = F^\pi \cup F^\delta$. As this notation suggests, we take π and δ to represent the two sides in a dispute, plaintiff and defendant, and where s is one of these sides, we let \bar{s} represent the other: $\bar{\pi} = \delta$ and $\bar{\delta} = \pi$.

Given this collection F of factors, a *fact situation* X , of the sort presented to the court for judgment, can then be defined simply as some particular subset of these factors: $X \subseteq F$. And where X is a fact situation of this kind, we let X^s represent the factors from X that support the side s , so that: $X^\pi = X \cap F^\pi$ and $X^\delta = X \cap F^\delta$. Of course, any interesting situation will contain factors favoring both sides of a given dispute. For example, the situation $X_1 = \{f_1^\pi, f_2^\pi, f_1^\delta, f_2^\delta\}$ contains two factors each favoring the plaintiff and the

traditional methods of instruction in controlled studies; see Aleven and Ashley (1997).

defendant, with those factors favoring the plaintiff contained in $X_1^\pi = \{f_1^\pi, f_2^\pi\}$ and those favoring the defendant contained in $X_1^\delta = \{f_1^\delta, f_2^\delta\}$.

2.2 Reasons and rules

When presented with a fact situation, the court's primary task is to reach a decision, or determine an outcome. Given our assumption that reasoning proceeds in a single step, we can suppose that the *outcome* of a case is a decision either in favor of the plaintiff or in favor of the defendant, with these two outcomes represented as π or δ respectively.

In addition to deciding for one side or the other, we generally expect the court to supply a rule, or principle, to serve as justification for its decision. Rules of this kind will be characterized in terms of reasons, where a *reason for a side* is some set of factors uniformly favoring that side; a *reason* can then be defined as a set of factors uniformly favoring one side or another. To illustrate: $\{f_1^\pi, f_2^\pi\}$ is a reason favoring the plaintiff, and so a reason, while $\{f_1^\delta, f_2^\delta\}$ is a reason favoring the defendant, and so likewise a reason; but the set $\{f_1^\pi, f_1^\delta\}$ is not a reason, since it contains factors favoring opposite sides. Reasons are to be interpreted conjunctively, so that, for example, the reason $\{f_1^\pi, f_2^\pi\}$ represents the conjunction of the propositions represented by the factors f_1^π and f_2^π , and the reason $\{f_1^\pi\}$ carries the same meaning as the factor f_1^π .

Since reasons, like fact situations, are sets of factors, we can stipulate that a reason U *holds* in a situation X just in case each factor from U belongs to X , so that U is a subset of X , or $U \subseteq X$. And we can also define a relation of strength among reasons for a side according to which, where U and V are reasons for the same side, then V *is at least as strong a reason as U for that side* just in case U is a subset of V , or $U \subseteq V$. To illustrate: We can see, first, that the reason $\{f_1^\pi\}$ holds in the previous fact situation $X_1 = \{f_1^\pi, f_2^\pi, f_1^\delta, f_2^\delta\}$, since $\{f_1^\pi\} \subseteq X_1$. And we can see that, of the two reasons $\{f_1^\pi\}$ and $\{f_1^\pi, f_2^\pi\}$, the second favors the plaintiff at least as strongly as the first, since $\{f_1^\pi\} \subseteq \{f_1^\pi, f_2^\pi\}$.

Given this notion of a reason, a *rule* can now be defined as a statement of the form $U \rightarrow s$, where U is a reason supporting the side s . For convenience, we introduce two

auxiliary functions—*Premise* and *Conclusion*—picking out the premise and conclusion of a rule, so that, if r stands for the rule just mentioned, we would have $Premise(r) = U$ and $Conclusion(r) = s$. And we will say that a rule is *applicable* in a situation whenever the reason that forms its premise holds in that situation. To illustrate: The statement $\{f_1^\pi\} \rightarrow \pi$ is a rule, since $\{f_1^\pi\}$ is a reason supporting the plaintiff. If we take r_1 to stand for this rule, we would have $Premise(r_1) = \{f_1^\pi\}$ and $Conclusion(r_1) = \pi$. And r_1 is applicable in the situation X_1 above, since $Premise(r_1)$ holds in this situation.

The rules defined here are to be interpreted as defeasible, telling us that their premises entail their conclusions, not as a matter of necessity, but only by default. Continuing with our illustration: What the rule $r_1 = \{f_1^\pi\} \rightarrow \pi$ means, very roughly, is that, whenever the premise $\{f_1^\pi\}$ of the rule holds in some situation, then, as a default, the court ought to decide that situation for the conclusion π of the rule—or perhaps more simply, that the premise of the rule provides the court with a pro tanto reason for deciding in favor of its conclusion.

2.3 Cases

Given the concepts introduced so far—fact situations, rules, outcomes—a *case* can now be defined as a situation together with an outcome and a rule through which that outcome is justified: such a case can be specified as a triple of the form $c = \langle X, r, s \rangle$, where X is a situation containing the factors presented to the court, r is a rule, and s is an outcome.¹³

To aid our discussion of cases, we introduce three more auxiliary functions—*Facts*, *Rule*, and *Outcome*—mapping cases into their component parts, so that, in the case c above, we would have $Facts(c) = X$, $Rule(c) = r$, and $Outcome(c) = s$. And in order for the concept of a case to make sense, we impose two coherence conditions: first, that the rule of the case is

¹³Our representation of cases embodies the simplifying assumption that the particular rule underlying a court’s decision is plain, ignoring the extensive literature on methods for determining the *ratio decidendi* of a case; and we suppose, as a further simplification, that a case always contains a single rule, ignoring situations in which a court might offer several rules for a decision, or in which a court reaches a decision by majority, with different members of the court offering different rules, or in which a court might simply render a decision in a case without setting out any general rule at all.

applicable to the fact situation of that case, or equivalently, that the reason for the decision must hold in that fact situation, and second, that the conclusion of the case rule must match the outcome of the case itself.

The various concepts and conditions can be illustrated with the case $c_1 = \langle X_1, r_1, s_1 \rangle$, where the fact situation of this case is the familiar $X_1 = \{f_1^\pi, f_2^\pi, f_1^\delta, f_2^\delta\}$, where the case rule is the familiar $r_1 = \{f_1^\pi\} \rightarrow \pi$, and where the outcome of the case is $s_1 = \pi$, a decision for the plaintiff. Evidently, the case satisfies our two coherence conditions. The case rule is applicable to the facts, since $Premise(r_1) \subseteq X_1$ and the conclusion of the case rule matches the case outcome, since $Conclusion(r_1) = \pi$. This particular case, then, represents a situation in which the court, when confronted with the fact situation X_1 , decided for the plaintiff by applying or introducing the rule r_1 , according to which the presence of the factor f_1^π —that is, the reason $\{f_1^\pi\}$ —leads, by default, to a decision for the plaintiff.

Finally, with this notion of a 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—a set of precedent cases—that will be taken to represent the common law in some area, and to constrain the decisions of future courts.

3 Constraint by reasons

How do the cases from a case base constrain future decisions? According to the reason model, what matters about a precedent case is the precedent court's assessment of the relative importance of the reasons presented by that case for each of the opposing sides; this assessment can be represented as a priority ordering on reasons. Later courts are then required to reach decisions that are consistent with the priority ordering derived from the decisions of earlier courts, thereby respecting the assessment of relative importance among reasons determined by these earlier, or precedent, courts.

3.1 A priority ordering on reasons

In order to develop this idea, we need to explain how a priority ordering on reasons can be derived from the decisions of earlier courts, and then what it means for the decision of a

later court to be consistent with that ordering.

To begin with, then, let us return to the case $c_1 = \langle X_1, r_1, s_1 \rangle$ —where $X_1 = \{f_1^\pi, f_2^\pi, f_1^\delta, f_2^\delta\}$, where $r_1 = \{f_1^\pi\} \rightarrow \pi$, and where $s_1 = \pi$ —and ask what information is carried by this case; what is the court telling us with its decision? Well, two things. First of all, with its decision for the plaintiff on the basis of the rule r_1 , the court is registering its judgment that $Premise(r_1)$, the reason for its decision, is more important—or has higher *priority*—than any reason for the defendant that holds in X_1 , the fact situation of the case.¹⁴ How do we know this? Because if the court had viewed some reason for the defendant that held in the situation X_1 as more important, or higher in priority, than $Premise(r_1)$, the court would have found for the defendant on the basis of that reason, rather than for the plaintiff on the basis of $Premise(r_1)$. And second, if the court is telling us explicitly that the reason $Premise(r_1)$ itself has higher priority than any reason for the defendant that holds in X_1 , then the court must also be telling us, at least implicitly, that any other reason for the plaintiff that is at least as strong as $Premise(r_1)$ must likewise have a higher priority than any reason for the defendant that holds in this situation.

We can recall that a reason U for the defendant holds in the situation X_1 just in case $U \subseteq X_1$. And a reason V for the plaintiff is at least as strong for the plaintiff as the reason $Premise(r_1)$ just in case $Premise(r_1) \subseteq V$. If we let the relation $<_{c_1}$ represent the priority ordering on reasons derived from the particular case c_1 , then, the force of the court's decision in this case is simply that: where U is a reason favoring the defendant and V is a reason favoring the plaintiff, we have $U <_{c_1} V$ just in case $U \subseteq X_1$ and $Premise(r_1) \subseteq V$. To illustrate: Consider the reason $\{f_1^\delta\}$ for the defendant and the reason $\{f_1^\pi, f_2^\pi, f_3^\pi\}$ for the plaintiff. Here, we have $\{f_1^\delta\} \subseteq X_1$ as well as $Premise(r_1) \subseteq \{f_1^\pi, f_2^\pi, f_3^\pi\}$. It therefore

¹⁴When comparing the relative importance of reasons, it is more common to say that one carries greater weight than the other, or that one is weightier than the other. I prefer to speak in terms of priority, rather than weight, for two reasons: first, the priority ordering on reasons to be defined here is nonlinear, while the concept of weight tends to suggest linearity; second, the ordering to be defined here allows only ordinal comparisons among reasons, while the concept of weight suggests that cardinal comparisons must be available as well.

follows that $\{f_1^\delta\} <_{c_1} \{f_1^\pi, f_2^\pi, f_3^\pi\}$ —the court’s decision in the case c_1 entails that the reason $\{f_1^\pi, f_2^\pi, f_3^\pi\}$ favoring the plaintiff is to be assigned a higher priority than the reason $\{f_1^\delta\}$ favoring the defendant.

Generalizing from this example, we reach the following definition of the priority ordering among reasons derived from a single case:

Definition 1 (Priority ordering derived from a case) Let $c = \langle X, r, s \rangle$ be a case, and let U and V be reasons favoring the sides \bar{s} and s respectively. Then the relation $<_c$ representing the priority ordering on reasons derived from the case c is defined by stipulating that $U <_c V$ if and only if $U \subseteq X$ and $Premise(r) \subseteq V$.

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

Definition 2 (Priority ordering derived from a case base) Let Γ be a case base, and let U and V be reasons. Then the relation $<_\Gamma$ representing the priority ordering on reasons derived from the case base Γ is defined by stipulating that $U <_\Gamma V$ if and only if $U <_c V$ for some case c from Γ .

And using this concept of the priority ordering derived from a case base, we can now define a case base itself as inconsistent if its derived ordering yields conflicting information about the priority among reasons—telling us, for some pair of reasons, that each has a higher priority than the other—and consistent otherwise:

Definition 3 (Inconsistent and consistent case bases) Let Γ be a case base with $<_\Gamma$ its derived priority ordering. Then Γ is inconsistent if and only if there are reasons U and V such that $U <_\Gamma V$ and $V <_\Gamma U$, and consistent otherwise.

3.2 Constraint

We now present the reason model of constraint itself, building on the concept of case base consistency. The guiding idea, once again, is that, in deciding a case, a constrained court is required simply to preserve the consistency of the background case base. Suppose, more exactly, that a court constrained by a consistent background case base Γ is confronted with a new fact situation X . Then what the reason model requires is that the court should reach a decision in the situation X that is itself consistent with Γ —that is, a decision that does not introduce inconsistency into the background case base.

The reason model applies, in the first instance, to the rules on the basis of which a court can justify its decisions:

Definition 4 (Reason model constraint on rule selection) Let Γ be a consistent case base and X a fact situation confronting the court. Then the reason model of constraint on rule selection requires the court to base its decision in the situation X on some rule r supporting an outcome s such that the case base $\Gamma \cup \{ \langle X, r, s \rangle \}$ is consistent.

But of course, once this constraint on rule selection is in place, the reason model can naturally be interpreted as requiring a decision for a particular side just in case every rule satisfying the constraint on rule selection supports that side:

Definition 5 (Reason model constraint on decision) Let Γ be a consistent case base and X a fact situation confronting the court. Then the reason model of constraint on decision requires the court to reach a decision in the situation X for the side s if and only if every rule satisfying the reason model constraint on rule selection supports the side s .

These definitions can be illustrated by imagining that the background case base is $\Gamma_1 = \{c_1\}$, containing as its single member the familiar case $c_1 = \langle X_1, r_1, s_1 \rangle$ —where, again, $X_1 = \{f_1^\pi, f_2^\pi, f_1^\delta, f_2^\delta\}$, where $r_1 = \{f_1^\pi\} \rightarrow \pi$, and where $s_1 = \pi$. Suppose that, against this background, the court confronts the fresh situation $X_2 = \{f_1^\pi, f_2^\pi, f_1^\delta, f_2^\delta, f_3^\delta\}$ and considers finding for the defendant in this situation on the basis of the reason $\{f_1^\delta, f_2^\delta\}$, leading to the

decision $c_2 = \langle X_2, r_2, s_2 \rangle$, where X_2 is as above, where $r_2 = \{f_1^\delta, f_2^\delta\} \rightarrow \delta$, and where $s_2 = \delta$. Can the court carry through with this plan of action, according to the reason model?

Well, as we can see, $Premise(r_1) = \{f_1^\pi\}$, the reason for the decision in the initial case, holds in the new situation X_2 as well, since $\{f_1^\pi\} \subseteq X_2$. And of course, the new reason $Premise(r_2) = \{f_1^\delta, f_2^\delta\}$ favors the defendant at least as strongly as itself—that is, $Premise(r_2) \subseteq Premise(r_2)$, or $Premise(r_2) \subseteq \{f_1^\delta, f_2^\delta\}$. It therefore follows from Definition 1 that c_2 , the court’s envisaged decision, would assign the reason $\{f_1^\delta, f_2^\delta\}$ for the defendant a higher priority than the reason $\{f_1^\pi\}$ for the plaintiff—that is, $\{f_1^\pi\} <_{c_2} \{f_1^\delta, f_2^\delta\}$. But Γ_1 already contains the case c_1 , from which, in a similar fashion, we can derive the priority relation $\{f_1^\delta, f_2^\delta\} <_{c_1} \{f_1^\pi\}$, telling us exactly the opposite. Since the augmented case base

$$\begin{aligned}\Gamma_2 &= \Gamma_1 \cup \{c_2\} \\ &= \{c_1, c_2\}\end{aligned}$$

resulting from the court’s envisaged decision contains both these cases, we would then have both $\{f_1^\delta, f_2^\delta\} <_{\Gamma_2} \{f_1^\pi\}$ and $\{f_1^\pi\} <_{\Gamma_2} \{f_1^\delta, f_2^\delta\}$ by Definition 2, so that, by Definition 3, this augmented case base would be inconsistent. By Definition 4, then, we can conclude that the court is not permitted to carry through with its plan of deciding for the defendant in the situation X_2 on the basis of the rule r_2 , since c_2 , the resulting decision, would introduce an inconsistency into the background case base, but the reason model requires decisions to preserve case base consistency.

Of course, it does not follow from the fact that the court cannot decide for the defendant in the situation $X_2 = \{f_1^\pi, f_2^\pi, f_1^\delta, f_2^\delta, f_3^\delta\}$ on the basis of the particular rule r_2 that it cannot decide for the defendant in this situation at all. Suppose, for example, that the court appeals to the reason $\{f_1^\delta, f_3^\delta\}$ to justify its decision for the defendant, leading to the case $c_3 = \langle X_3, r_3, s_3 \rangle$, where $X_3 = X_2$, where $r_3 = \{f_1^\delta, f_3^\delta\} \rightarrow \delta$, and where $s_3 = \delta$. The augmented case base

$$\begin{aligned}\Gamma_3 &= \Gamma_1 \cup \{c_3\} \\ &= \{c_1, c_3\}\end{aligned}$$

resulting from this decision would then be consistent. As before, the previous case c_1 supports

the priority $\{f_1^\delta, f_2^\delta\} <_{c_1} \{f_1^\pi\}$, and the new decision c_3 would now support the priority $\{f_1^\pi\} <_{c_3} \{f_1^\delta, f_3^\delta\}$, so that we would then have both the case base priorities $\{f_1^\delta, f_2^\delta\} <_{\Gamma_3} \{f_1^\pi\}$ and $\{f_1^\pi\} <_{\Gamma_3} \{f_1^\delta, f_3^\delta\}$. But there is nothing inconsistent about this pair of priorities, as we can see, informally at least, with another homely example: one can easily imagine a teenager thinking, and thinking consistently, that going to the movies is more fun than going to the beach with her parents, but that going to the beach with her friends is more fun than going to the movies.

Suppose the court does, in fact, decide the situation X_2 in this way, supplementing the background case Γ_1 with the new decision c_3 , leading to the augmented case base $\Gamma_3 = \Gamma_1 \cup \{c_3\}$. According to the reason model, this decision would then represent a step in the normal development of a legal system, which proceeds more generally as follows: A court confronts a new situation X against the background of a consistent case base Γ , with an associated ordering $<_\Gamma$ on reasons. The court is required to base its decision on some rule r supporting an outcome s such that the case base $\Gamma' = \Gamma \cup \{\langle X, r, s \rangle\}$ is consistent, with the result that the background case base is augmented with this new decision. The next court confronting the next new situation Y must then work against the background of the augmented case base Γ' , which gives rise to the strengthened ordering $<_{\Gamma'}$ on reasons. This new court is likewise required to base its decision on some rule r' supporting an outcome s' such that the case base $\Gamma'' = \Gamma' \cup \{\langle Y, r', s' \rangle\}$ is consistent, thus further augmenting the case base, further strengthening the underlying priority ordering on reasons, and the process continues.

The hypothesis of the reason model is that this is how the common law develops in the normal, incremental case—by building up a stronger and stronger priority ordering on reasons through a series of decisions that are, at each stage, consistent with the existing case base.

3.3 The domestic scenario

All of this has been very abstract. For a more concrete illustration of legal development according to the reason model, we return to the domestic example set out in the Introduction. The example centered around a situation in which Jack and Jo have two children: Emma, who has just turned nine, did her chores, failed to finish dinner, but completed her homework, and Max, age twelve, who also did his chores, but neither finished dinner nor completed homework. Both children wanted to stay up and watch TV. We imagined that Emma first asked Jo, who granted the request to watch TV, justifying her decision with the rule, “Children age nine or greater can stay up and watch TV.” Next, we imagined, Max asked Jack, who denied the request to watch TV, distinguishing this case from that of Emma by appeal to the fact that Max failed to complete his homework, and introducing the new rule “Children who have not completed their homework cannot stay up and watch TV.”

With Max and Emma as plaintiffs, and with Jack and Jo functioning—as parents do—both as defendants and as adjudicators, or courts, this scenario can be cast in our framework by letting the factor f_1^π represent the fact that the child in question is at least nine years old, by letting f_2^π represent the fact that the child in question completed chores, and then letting f_1^δ and f_2^δ represent, respectively, the facts that the child failed to finish dinner and failed to complete homework. The initial situation presented by Emma to Jo can then be represented as $X_4 = \{f_1^\pi, f_2^\pi, f_1^\delta\}$, which Jo then decided for Emma on the basis of the rule $r_4 = \{f_1^\pi\} \rightarrow \pi$, leading to the decision $c_4 = \langle X_4, r_4, s_4 \rangle$, where X_4 and r_4 are as above, and where $s_4 = \pi$. As a result of this initial decision, the case base representing the common law of the household, at least as it pertains to staying up and watching TV, is $\Gamma_4 = \{c_4\}$, with \langle_{Γ_4} as its associated ordering on reasons.

Next, the situation presented by Max to Jack can be represented as $X_5 = \{f_1^\pi, f_2^\pi, f_1^\delta, f_2^\delta\}$. In keeping with our story, we can suppose that Jack would like to decide against Max on the basis of the rule $r_5 = \{f_2^\delta\} \rightarrow \delta$, leading to the decision $c_5 = \langle X_5, r_5, s_5 \rangle$, where X_5 and r_5 are as above, and where $s_5 = \delta$. Can he do so, according to the reason model, against the background of the case base Γ_4 ? The answer is Yes. From Jo’s earlier decision,

we can conclude that the reason $\{f_1^\pi\}$ is to be assigned a higher priority than the reason $\{f_1^\delta\}$ —that $\{f_1^\delta\} <_{c_4} \{f_1^\pi\}$, so that $\{f_1^\delta\} <_{\Gamma_4} \{f_1^\pi\}$ as well. And Jack’s decision would force us to conclude also that the reason $\{f_2^\delta\}$ must be assigned a higher priority than the reason $\{f_1^\pi\}$ —that $\{f_1^\pi\} <_{c_5} \{f_2^\delta\}$. But there is no conflict between this priority statement and the previous priority statement, derived from Jo’s decision—a reasonable individual might, for example, prefer chocolate ice cream to vanilla and vanilla to strawberry. And because the background case base Γ_4 currently contains only Jo’s decision, it follows that Jack’s decision in the case of Max is consistent with this case base as well. The reason model thus allows Jack to carry through with his decision, resulting in

$$\begin{aligned}\Gamma_5 &= \Gamma_4 \cup \{c_5\} \\ &= \{c_4, c_5\}\end{aligned}$$

as the updated case base now representing the household common law, with $<_{\Gamma_5}$ as its strengthened ordering on reasons.

4 Discussion

I now discuss the relation of the reason model of constraint to Lamond’s original account, try to clear up what might be some misunderstandings of the reason model, and address some related issues.

4.1 Defeat and consistency

First, what is the relation of the reason model to Lamond’s original “reasons account”? My view is that the reason model provides one particular interpretation of Lamond’s account, arrived at by defining the key concept of consistency in one particular way.

Let us begin with Lamond’s initial statement of his position:

The starting point for a reason-based account, then, is a reflection upon what is, in truth, the fundamental requirement of the common-law doctrine of precedent, namely, that later courts treat earlier cases as correctly decided on their facts.

What flows from this requirement is the necessity of reaching a decision in a later case that is consistent with the correctness of the earlier decision. The question is what sort of consistency does this entail? The conventional view interprets consistency in terms of the use of the same “rule” as that relied upon in the precedent case. Those who demur on the significance of *rationes* interpret it as entailing the application of the same set of principles that justified the preceding case(s). The reason-based view requires that later courts accept the correctness of the precedent courts assessment of the balance of reasons on the facts of the precedent case.¹⁵

There is a lot going on here: Precedent requires later courts to regard earlier cases as correctly decided, and what this comes to is simply that later courts are required to reach decisions that are consistent with the decisions reached in earlier cases. Consistency is to be interpreted, not in terms of rules or principles, but as acceptance of the precedent court’s assessment of a certain relation among reasons.

This passage of Lamonds’s is so interesting and innovative that it seems churlish to note that it does not provide a precise definition of what it means for a later decision to be consistent with an earlier decision. Nevertheless, that is what I want to call attention to, since it is exactly this gap that the reason model fills, and fills in a particular way.

In order to see how, we turn first to Lamond’s characterization of the meaning of a particular decision. He imagines—using slightly different notation than that employed here—that a court faced with a fact situation of the form

$$\{g_1, h_1, i_1, j_1, k_1, l_1\}$$

reaches a decision for C on the basis of the rule

$$\text{If } \{J, K, L\}, \text{ then } C,$$

and then describes the significance of this decision as follows:

¹⁵Lamond (2005, p. 3).

What the court decided was that in the context of $\{g_1, h_1, i_1, j_1, k_1, l_1\}$, the features $\{J, K, L\}$ justified the conclusion C . There are two aspects to what the court decided. The first is that the presence of features $\{J, K, L\}$ provided a sufficient reason to conclude C . The second is that the presence of the other features $\{G, H, I\}$ did not defeat the reason provided by $\{J, K, L\}$.¹⁶

What does it mean to say that one reason defeats another? In philosophy, there are two central notions of defeat, a contrast due to John Pollock, but more familiar in legal theory from the work of Joseph Raz.¹⁷ There is, first of all, *rebutting* defeat, which occurs when a reason for one action or conclusion is defeated by a stronger, or more important, reason for an inconsistent action or conclusion. To illustrate: Suppose I promise to meet Emma for lunch but encounter a drowning child along the way. Then my promise, which provides a reason for having lunch with Emma, is rebutted by the fact that there is a drowning child, which provides a stronger reason to rescue the child instead, an action that is inconsistent with meeting Emma for lunch. In addition to rebutting defeat, however, there is also, as Pollock points out, the phenomenon of *undercutting* defeat—or in Raz’s terminology,

¹⁶Lamond (2005, pp. 15–16). There are two differences between Lamond’s notation and that employed here. The first is that, for Lamond, situations are described in terms of particular facts, or tokens, while rules are described in terms of kinds of particular facts, or types. Thus, for example, a situation might contain x_i to represent the fact that the particular child Emma, say, has not finished dinner, while a rule might contain X to represent the general fact that the arbitrary child in question has not finished dinner—this notation, which is similar to that used in Raz (1979), is then governed by the convention that the token x_i falls under the type X . By contrast, in the current framework, the same factors occur in both situations and rules; the assumption is that situations are already described in terms of types, rather than tokens—that degree of abstraction is assumed to precede our analysis. This first notational difference is, I believe, inconsequential. The second is more significant: Lamond’s situations are described in terms of facts that do not, intrinsically, favor one side or the other; he refers to these in Lamond (2021) as “features.” In the current framework, by contrast, situations are described in terms of factors with a polarity, favoring some particular side—again, it is assumed that the classification of a mere feature as a factor, favoring a particular side, precedes our analysis. We return to this second difference below.

¹⁷See, for example, Pollock (1970) and Raz (1975); the distinction received its first formal treatment in Pollock (1987).

exclusionary defeat—which occurs when what would ordinarily count as a reason for an action or conclusion is taken out of consideration, or deprived of its force, by another reason. To illustrate: Imagine that Max contracts to deliver 150 bales of hay to a local farmer. Generally, such a contract would be taken as a reason to deliver the hay. But suppose it is discovered that Max is fourteen years old, and so a minor. The fact that Max is a minor does not, of course, provide any reason for him not to deliver the hay. But, since minors cannot enter into contracts, it takes the contract out of consideration—it undercuts, or excludes, the contract—as a reason to deliver the hay.

The reason model set out here can be seen as resulting from Lamond’s general description in two steps. First, we suppose that, in examples such as that set out by Lamond above, the kind of defeat involved is rebutting defeat, rather than undercutting or exclusionary defeat. On Lamond’s view, recall, what the decision he considers tells us is that $\{J, K, L\}$ provides a reason for C , and that this reason is not defeated by $\{G, H, I\}$. Now, if the notion of defeat involved is rebutting defeat, then $\{G, H, I\}$ defeats $\{J, K, L\}$ just in case $\{G, H, I\}$ provides a stronger reason for the opposite conclusion, say $\neg C$. So, if we can suppose that $\{G, H, I\}$ is, in fact, a reason for the opposite conclusion—and otherwise, how could it defeat $\{J, K, L\}$ at all—then what we can conclude from the fact that $\{G, H, I\}$ does not defeat $\{J, K, L\}$ is simply that $\{G, H, I\}$ does not provide a stronger reason for $\neg C$ than $\{J, K, L\}$ provides for C —in other words, and given certain natural assumptions about the ordering relation, that $\{J, K, L\}$ is stronger, or has higher priority, than $\{G, H, I\}$. On this view, then, what we derive from a precedent decision is a priority ordering on reasons, exactly as in Definition 1 of the current reason model.

So that is the first step in moving from Lamond’s general description to the specific interpretation provided by the reason model: taking the kind of defeat involved as rebutting defeat, so that a previous decision can be understood as generating a priority ordering on reasons. The second step is even simpler. We can at last define a precise notion of consistency to fill the gap in Lamond’s account: using the priority just introduced, a new decision can be characterized as consistent with a previous set of cases just in case the priority ordering

derived from the new decision does not conflict with that derived from any of the previous cases—just in case, that is, there are no two reasons U and V such that the new decision assigns U a higher priority than V while some previous case assigns V a higher priority than U . Evidently, just as our first step runs parallel to Definition 1 of the current reason model, this second step, runs parallel to the current Definitions 2 and 3. The requirement that new decisions preserve consistency is now captured in Definitions 4 and 5.

I have traced one interpretative path, leading from Lamond’s account to my own reason model of constraint, which results from treating the idea of defeat as rebutting defeat, and then using the associated priority ordering to define consistency. But there are other paths as well. In particular, Mullins’s interpretation, mentioned above—which he describes as a *protected* reason model, to distinguish it from the present *prioritized* reason model—can be arrived at in a similar fashion, except that it is based on the notion of undercutting or exclusionary, rather than rebutting, defeat, and then defines a different notion of consistency based on those ideas.

4.2 Reasons and rules

I have argued that Lamond’s central idea is that precedential constraint can be defined entirely in terms of relations among reasons, without any appeal to rules, and also that my reason model is an interpretation of Lamond’s account. Yet the reason model presented here is constructed on the bases of a formal representation of cases as containing rules: for example, the case $c_1 = \langle X_1, r_1, s_1 \rangle$ considered earlier contains the defeasible rule $r_1 = \{f_1^\pi\} \rightarrow \pi$. How, then, can I say that the reason model does not depend on rules?

To this question I have a simple answer, and then a qualification. The simple answer is that, although rules are indeed present in the representation of cases, they do absolutely no work in our definition of the reason model. With suitable adjustments elsewhere in our series of definitions, the case c_1 , for instance, could just as easily have been represented as $c_1' = \langle X_1, \{f_1^\pi\}, s_1 \rangle$ containing the reason that forms the premise of the rule r_1 , in place of the rule itself. It is now natural to ask: If case rules are not actually used in the definition of

the reason model, why are they present in our representation of cases at all? The answer is that, although rules are not used in the definition of the reason model itself, they can, in fact, play an illuminating role in further development of the theory. It can be shown, for example, that, when case rules are actually put to use as rules, the reason model of constraint can then be reconstructed within a defeasible logic, and also that the reason model is equivalent to a version of the standard model, based on rule modification.¹⁸

So that is the simple answer: Rules play no essential role. The reason model can be formulated in terms of reasons alone. The qualification—which I hesitate to mention, but will—is that, in other work, I have argued that reasons themselves can be identified as the premises of appropriate default rules.¹⁹ What does it mean to say that my promise to have lunch with Emma is a reason for doing so? It means that my reasoning—or the proper evaluation of my reasoning—is guided by a rule according to which, if I promise to have lunch with Emma, then, by default, I ought to do so. Why would my reasoning, or the proper evaluation of my reasoning, be guided by such a default? Either because I take, or because I should take, promising to have lunch with Emma as a reason for doing so.

4.3 Features, factors, disjunction, negation

I want to close by mentioning two other issues that come up in Lamond’s more recent discussion.

First, in contrasting his reasons account of precedent with my reason model, Lamond emphasizes as a significant difference that his own account starts with “features” of situations while the present reason model starts with “factors”—where features are mere factual aspects of a situation, and factors are legally significant patterns of facts that are taken to favor one side or the other.²⁰ Lamond emphasizes that important legal reasoning is involved in understanding a situation, presented as a set of features, in terms of factors. I agree. Indeed,

¹⁸See Horty (2016) for a reconstruction of the rule model within a defeasible logic, and Horty (2015) for connections between the rule model and the standard model.

¹⁹See the Introduction to Horty (2012).

²⁰See Lamond (2021, Section 2.3).

I wrote in my initial presentation of the reason model that:

... the mere ability to understand a case in terms of the factors it presents itself requires a significant degree of legal expertise, which is presupposed here. Our theory thus starts with cases to which we must imagine that this expertise has already been applied, so that they can be represented directly in terms of the factors involved; we are concerned here only with the subsequent reasoning.²¹

The reason model, then, is entirely concerned with issues that arise “downstream” of the representation of a situation as a particular set of factors, and not with issues involved in arriving at the appropriate representation—and this holds not just of the reason model, but of the entire body of work in artificial intelligence and law that takes for granted the factor-based representation of legal information.

I do feel, though, that Lamond raises an important question here, and that is whether the overall process of legal understanding, argumentation, and decision can legitimately be divided into an “upstream” part that is devoted to arriving at an appropriate factor-based representation of the set of features present in some situation, and then a downstream part that is devoted to arriving at an appropriate decision, supported by appropriate arguments, based on that representation. It may well be, as Lamond suggests, that treating these two stages of the overall reasoning process as entirely separate modules—in the way that the factor-based approach suggests—forces us to ignore important interactions between these two modules.

A second issue that Lamond mentions concerns the representation, within accounts such as the reason model, of disjunctive and negative information.²² This is a problem I have thought about as well, with little to show for it. My own lack of progress is due to two obstacles: I am not sure, first, exactly how to represent information like this within the reason model framework, and second, I wonder about the extent to which we can do so consistent with the traditional common law distinction between *ratio decidendi* and *obiter*

²¹Horty (2012, p. 5).

²²See Lamond (2021, Sections 2.2 and 2.3).

dicta.²³

To illustrate the problem, let us return to the situation presented by Emma to Jo. As we recall, in representing this situation, we took the factor f_1^π to represent the fact that the child in question is at least nine years old, the factor f_2^π to represent the fact that the child in question completed chores, and the factor f_1^δ to represent the fact that the child failed to finish dinner. Since Emma had reached the age of nine and completed chores, then, but did not finish dinner, we imagined that the situation she presented to Jo was $X_4 = \{f_1^\pi, f_2^\pi, f_1^\delta\}$. Now suppose, as before, that Jo wants to decide for Emma. What are her options for justifying her decision? According to the reason model, she has three options. She can justify her decision on the basis of the fact that Emma is nine years old, on the basis of the fact that Emma has completed chores, or on the basis of the fact that Emma is nine years old and has completed chores—on the basis, that is, of the reasons $\{f_1^\pi\}$, $\{f_2^\pi\}$, or $\{f_1^\pi, f_2^\pi\}$.

Suppose, however, that what Jo thinks is that either Emma’s age or her completion of chores would justify a decision in her favor. With a rule, this would be easy to express, as, say,

$$\{f_1^\pi\} \vee \{f_2^\pi\} \rightarrow \pi,$$

where \vee is the standard logical symbol for disjunction. But within the reason model, this is a much more difficult idea to express. One option would be to liberalize the ways in which Jo, through her judgments, can modify the background case base. Perhaps Jo could register her disjunctive judgment by supplementing the existing case base, not just with a single new case, but with both the new cases $c_6 = \langle X_6, r_6, s_6 \rangle$ and $c_7 = \langle X_7, r_7, s_7 \rangle$, where $X_6 = X_7 = X_4$, where $r_6 = \{f_1^\pi\} \rightarrow \pi$ and $r_7 = \{f_2^\pi\} \rightarrow \pi$, and where $s_6 = s_7 = s_4$. As the reader can verify, this would then have the effect of prioritizing both the reason $\{f_1^\pi\}$, and the reason $\{f_2^\pi\}$ over the opposing reason $\{f_1^\delta\}$.

Or suppose what Jo thinks is that Emma’s age justifies a decision in her favor but that,

²³See Chapter 2 of Cross and Harris (1991) for a discussion of this distinction in traditional terms, then Branting (1993) for an analysis within artificial intelligence and law

taken on its own, her completion of chores does not. It is easy enough to express the idea that Emma’s age justifies a decision in her favor—the previous $c_4 = \langle X_4, r_4, s_4 \rangle$ did that—but even if we allow Jo to supplement the background case base with multiple new cases, how could she express the negative information that $\{f_2^\pi\}$ itself does not have higher priority than $\{f_1^\delta\}$? Jo might try to express this information with the case $c_8 = \langle X_8, r_8, s_8 \rangle$, where $X_8 = \{f_2^\pi, f_1^\delta\}$, where $r_8 = \{f_1^\delta\} \rightarrow \delta$, and where $s_8 = \delta$. As the reader can verify, once again, this would allow us to conclude that $\{f_1^\delta\}$ is assigned a higher priority than $\{f_2^\pi\}$. But that seems too strong. Jo’s intention is only to deny that $\{f_2^\pi\}$ carries a higher priority than $\{f_1^\delta\}$, not necessarily to affirm that $\{f_1^\delta\}$ carries a higher priority than $\{f_2^\pi\}$.

Or suppose what Jo wants to express is that Emma’s age justifies a decision in her favor, but also—perhaps anticipating the case later presented by Max to Jack, and hoping to block Jack’s decision—that Emma’s age would still justify such a decision even if she had not done her homework. Again, the positive component of Jo’s judgment could be carried by the previous $c_4 = \langle X_4, r_4, s_4 \rangle$, but the negative component would have to draw on factors not even present in the situation presented by Emma to Jo.

These examples are meant to illustrate the technical problems that arise in representing disjunctive or negative information within the reason model, but I hope that they illustrate the conceptual problems as well. The negative information involved in our second two examples seems to force us to consider situations other than that presented by Emma to Jo, and although the distinction between *ratio* and *dicta* is notoriously hard to formulate, it is generally agreed that judgments concerning hypothetical situations, not explicitly before the court, are to be classified as *dicta*.

5 Conclusion

My primary goal in this paper has been to present, as simply and clearly as possible, the reason model of precedential constraint, which I developed as an interpretation of Lamond’s idea that precedent in the common law could usefully be understood without any appeal to rules at all. This idea of Lamond’s has given rise to other interpretations as well, and

suggests a number of new paths for research.

I close simply by mentioning one of these, which arises when we reflect on the process of common law development—as the reason model invites us to do—from the standpoint of social choice theory. What the common law yields, according to the reason model, is a social ordering on reasons, constructed as a result of decisions by individual courts, presumably on the basis of their own individual orderings on reasons, their own preferences. The common law can thus be seen as the realization of a particular kind of preference aggregation procedure—a particular way of aggregating individual preferences, or priorities, among reasons into a group preference. Unlike the preference aggregation functions usually studied in social choice theory, however, the procedure realized by the common law does not simply take a collection of individual preferences as inputs and output a group preference all at once. Instead, the common law constructs its group preference ranking on reasons through a procedure that is piecemeal, distributed, and responsive to particular circumstances. Especially given the challenges to common law reasoning that have been advanced over the years, it will be interesting to explore the extent to which the preference aggregation procedure realized by the common law can be justified using tools and ideas from social choice theory.²⁴

Acknowledgments

I am grateful for comments from Ilaria Canavotto, Mark McBride, and Adam Rigoni on an earlier draft of this paper.

References

- Aleven, Vincent (1997). *Teaching Case-Based Argumentation Through a Model and Examples*. PhD thesis, Intelligent Systems Program, University of Pittsburgh.
- Aleven, Vincent and Ashley, Kevin (1997). Evaluating a learning environment for case-based argumentation skills. In *Proceedings of the Sixth International Conference on Artificial In-*

²⁴For some initial explorations, see Horty and Pacuit (2021).

- telligence and Law (ICAAIL-97)*, pages 170–179. The Association for Computing Machinery Press.
- Alexander, Larry (1989). Constrained by precedent. *Southern California Law Review*, 63:1–64.
- Alexander, Larry and Sherwin, Emily (2001). *The Rule of Rules: Morality, Rules, and the Dilemmas of Law*. Duke University Press.
- Alexander, Larry and Sherwin, Emily (2008). *Demystifying Legal Reasoning*. Cambridge University Press.
- Ashley, Kevin (1990). *Modeling Legal Argument: Reasoning with Cases and Hypotheticals*. The MIT Press.
- Branting, L. Karl (1993). A reduction-graph model of ratio decidendi. In *Proceedings of the Fourth International Conference on Artificial Intelligence and Law (ICAAIL-93)*, pages 40–49. ACM Press.
- Broughton, Gabriel (2019). Vertical precedents in formal models of precedential constraint. *Artificial Intelligence and Law*, 27:253–307.
- Cross, Rupert and Harris, J. W. (1991). *Precedent in English Law*. Oxford University Press. Fourth edition.
- Dancy, Jonathan (2004). *Ethics Without Principles*. Oxford University Press.
- Eisenberg, Melvin (1988). *The Nature of the Common Law*. Harvard University Press.
- Horty, John (2011). Rules and reasons in the theory of precedent. *Legal Theory*, 17:1–33.
- Horty, John (2012). *Reasons as Defaults*. Oxford University Press.
- Horty, John (2015). Constraint and freedom in the common law. *Philosopher’s Imprint*, 15(25).

- Horty, John (2016). Reasoning with precedent as constrained natural reasoning. In Lord, Errol and Maguire, Barry, editors, *Weighing Reasons*, pages 193–212. Oxford University Press.
- Horty, John (2019). Reasoning with dimensions and magnitudes. *Artificial Intelligence and Law*, 27:307–345.
- Horty, John and Bench-Capon, Trevor (2012). A factor-based definition of precedential constraint. *Artificial Intelligence and Law*, 20:181–214.
- Horty, John and Pacuit, Eric (2021). The common law as the solution to a social choice problem. Manuscript.
- Lamond, Grant (2005). Do precedents create rules? *Legal Theory*, 11:1–26.
- Lamond, Grant (2021). Revisiting the reasons account of precedent. Unpublished manuscript.
- Lance, Mark and Little, Margaret (2004). Defeasibility and the normative grasp of context. *Erkenntnis*, 61:435–455.
- Lance, Mark and Little, Margaret (2007). Where the laws are. In Shafer-Landau, Russ, editor, *Oxford Studies in Metaethics*, volume 2, pages 149–171. Oxford University Press.
- Levi, Edward (1949). *An Introduction to Legal Reasoning*. The University of Chicago Press.
- Mullins, Robert (2020). Protected reasons and precedential constraint. *Legal Theory*, 26:40–61.
- Pollock, John (1970). The structure of epistemic justification. In *Studies in the Theory of Knowledge*, American Philosophical Quarterly Monograph Series, number 4, pages 62–78. Basil Blackwell Publisher, Inc.
- Pollock, John (1987). Defeasible reasoning. *Cognitive Science*, 11:481–518.

- Prakken, Henry (20xx). A formal analysis of some factor- and precedent-based accounts of precedential constraint. *Artificial Intelligence and Law*, forthcoming.
- Prakken, Henry and Sartor, Giovanni (1998). Modelling reasoning with precedents in a formal dialogue game. *Artificial Intelligence and Law*, 6:231–287.
- Raz, Joseph (1975). *Practical Reason 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, Joseph (1979). *The Authority of Law*. Oxford University Press.
- Rigoni, Adam (2014). Common-law judicial reasoning and analogy. *Legal theory*, 20:133–156.
- Rigoni, Adam (2015). An improved factor based approach to precedential constraint. *Artificial Intelligence and Law*, 23:133–160.
- Rissland, Edwina (1990). Artificial intelligence and law: stepping stones to a model of legal reasoning. *The Yale Law Journal*, 99:1957–1981.
- Rissland, Edwina and Ashley, Kevin (1987). A case-based system for trade secrets law. In *Proceedings of the First International Conference on Artificial Intelligence and Law (ICAIL-87)*, pages 60–66. The Association for Computing Machinery Press.
- Rissland, Edwina and Ashley, Kevin (2002). A note on dimensions and factors. *Artificial Intelligence and Law*, 10:65–77.
- Schauer, Frederick (1989). Is the common law law? *California Law Review*, 77:455–471.
- Schauer, Frederick (1991). *Playing by the Rules: A Philosophical Examination of Rule-Based Decision-Making in Law and Life*. Oxford University Press.
- Simpson, A. W. B. (1961). The *ratio decidendi* of a case and the doctrine of binding precedent. In Guest, A. G., editor, *Oxford Essays in Jurisprudence*, pages 148–175. Oxford University Press.