

Introduction

This book presents an account of reasoning, or decision making, against a background of constraints derived from a body of previous authoritative decisions, or precedent cases. The nature and operation of the constraints derived from precedent cases is a central concern of legal theory in the common law tradition and has received its most careful analysis there. In developing my account, I will draw extensively on the literature from legal theory, and the account I present is intended to capture important aspects of precedential constraint in the common law.

At the same time, I want to suggest that the process of reasoning against a background of constraints derived from precedent cases, as in the common law, is not a special style of reasoning learned in law school and practiced only by lawyers. It is, instead, a familiar form of reasoning that we all engage in all the time – in our development of small-scale normative systems, for example, or in our use of open-textured predicates. An important goal of the project is to isolate this form of reasoning as an interesting topic of investigation in its own right, both for philosophy and also for logic, even apart from its varied applications within particular legal institutions. Finally, I will argue that the entire framework – the constraints derived from precedent cases, the associated process of reasoning – is rational, in the sense that it provides, if not a unique optimal solution, then at least one sensible solution to the problem of combining individual preferences and values into a shared priority ordering on reasons that can be used to support decisions affecting members of larger communities. In this sense, the account presented here can be seen as offering a justification for practices in the common law that many writers have taken to be objectionable.

Traditionally, the common law is distinguished from statutory or regulatory law, which is based on a collection of rules crafted by legislators or administrative agencies. Although the proper understanding of statutes and regulations can present complex issues of interpretation, often centered

around questions of legislative meaning or intent, the underlying mechanism of constraint, in this case, is at least straightforward. Setting matters of interpretation aside, statutory and regulatory rules work just like any other rules. To take a mundane example, if you are sorting marbles according to a set of directives that includes the rule “Blue marbles must be placed in the blue bin,” what this means is that, if you come across a blue marble, you must place it in the blue bin. If, instead, you place a blue marble in the red bin, you fail to satisfy the constraints imposed by the rule – you have violated the rule. In exactly the same way, if a state’s motor vehicle code includes the regulation “Vehicles operated on public roads must be registered with the Department of Motor Vehicles,” what this means is that, if you operate a vehicle on public roads, that vehicle must be registered with the Department of Motor Vehicles, and you have violated the regulation if you operate a vehicle that is not registered.

The common law, by contrast, is organized around a collection of previous decisions, or precedent cases, rather than a collection of rules. Situations requiring decisions come before authoritative bodies – typically courts – and decisions are rendered. The common law itself is then thought to emerge from these decisions through *stare decisis*, a complex doctrine according to which, on the most common interpretation, decisions reached in earlier, or precedent, cases constrain the decisions available to courts in later cases, while still allowing these later courts a degree of freedom to respond in creative ways to fresh circumstances. In understanding the nature of common law constraint, the central challenge is to determine the precise character of this doctrine of *stare decisis* – or more simply, the doctrine of *precedent*.¹ Although techniques for arguing on the basis of precedent are studied early on in law schools, mastered with relative ease, and applied on a daily basis by legal practitioners, it has proved to be considerably more difficult to arrive at an adequate theoretical account of the doctrine itself. Such an account would have to answer two questions, at least: 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 imposed by precedent and the freedoms allowed to later courts for developing the law.

¹ The literature on precedent is immense, but useful introductions are provided in Schauer (1987) and Alexander (1989). An illuminating survey of the issues posed by precedent from a philosophical perspective can be found in Lamond (2006); a detailed discussion of precedent in one particular legal system can be found in Cross and Harris (1991).

Let us begin with the familiar, and very natural, position that constraint in the common law, like statutory or regulatory constraint, also depends on rules. A precedent case normally contains, not only a factual description of some situation together with a decision on the basis of those facts, but some particular rule through which that decision is justified, the *ratio decidendi* of the case. And according to the position under consideration, it is this rule that carries precedential constraint. On this view – which can be described as the *rule model* of precedent – the common law is not so different from statutory or regulatory law.² The mechanism of constraint, in each case, depends on previously formulated rules. Just as statutory and regulatory law is based on statutes and regulations, the common law is likewise based on rules, with the sole difference that common law rules are formulated by courts, rather than legislators or administrative agencies.

The general idea that common law constraint is based on rules can be developed in a number of different ways, depending, in the first instance, on the nature of the rules involved. Some writers argue that common law rules are best seen as defeasible – *prima facie*, or *pro tanto* – leading to what I will call the *defeasible rule model* of constraint. I have a good deal of sympathy with this position, which has been, I believe, underexplored in legal theory for the simple reason that, until recently, there was no adequate account available of reasoning on the basis of defeasible rules: how could we hope to explicate common law constraint using defeasible rules without understanding the logical behavior of defeasible rules themselves? Now, however, as a result of the sustained study of defeasible, or default, logics carried out during the past several decades within the fields of logic, philosophy, and computer science, we have a much better understanding of defeasible reasoning in general. In my own previous work, I have argued that a particular default logic can be interpreted as providing a precise, mathematical theory of the way in which reasons support conclusions in ethics and epistemology; and later in the book, I will extend this interpretation to show how the central account of precedential constraint developed here can be reformulated

² I follow Alexander (1989) and Alexander and Sherwin (2008) in classifying the positions available for understanding precedential constraint into a taxonomy of *models*; my taxonomy, however, differs from theirs both in allowing for more possibilities and in separating certain positions that they group together. In particular, while Alexander and Sherwin use the phrase “rule model” to refer only to the position described below as the “serious rule model,” I use this phrase as a more general classification to encompass a variety of positions organized around the manipulation of rules.

within that default logic, a result that I take as a vindication of the defeasible rule model of constraint.³

For the present, however, I want to start with the opposite position, that the generalizations expressed by common law rules are not defeasible at all, but strict, or exceptionless. On this view, common law rules are thought to mirror the kind of exceptionless generalizations studied in ordinary symbolic logic. The idea is that rules such as “Vehicles are not allowed in the park,” for example, or “Contracts with minors are voidable” are like the statements “All equilateral triangles are equiangular” or “All men are mortal” – there can be no exceptions.

As it turns out, the position that common law rules state exceptionless generalizations can itself be developed in two ways, this time depending, not on the nature, or meaning, of the rules themselves, but on the system of conventions governing their use. Some writers argue that, once a common law rule has been introduced in an earlier case, it must then control any later case in which it is applicable, unless the court in the later case wishes to overrule the earlier decision and has the authority to do so. This hardheaded position can be described as the *serious rule model* of constraint.⁴ Other writers favor a more flexible approach, according to which, although only certain courts, depending on their place in the judicial hierarchy, have the authority to overrule earlier decisions, all courts have the power to *distinguish* the fact situation of a case currently under consideration from that of some earlier case in which a precedent rule was formulated – the power, that is, to identify important, or material, differences between the two situations, which the court can then take as justification for modifying the earlier rule in order to avoid an inappropriate application of this rule to the current situation.

The idea that general rules must, on occasion, be modified to avoid inappropriate application in particular situations goes back to antiquity, and receives its most notable treatment in Book V of Aristotle’s *Nichomachean Ethics*. There, in the course of his discussion of equity, Aristotle emphasizes that – due to the open-ended variability of human circumstances – it is inevitable that legal generalizations will, at times, yield incorrect results:

³ See Horty (2012) for my previous work interpreting reasons within default logic.

⁴ The position has been developed and defended with great force by Alexander and Sherwin; see Alexander (1989) and then, especially, Alexander and Sherwin (2001, 2008).

The reason is that all law is universal but about some things it is not possible to make a universal statement which shall be correct. In those cases, then, in which it is necessary to speak universally, but not possible to do so correctly, the law takes the usual case, though it is not ignorant of the possibility of error.⁵

In situations like this, as Aristotle explains, when the straightforward application of legal generalizations leads to a flawed outcome, it is necessary to offer a correction:

When the law speaks universally, then, and a case arises on it which is not covered by the universal statement, then it is right, when the legislator fails us and has erred by over-simplicity, to correct the omission – to say what the legislator himself would have said had he been present, and would have put into his law if he had known.⁶

And he illustrates his conclusion that legal generalizations must be modified to fit particular cases with a striking analogy to the “Lesbian rule” – a ruler made of soft lead from the island of Lesbos, which could be adjusted to measure irregular shapes:

In fact this is the reason why all things are not determined by law, viz. that about some things it is impossible to lay down a law, so that a decree is needed. For when the thing is indefinite the rule also is indefinite, like the leaden rule used in making the Lesbian moulding; the rule adapts itself to the shape of the stone and is not rigid, and so too the decree is adapted to the facts.⁷

Turning to the common law proper, a canonical statement of the position that rules are malleable in this way can be found in the work of Edward Levi, who begins his monograph on legal reasoning by considering the hypothesis that the common law could be approached as if it depended on

⁵ NE 1137b12-17 (Ross translation, here and below).

⁶ NE 1137b19-24.

⁷ NE 1137b28-31. Although this concern with the rigidity of general rules is most closely associated with Aristotle, it goes back at least to Plato, who writes in the *Statesman* that law “cannot prescribe with perfect accuracy what is good and right for each member of the community,” since “[t]he differences of human personality, the variety of men’s activities, and the inevitable unsettlement attending all human experience make it impossible for any art whatsoever to issue unqualified rules holding good on all questions at all times” (294b1-6; Skemp translation, here and below). Unfortunately, as he goes on to say, “the law tends to issue just this invariable kind of rule,” and so “it is like a self-willed, ignorant man who lets no one do anything but what he has ordered and forbids all subsequent questioning” (294b8-c2).

a method of applying general rules of law to diverse facts – in short, as though the doctrine of precedent meant that general rules, once properly determined, remained unchanged, and then were applied, albeit imperfectly, in later cases.⁸

Levi immediately rejects this view, however, writing that, to the contrary, common law rules “change from case to case and are remade with each case,” and that “the rules change as the rules are applied.”⁹ The view that common law rules can be adapted to fit new situations is, arguably, the most prevalent position among legal theorists and provides what I will refer to as the *standard model* of precedential constraint.¹⁰ This standard model is often thought to offer the most accurate picture of ordinary, incremental legal development, through the gradual modification of common law rules in light of later cases. It could be illustrated by tracing the development of an actual common law doctrine, but it will be simpler – and also help to illustrate the point that common law reasoning is already at work in everyday situations – to begin with a more mundane example.¹¹

Suppose, then, 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 precedents. And imagine that, one night, Emma, who has completed both her chores and her homework, 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 completed chores and failed to finish dinner, but who has, in addition, failed to complete

⁸ Levi (1949, p. 2).

⁹ *Ibid.*, pp. 3–4.

¹⁰ Versions of this position have been developed, not just by Levi but by Simpson (1961), Raz (1979), Eisenberg (1988), and Schauer (1989, 1991), along with many others.

¹¹ For a legal example, see Levi’s (1949, pp. 8–27) discussion of the development, within the standard model, of the changing common law rules characterizing the notion of an “inherently dangerous” object. Other artificial examples of the sort considered here can be found in Burton (1985), and especially in Twining and Miers (2010).

homework, asks Jack whether he can stay up and watch TV. And suppose that, in this case, Jack refuses, on the grounds that Max has not completed his homework. Max might reasonably appeal Jack's decision with the complaint, "Ah, but given the precedent established last night, in the case of Emma, our household is now governed by a rule according to which children age nine or greater can stay up and watch TV." And of course, according to the serious rule model of constraint, Max would be right, since the rule established in the case of Emma is applicable in his situation as well. The standard model, however, allows Jack to defend his 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. The overall effect of Jack's decision, according to the standard model, is that the household legal system is changed in two ways. It now contains, first of all, a new rule to justify Jack's decision in the case of Max – perhaps the rule, "Children who have not completed their homework cannot stay up and watch TV." And second, the rule previously set out by Jo in the case of Emma has now been modified in order to avoid an unwanted application in the latter case – perhaps reading, "Children age nine or greater can stay up and watch TV, unless they have failed to complete their homework."

Various proposals have been offered about how, exactly, Jack's modification of Jo's previous rule might be justified. Maybe Jack thinks his modified rule provides a better representation than Jo's original formulation of what she had in mind to begin with, or that – whether or not this is what she had in mind – the new rule is the one Jo would have set out had she found herself in, or at least envisioned, the new situation presented by Max.¹² Or maybe Jack simply feels that the overall household regulatory system is sufficiently improved by his modification of Jo's rule. Regardless of justification, however, the fact that the standard model of constraint allows Jack to modify Jo's rule at all leads to a conceptual problem concerning the very notion of rule-based constraint: if Jack is indeed able to reformulate Jo's earlier rule in order to avoid unwanted application in a later case, how can he be thought of as constrained by that rule? Or more generally: if the constraints imposed

¹² The idea that later modifications of a common law rule are meant to capture what the "original court had in mind" is explored by Raz (1979, p. 188); Aristotle's picture is, evidently, based on the counterfactual condition – that a "decree," or rule modification, is meant to represent "what the legislator himself would have said had he been present, and would have put into his law if he had known."

by decisions of earlier courts are supposed to be carried by rules, but later courts are free to modify these rules in order to avoid unwelcome outcomes, how can these rules impose any real constraints at all – how can courts be constrained by rules that they are free to modify at will?

In fact, the literature on the standard model contains a response to this problem – first set out explicitly by Joseph Raz, although, as Raz notes, it owes much to the previous work of A. W. B. Simpson.¹³ The central idea is that, although later courts are indeed free to modify the rules set out by earlier courts, they are not free to modify these rules entirely at will. Any later modification of an earlier rule must satisfy two conditions: First, the modification can consist only in the addition of further restrictions, narrowing the original rule. And second, the modified rule must continue to support the original outcome in the case in which it was introduced, as well as in any further cases in which this rule was appealed to as a justification.

The force of these *Raz/Simpson conditions* on rule modification can be illustrated by returning to our domestic example, where Jo's initial rule, "Children age nine or greater can stay up and watch TV," introduced in the case of Emma, was later modified by Jack to read "Children age nine or greater can stay up and watch TV, unless they have failed to complete their homework," in order to block applicability to Max. Here, Jack's modification of the rule satisfies both of the two Raz/Simpson conditions: first, it simply narrows Jo's original rule with a further requirement for applicability, and second, it yields the same result as the original rule in the case in which it was introduced, that Emma can watch TV. Suppose, by contrast, Jack had modified Jo's original rule to read, "Children who are female can stay up and watch TV." Although this replacement would succeed in blocking applicability to Max, it violates the first of the two Raz/Simpson conditions: the new rule is not simply a narrowing of Jo's original rule but instead applies in some situations where the original rule would not – to a seven-year-old female child, for instance. Or suppose Jack had modified the original rule to

¹³ See Raz (1979, pp. 180–209) and Simpson (1961). Simpson himself was sharply aware of the problem described in the previous paragraph, describing it as a "paradox," which he presents as follows: "[t]he 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" (p. 172). His own approach to this problem was developed in the course of a dialogue with Goodhart and Montrose, going back to Goodhart's (1930) but vigorously pursued in the 1950s. See, in order: Montrose (1956), Montrose (1957a), Simpson (1957), Montrose (1957b), Simpson (1958), Goodhart (1959), and Simpson (1959).

read, “Children age nine or greater who have finished their dinner can stay up and watch TV.” The modification would again block applicability to Max, since he did not finish his dinner, but in this case it violates the second of the two conditions: it fails to justify the original outcome in the original case of Emma, since she did not finish her dinner either.

If we understand the standard model as including the Raz/Simpson conditions on rule modification, then a response to our initial problem concerning constraint is now available: even though later courts are free to modify the rules set out by earlier courts, they are nevertheless constrained by these rules, since they can modify them only in certain ways, those satisfying the Raz/Simpson conditions. This response to our initial problem, however, leads at once to another. Presumably, even if some modification of an earlier rule satisfies the Raz/Simpson conditions, a later court would, all the same, choose not to modify the rule in that way unless the court believed that it could actually improve the rule by doing so. But if a later court believes that it can improve an earlier rule through modification, why should it limit itself to modifications that satisfy the Raz/Simpson conditions? Why should the court not be free to modify the rule in any way at all that leads to an improvement, or if its modifications must be subject to conditions, then why these conditions and not others – in short: what is the justification for this particular set of conditions, the Raz/Simpson conditions, on rule modification?

This book offers an answer to that question, but the answer is arrived at in a roundabout way. My principal aim is to present an entirely different model of constraint – the *reason model* – according to which the mechanism of constraint does not depend on rules at all, whether defeasible or strict, but is instead defined in terms of reasons. The approach derives from an earlier proposal by Grant Lamond and is developed here using ideas from the field of artificial intelligence and law.¹⁴ On this view, what is important about an

¹⁴ See Lamond (2005) for his initial proposal and Lamond (2022) for more recent reflections. The first version of the account presented here occurs in Horty (2011), later developed in Horty (2015, 2016, 2017, 2019, 2021). This account has been related to research in artificial intelligence and law in Horty and Bench-Capon (2012), compared to arguments from analogy and enriched in various ways in Rigoni (2014, 2015), limited in scope in Broughton (2019), explored from a formal perspective in Prakken (2021), and compared to Lamond’s original proposal in Horty (2022). More recently, a different interpretation of Lamond’s proposal, and one that connects it more closely with traditional ideas from legal theory, has been presented in Mullins (2020); interestingly, Mullins shows that his alternative interpretation is, in a precise sense, equivalent to that presented here.

earlier decision is the earlier court's assessment of the importance among the various reasons presented by that case, an assessment that is expressed here as a priority ordering among reasons. Later courts are then constrained, not to follow the rules set out in precedent cases, or even to modify those rules only in certain ways, but simply to reach decisions that are consistent with the priority ordering among reasons that has already been established in earlier decisions. The development of the common law is pictured, not as the elaboration over time of an increasingly complex system of rules, but instead as the gradual construction of an ordering relation among reasons, reflecting their importance, or priority.

A principal advantage of the reason model is that it shows how several of the existing models of precedential constraint can be unified, helping us to see what is correct in each, and how they are related. Because of the link between reasons and default rules, it is possible, first of all, as mentioned earlier, to reinterpret the reason model as a defeasible rule model of constraint – and then to generalize this interpretation, showing how the reason model fits into the framework of defeasible reasoning more broadly. Second, once the reason model of constraint has been defined, it can be shown – somewhat surprisingly – to supply the necessary support for the Raz/Simpson conditions on rule modification; indeed, as we will see, the standard model and the reason model are, in a straightforward sense, equivalent.

A different conception of precedential constraint is provided by the *a fortiori* model, according to which – regardless of any rules or opinions contained in earlier decisions – a later court is constrained at all only if it is facing a situation at least as strong for the winning side of some precedent case as the fact situation of that precedent case itself.¹⁵ Because it disregards case rules entirely, the *a fortiori* model can be seen as advancing a radical notion of precedential constraint. Nevertheless, we will see that, at least on a charitable interpretation, this conception of constraint can be understood as a special case of the more general reason model.

¹⁵ Alexander (1989) characterizes the model of constraint based on a *fortiori* reasoning as the “result model,” and groups it together with the standard model. Although I adopted this “result model” terminology in Horty (2004), I now think it is better to label this particular conception as the “*a fortiori* model,” in order to make it clear that the standard model and the *a fortiori* model are distinct.

Finally, there is the *natural model* of constraint, according to which a court confronting a new situation is imagined to reach a decision by engaging in a process of ordinary, or natural, reasoning – based on the various reasons bearing on that situation, possibly including previous cases, and assigning to these reasons the priorities they seem to deserve.¹⁶ On this view, the decisions reached in previous cases are treated as nothing but ordinary events in the natural world, to be taken into account, just like any other ordinary events, in the course of natural reasoning, but not as special sources of law. Here too, it turns out that the natural model can be related to the reason model in a helpful way if, as I have advocated, we understand default logic as providing a precise account of natural reasoning. For we can then see exactly how the process of natural reasoning can itself be constrained – shaped, or modified – by precedent cases to yield results that accord with the reason model.

The book contains seven chapters and an appendix.

Chapter 1 introduces the reason model of constraint, setting out basic definitions and illustrating these definitions with a number of examples, including an application to statements formed from open-textured predicates. Chapter 2 then explores some of the concepts underlying the reason model in more detail.

Chapter 3 presents – for the first time, I believe – a precise formulation of the standard model of constraint. Based on this formulation, it is then established that the standard model and the reason model are equivalent. This chapter also presents a precise formulation of the *a fortiori* model of constraint, which allows us to see it as a special case of the reason model.

The equivalence between the standard model and the reason model of constraint, though perhaps surprising, is also reassuring, in the way that it is always reassuring when two different analyses of a concept, starting from different initial points and relying on different ideas, agree in outcome. Nevertheless, and in spite of this equivalence, Chapter 4 argues that the reason model is preferable to the standard model, for three reasons: First, the reason model supports a satisfying account of the process of reasoning under the constraints of precedent – described informally in this chapter, and characterized as the process of constrained natural reasoning. Second, the reason model allows us to draw a sharp and principled distinction between the two common law operations of distinguishing and overruling a previous

¹⁶ See Alexander (1989, Section 2) and Alexander and Sherwin (2008, Chapter 2), as well as the references cited there, for a more complete discussion of the natural model.

decision. And third, the reason model allows a deeper understanding of the real mechanism through which common law decisions constrain later courts.

These first four chapters, which presuppose nothing but an acquaintance with set-theoretic notation, provide a thorough introduction to the reason model of constraint. Taken together, they define the reason model, illustrate the model with applications, including an application to open-textured predicates, define the standard and a *fortiori* models, show that the first of these is equivalent to the reason model while the second is a special case, and finally, defend the reason model. For many readers, this may be all they wish to know about the reason model – and certainly, for some readers, it will be much more than they wish to know. The remaining three chapters suggest some directions in which the reason model can be developed. These chapters are more exploratory, more tentative, and involve formal material that is, although not actually more advanced, at least more intricate than that employed earlier.

Chapters 5 and 6 show how the notion of constrained natural reasoning, described informally in Chapter 4, can be characterized more precisely. Chapter 5 first reviews default logic as a formal theory of natural reasoning and extends this theory to an account of natural reasoning about the problems presented to courts.

Chapter 6 then shows how natural reasoning of this kind can be adapted to respect the constraints derived from precedent cases, resulting, at last, in a precise analysis of constrained natural reasoning. The primary goal of the chapter is to establish that natural reasoning adapted to respect hard, or absolute, constraints from a background case base leads to decisions that satisfy the reason model – this result, as mentioned earlier, can be interpreted as a demonstration that the reason model can be reformulated as a defeasible rule model of constraint. But the chapter has other goals as well: It explores the more complex topic of reasoning against the background of, not only hard constraints but also constraints that are softer – persuasive, but not absolute – and offers, within this setting, an account of overruling a precedent decision. It illustrates the way in which the practice of precedent merges the values of individual courts into a shared priority ordering on reasons. And it shows how the tools of constrained natural reasoning can be employed, once again, to help us understand the phenomenon of open texture.

Finally, Chapter 7 broadens the scope of the reason model to a richer representational setting, where situations are described, not in terms of

statements that either do or do not hold but in terms of characteristics that may be present, or not, to a particular degree. The original definition of the reason model seems, at first, to generalize very naturally to this setting but then leads to a new wrinkle, which shows that, in this richer setting, the priority ordering among reasons must itself be refined. Chapter 7 is independent of the earlier Chapters 5 and 6 and can reasonably be read without them.

The Appendix provides, first, a list of fact situations and cases used as examples throughout the book, and then a verification of the various observations noted in the text.