

Reasoning with Precedents as Constrained Natural Reasoning

John Horty

www.umiacs.umd.edu/users/horty

Version of: August 14, 2014

Contents

1	Introduction	1
2	A prioritized default logic	3
3	Reasoning about priorities	10
4	Precedent	13
5	Constrained natural reasoning	19
6	Conclusion	26

1 Introduction

One of the things that makes reasoning with precedent in the common law so difficult to understand is that it seems to slip between two familiar models of decision making, with familiar advantages and disadvantages. According to the first model, a court reasoning with precedent would be required simply to follow existing rules, exactly as they have been laid down by earlier courts, unless it happens to confront a situation to which no previous rule applies, in which case it can formulate a new rule to justify its decision. The central advantage of this model is predictability; its central disadvantage is rigidity. According to the second model, a court could be thought of as reaching its decision in the natural way, just as most of us reach decisions about most things—by surveying the reasons that bear on the situation at hand, assigning these reasons the weights they seem to deserve, and then reaching whatever results these reasons, together with their assigned weights, seem to support. The central advantage of this model is the flexibility it allows courts to adapt their reasoning to particular circumstances; its central disadvantage is a lack of predictability.

In the face of these two familiar models of decision making, the literature contains two broad reactions to the suggestion that reasoning with precedent in the common law occupies, more properly, a middle position between them—allowing more freedom than reasoning driven entirely by rules, but demanding more predictability than natural reasoning alone. Some writers argue that there is not, in fact, any defensible middle position at all, so that we are forced to assimilate common law reasoning to one of the extreme models or the other, with no further options.¹ Others try to define an appropriate middle position, involving both rules and reasons, and to argue that it is defensible.²

¹See, for example, Alexander (1989) and, especially, Alexander and Sherwin (2001) and (2009) for a powerful argument that reasoning with precedent should be assimilated to rule-based reasoning.

²An example is the “presumptive positivism” developed by Schauer in (1989, p. 117n, pp. 196–206) and

This paper presents a third option. I do not an attempt to locate common law reasoning somewhere in the territory between rule-based and natural reasoning. Instead, I treat decision making in the common law as entirely based on reasons, just as in the natural model, but with the sole difference that a common law decision maker, driven by the requirements of precedent, must adapt his or her own weighting, or priority ordering, on reasons so that it coheres with a priority ordering derived from a background case base. More exactly, I suggest that such a decision maker engages in a form of *constrained natural reasoning*, in which precedent cases are treated as reasons for altering the weights, or priorities, that would normally be assigned to other reasons.

Any precise model of constrained natural reasoning would have to be built on top of a precise model of natural reasoning, so that we can see exactly what is being constrained and how the relevant constraints are supposed to operate. I will rely here on my own proposal—that natural reasoning can usefully be analyzed in terms of default logic, with reasons themselves treated as components of default rules, interacting as specified by the logic to support whatever conclusions they do.³ Section 2 of this paper provides an overview of the basic prioritized default logic underlying this proposal, where the priorities among default rules represent the weight of reasons. Section 3 then sketches one elaboration of the basic theory, in which the reasoner can be thought of as reasoning, also by default, about the priorities that guide his or her own default reasoning, and so evaluating reasons for assigning weights to other reasons.

Section 4 of the paper reviews a factor-based representation of legal cases, derived from

³This proposal is developed in my (2012); it is based on the logic for default reasoning originally presented by Reiter (1980). A different, though related, approach can be found in the work of Pollock; see especially Pollock (1995).

research in the field of artificial intelligence and law, as well as my own suggestion that earlier cases can themselves be thought of as determining a priority ordering on legal reasons, and that precedent then requires later courts only to reach a decision that is consistent with this ordering.⁴

Finally, Section 5 ties the two earlier parts of the paper together. Taking default logic as a model of ordinary natural reasoning, I show in this section how the constraints imposed by precedent can be coded into a default theory—with defaults representing the court’s own reasons, further defaults representing the reasons provided by previous cases, and still further defaults representing the requirement that the reasons provided by previous cases are to be assigned a higher priority than the court’s own reasons. It then turns out that a court reasoning on the basis of a default theory like this—engaging, that is, in constrained natural reasoning—will arrive at a decision satisfying the requirements of precedent.

The substantive goal of this exercise, once again, is to sketch a new model of reasoning with precedent—not as reasoning with rules, nor as natural reasoning, nor as something in between, but as a form of natural reasoning in which the weights, or priorities, that would normally be assigned to reasons may be altered in accord with reasons provided by the existing case base. There is also, however, a methodological goal, especially relevant to readers of this volume, and that is to demonstrate the kind of work that can be carried out on the basis of a detailed theory of reasons, their weights or priorities, and their interactions.

2 A prioritized default logic

We begin with a brief overview of a simple prioritized default logic, taking as background an ordinary logical system in which \wedge , \vee , \supset , and \neg are the operations of conjunction,

⁴This proposal is developed in my (2011); see also Lamond (2005) for a related, and earlier, suggestion.

disjunction, implication, and negation, and in which \top is the trivially true proposition. The turnstile \vdash indicates ordinary logical consequence, so that $\mathcal{E} \vdash X$ means that the proposition X follows from the set of propositions \mathcal{E} .

Where X and Y are propositions, we take $X \rightarrow Y$ as the *default rule* that allows us to conclude Y , by default, once X has been established. To illustrate: if we suppose that B is the proposition that Tweety is a bird and F the proposition that Tweety can fly, then $B \rightarrow F$ is the rule that allows us to conclude that Tweety can fly, by default, once it has been established that Tweety is a bird. We assume two functions—*Premise* and *Conclusion*—that pick out the premises and conclusions of default rules: if r is the default $X \rightarrow Y$, then $Premise(r)$ is the proposition X and $Conclusion(r)$ is the proposition Y . The second of these functions is lifted from individual defaults to sets of defaults in the obvious way: where \mathcal{S} is a set of defaults, $Conclusion(\mathcal{S}) = \{Conclusion(r) : r \in \mathcal{S}\}$ is the set of conclusions of those defaults belonging to \mathcal{S} .

Default rules are to be thought of as expressing the *reason relation*. In the case of our example, then, what the default $B \rightarrow F$ indicates is that the premise that Tweety is a bird functions as a reason for the conclusion that Tweety flies.⁵

Some defaults, as well as their corresponding reasons, have greater weight, or higher priority, than others.⁶ This information is represented through an ordering relation $<$ on default rules, where the statement $r < r'$ means that the default r' has a higher priority than

⁵I argue in my (2012) that reasons are provided, not by defaults in general, but only by defaults that are *triggered*—a concept that will be defined shortly. This refinement is not necessary for the current discussion, and I will ignore it.

⁶Though I will use both terms, I prefer to speak in terms of a priority ordering on reasons, rather than of weight, for two reasons: first, I allow for nonlinear priority orderings, while the concept of weight tends to suggest linearity; second, I allow for the possibility that there may be only an ordinal ranking among reasons, while the concept of weight suggests that cardinal comparisons must also be available.

the default r . Suppose, for example, that P is the proposition that Tweety is a penguin, so that $P \rightarrow \neg F$ is the default allowing us to conclude that Tweety cannot fly once it is established that Tweety is a penguin. Then if we take r_1 as the earlier default $B \rightarrow F$ and r_2 as this new default, it is natural to assume that $r_1 < r_2$.

We will focus to begin with on *fixed priority* default theories—theories, that is, in which all priorities among default rules are fixed in advance. Such a theory is a structure of the form $\Delta = \langle \mathcal{W}, \mathcal{D}, < \rangle$, in which \mathcal{W} is a set of ordinary statements, \mathcal{D} is a set of default rules, and $<$ is a strict partial ordering on \mathcal{D} , representing priority.

Defaults are often thought of as special rules of inference that can be used to extend the conclusions derivable from a body of hard information beyond its ordinary logical consequences, and for this reason, the conclusion sets supported by default theories are generally referred to as *extensions*. We will concentrate here, however, not on extensions themselves, but on the sets of defaults through which they can be generated. To begin with, then, let us define a *scenario* based on a default theory $\Delta = \langle \mathcal{W}, \mathcal{D}, < \rangle$ simply as some subset \mathcal{S} of the set \mathcal{D} of defaults contained in that theory. From an intuitive standpoint, a scenario is supposed to represent the particular subset of default rules that have actually been selected by the reasoning agent as providing sufficient support for their conclusions—those rules that can then be applied to the hard information from \mathcal{W} to generate an extension. Not every scenario based on a default theory is intuitively acceptable, of course; some might contain what seems to be the wrong selection of defaults. The goal, therefore, is to characterize the *proper scenarios*—those sets of defaults that could be accepted by an ideal reasoning agent based on the information contained in the original theory.

These ideas can be illustrated by returning to our example, with r_1 and r_2 as the defaults $B \rightarrow F$ and $P \rightarrow \neg F$. If we suppose that Tweety is both a bird and a penguin, then

the information from this example can be captured by the default theory $\Delta_1 = \langle \mathcal{W}, \mathcal{D}, < \rangle$, where $\mathcal{W} = \{P, P \supset B\}$, where $\mathcal{D} = \{r_1, r_2\}$, and where $r_1 < r_2$. The set \mathcal{W} contains the basic information that Tweety is a penguin, and that this entails the fact that he is a bird; the set \mathcal{D} contains the two defaults; and the ordering tells us that the default about penguins has higher priority than the default about birds. This theory allows four possible scenarios— $\mathcal{S}_1 = \emptyset$, $\mathcal{S}_2 = \{r_1\}$, $\mathcal{S}_3 = \{r_2\}$, or $\mathcal{S}_4 = \{r_1, r_2\}$ —corresponding to the situations in which the reasoning agent endorses neither of the two available defaults, only the first default, only the second, or both. From an intuitive standpoint, though, it seems that the agent should endorse the default r_2 , and only that default, leading to the conclusion that Tweety does not fly. Therefore, only the third of these four scenarios, $\mathcal{S}_3 = \{r_2\}$, should be classified as proper.

How, then, can we define the proper scenarios? The definition I offer depends on three initial concepts—triggering, conflict, and defeat.

The *triggered* defaults represent those that are applicable in the context of a particular scenario; they are defined as the defaults whose premises are entailed by that scenario—those defaults, that is, whose premises follow from the hard information belonging to the default theory together with the conclusions of the defaults already endorsed. More exactly, if \mathcal{S} is a scenario based on the theory $\Delta = \langle \mathcal{W}, \mathcal{D}, < \rangle$, the defaults triggered in this scenario are those belonging to the set

$$Triggered_{\mathcal{W}, \mathcal{D}}(\mathcal{S}) = \{r \in \mathcal{D} : \mathcal{W} \cup Conclusion(\mathcal{S}) \vdash Premise(r)\}.$$

To illustrate by returning to the Tweety example, suppose $\mathcal{S}_1 = \emptyset$. In the context of this scenario, both r_1 and r_2 are triggered, since $\mathcal{W} \cup Conclusion(\mathcal{S}_1) \vdash Premise(r_1)$ and $\mathcal{W} \cup Conclusion(\mathcal{S}_1) \vdash Premise(r_2)$.

A default will be classified as *conflicted* in the context of a scenario if the agent is already

committed to the negation of its conclusion—that is, the conflicted defaults in the context of the scenario \mathcal{S} , based on the theory $\Delta = \langle \mathcal{W}, \mathcal{D}, < \rangle$, are those belonging to the set

$$Conflicted_{\mathcal{W}, \mathcal{D}}(\mathcal{S}) = \{r \in \mathcal{D} : \mathcal{W} \cup Conclusion(\mathcal{S}) \vdash \neg Conclusion(r)\}.$$

This idea can be illustrated through another example. Suppose that Q , R , and P are the statements that Nixon is a Quaker, that Nixon is a Republican, and that Nixon is a pacifist; and let r_1 and r_2 be the defaults $Q \rightarrow P$ and $R \rightarrow \neg P$, instances for Nixon of the generalizations that Quakers tend to be pacifists and that Republicans tend not to be pacifists. Then, since Nixon was, in fact, both a Quaker and a Republican, we can represent an agent's information through the theory $\Delta_2 = \langle \mathcal{W}, \mathcal{D}, < \rangle$, where $\mathcal{W} = \{Q, R\}$, where $\mathcal{D} = \{r_1, r_2\}$, and where $<$ is empty, since neither default has a higher priority than the other. Now imagine that, on whatever grounds, the agent decides to endorse one of these two defaults—say r_1 , supporting the conclusion P —and is therefore reasoning in the context of the scenario $\mathcal{S}_1 = \{r_1\}$. In this context, the other default— r_2 , supporting the conclusion $\neg P$ —will be conflicted, since $\mathcal{W} \cup Conclusion(\mathcal{S}_1) \vdash \neg Conclusion(r_2)$.

Finally, defeat. Although this concept is surprisingly difficult to define in full generality, the basic idea is simple enough, and can serve as the basis of a preliminary definition.⁷ Very roughly, a default will be classified as *defeated* in the context of a scenario \mathcal{S} , based on the theory $\Delta = \langle \mathcal{W}, \mathcal{D}, < \rangle$, whenever there is a stronger triggered default that supports a conflicting conclusion—whenever, that is, the default belongs to the set

$$\begin{aligned} Defeated_{\mathcal{W}, \mathcal{D}, <}(\mathcal{S}) &= \{r \in \mathcal{D} : \text{there is a default } r' \in Triggered_{\mathcal{W}, \mathcal{D}}(\mathcal{S}) \text{ such that} \\ &\quad (1) \ r < r', \\ &\quad (2) \ \mathcal{W} \cup \{Conclusion(r')\} \vdash \neg Conclusion(r)\}. \end{aligned}$$

⁷See Chapter 8 of my (2012) for a discussion of the difficulties involved in arriving at a general definition.

This idea can be illustrated by returning to the Tweety example, the theory $\Delta_1 = \langle \mathcal{W}, \mathcal{D}, < \rangle$, where $\mathcal{W} = \{P, P \supset B\}$, where $\mathcal{D} = \{r_1, r_2\}$ with r_1 and r_2 as the defaults $B \rightarrow F$ and $P \rightarrow \neg F$, and where $r_1 < r_2$. We can suppose once again that the agent has not yet endorsed either of the two defaults, so that the initial scenario is $\mathcal{S}_1 = \emptyset$, and neither default is conflicted. Still, the default r_1 is defeated, since r_2 is triggered, and we have both (1) $r_1 < r_2$ and (2) $\mathcal{W} \cup \{Conclusion(r_2)\} \vdash \neg Conclusion(r_1)$.

Once the underlying notions of triggering, conflict, and defeat are in place, we can define the notion of a default that is *binding* in the context of the scenario \mathcal{S} , based on the theory $\Delta = \langle \mathcal{W}, \mathcal{D}, < \rangle$, as one that is triggered in this context, but neither conflicted nor defeated—as a default, that is, belonging to the set

$$\begin{aligned} Binding_{\mathcal{W}, \mathcal{D}, <}(\mathcal{S}) = \{ & r \in \mathcal{D} : \quad r \in Triggered_{\mathcal{W}, \mathcal{D}}(\mathcal{S}), \\ & r \notin Conflicted_{\mathcal{W}, \mathcal{D}}(\mathcal{S}), \\ & r \notin Defeated_{\mathcal{W}, \mathcal{D}, <}(\mathcal{S}) \}. \end{aligned}$$

The concept can again be illustrated with Tweety, under the assumption that the agent's scenario is $\mathcal{S}_1 = \emptyset$. Here, the default r_1 , supporting the conclusion F , is triggered in the context of this scenario, and it is not conflicted, but as we have just seen, it is defeated by the default r_2 ; and so it is not binding. By contrast, the default r_2 , supporting the conclusion $\neg F$, is likewise triggered, not conflicted, and in this case not defeated either. This default is, therefore, binding.

We can now turn, at last, to the notion of a proper scenario. There are again some complexities involved in the full definition of this idea, though these need not concern us here.⁸ We can therefore work with a preliminary definition of a proper scenario, based on a theory $\Delta = \langle \mathcal{W}, \mathcal{D}, < \rangle$, as one containing all and only the defaults that are binding in

⁸See Appendix A.1 of my (2012) for the full definition.

the context of that very scenario—a definition, that is, according to which a scenario \mathcal{S} is classified as proper just in case

$$\mathcal{S} = \textit{Binding}_{\mathcal{W}, \mathcal{D}, <}(\mathcal{S}).$$

We can think of the defaults from a proper scenario as presenting, not just reasons, but good reasons, in the context of that scenario. An agent who has accepted a set of defaults forming a proper scenario, then, is in an enviable position. Such an agent has already accepted all and only those defaults that it recognizes as presenting good reasons; the agent, therefore, has no incentive either to abandon any of the defaults already accepted, or to accept any others.

To illustrate, we return one last time to our examples. The first, Tweety, was the theory $\Delta_1 = \langle \mathcal{W}, \mathcal{D}, < \rangle$, where $\mathcal{W} = \{P, P \supset B\}$, where $\mathcal{D} = \{r_1, r_2\}$ with r_1 as $B \rightarrow F$ and r_2 as $P \rightarrow \neg F$, and where $r_1 < r_2$. We noted earlier that, of the four possible scenarios based on this theory—that is, $\mathcal{S}_1 = \emptyset$, $\mathcal{S}_2 = \{r_1\}$, $\mathcal{S}_3 = \{r_2\}$, and $\mathcal{S}_4 = \{r_1, r_2\}$ —only the third seemed attractive from an intuitive point of view; and with our definitions in place, the reader can now verify that only \mathcal{S}_3 is proper. Our second example, Nixon, was the theory $\Delta_2 = \langle \mathcal{W}, \mathcal{D}, < \rangle$, where $\mathcal{W} = \{Q, R\}$, where $\mathcal{D} = \{r_1, r_2\}$ with r_1 as $Q \rightarrow P$ and r_2 as $R \rightarrow \neg P$, and where $<$ is empty. This theory allows two proper scenarios, both $\mathcal{S}_1 = \{r_1\}$ and $\mathcal{S}_2 = \{r_2\}$. There are a number of ways of interpreting situations like this, in which default theories allow multiple proper scenarios; one is to suppose that our standards of rationality certify each of these scenarios as acceptable outcomes, without favoring one over the other.⁹

⁹See Chapters 1 and 7 my (2012) for a discussion of different ways of interpreting default theories allowing multiple proper scenarios.

3 Reasoning about priorities

We have concentrated thus far on fixed priority default theories, in which priority relations among default rules are fixed in advance. In fact, however, some of the most important things we reason about, and reason about by default, are the priorities among the very default rules that guide our default reasoning—we offer reasons for taking some of our reasons more seriously than others. Although this process may sound complicated, it turns out that the basic theory can be extended to account for this kind of reasoning in four simple steps.

The first step is to enrich our background language with the resources to enable formal reasoning about priorities among defaults: a new set of individual constants, to be interpreted as names of defaults, together with a relation symbol representing priority. For the sake of simplicity, we will assume that each of these new constants has the form n_X , for some subscript X , and that each such constant refers to the default r_X , or in schematic contexts, that the constants n, n', n'' and so on refer to the default rules r, r', r'' and so on. And we will assume also that our language now contains the relation symbol \prec , representing priority among defaults.

To illustrate, suppose that r_1 and r_2 are the defaults $A \rightarrow B$ and $C \rightarrow \neg B$, respectively, and that r_3 is the priority default $D \rightarrow n_1 \prec n_2$. Then since n_1 refers to r_1 and n_2 refers to r_2 , what r_3 says is that D functions as a reason for assigning r_2 a higher priority than r_1 . As a result, we would expect that, when all three of these defaults are triggered—that is, when A , C , and D all hold—the default r_1 will generally be defeated by r_2 , since the two defaults have conflicting conclusions. Of course, since r_3 is itself a default, the information it provides concerning the priority between r_1 and r_2 is defeasible, and could likewise be defeated.

The second step is to shift our attention from theories of the form $\langle \mathcal{W}, \mathcal{D}, \prec \rangle$ —that is, from fixed priority default theories—to theories containing a set \mathcal{W} of ordinary propositions

as well as a set \mathcal{D} of defaults, but no priority relation on the defaults that is fixed in advance. Instead, both \mathcal{W} and \mathcal{D} may contain initial information concerning priority relations among defaults, and then conclusions about these priorities, like any other conclusions, are arrived at through default reasoning. Because conclusions about the priorities among defaults might themselves vary depending on other conclusions drawn by the reasoning agent, theories like this, of the form $\Delta = \langle \mathcal{W}, \mathcal{D} \rangle$, are known as *variable priority* default theories; it is stipulated as part of the definition that the set \mathcal{W} of hard information in such a theory must contain each instance of the *asymmetry* schema $(n < n') \supset \neg(n' < n)$ in which the variables are replaced with names of the defaults from \mathcal{D} .¹⁰

Now suppose the agent accepts some particular scenario \mathcal{S} based on a variable priority theory; the third step, then, is to lift the priority ordering implicit in the agent's scenario to an explicit ordering that can be used in default reasoning. This is done in the simplest possible way, through the introduction of a *derived* priority ordering $<_{\mathcal{S}}$, defined as follows:

$$r <_{\mathcal{S}} r' \text{ just in case } \mathcal{W} \cup \text{Conclusion}(\mathcal{S}) \vdash n \prec n'.$$

The statement $r <_{\mathcal{S}} r'$ is taken to mean that r' has a higher priority than r according to the scenario \mathcal{S} . The force of the definition, then, is that this relation holds just in case $n \prec n'$ can be derived from the conclusions of the defaults belonging to \mathcal{S} , taken together with the hard information from \mathcal{W} . Because \mathcal{W} contains all instances of asymmetry, the derived priority relation $<_{\mathcal{S}}$ is guaranteed to be asymmetric.

The fourth and final step is to define the notion of a proper scenario for a variable priority default theory. This is accomplished by leveraging our previous definition of proper

¹⁰It is more common to stipulate that \mathcal{W} should contain instances of transitivity and irreflexivity, rather than asymmetry. We adopt the weaker assumption of asymmetry alone in this paper because the ordering relation on reasons derived from case bases, to be defined shortly, and which I propose to link to default theories, is not transitive; see my (2011, pp. 15-17) for a discussion.

scenarios for fixed priority theories of the form $\langle \mathcal{W}, \mathcal{D}, < \rangle$, where $<$ can be any strict partial ordering whatsoever. Using this previous definition, we can now stipulate that \mathcal{S} is a *proper scenario* for the variable priority theory $\Delta = \langle \mathcal{W}, \mathcal{D} \rangle$ just in case \mathcal{S} is a proper scenario, in the previous sense, for the particular fixed priority theory $\langle \mathcal{W}, \mathcal{D}, <_{\mathcal{S}} \rangle$, where \mathcal{W} and \mathcal{D} are carried over from the variable priority theory Δ , and where $<_{\mathcal{S}}$ is the priority relation derived from the scenario \mathcal{S} itself. The intuitive picture is this. In searching for a proper scenario, the agent arrives at some scenario \mathcal{S} , which then entails conclusions about various aspects of the world, including priority relations among the agent's own defaults. If these derived priority relations can be used to justify the agent in accepting exactly the scenario \mathcal{S} that the agent began with, then that scenario is proper.

These various definitions can be illustrated through a variant of the previous Nixon example in which it is useful to adopt, not the epistemic perspective of a third party trying to decide whether or not Nixon is a pacifist, but instead, the practical perspective of a young Nixon trying to decide whether or not to become a pacifist. As before, we take r_1 and r_2 as the defaults $Q \rightarrow P$ and $R \rightarrow \neg P$, where P , Q , and R are the propositions that Nixon is a pacifist, a Quaker, and a Republican. Given our current perspective, these two defaults should now be interpreted as providing practical, rather than epistemic, reasons: r_1 corresponds to the fact that, as a Quaker, Nixon has reason to become a pacifist, and r_2 to the fact that, as a Republican, he has reason not to become a pacifist.

In light of his conflicting reasons, let us imagine that Nixon seeks advice, first, from an elder of his Friends Meeting, who tells him that his religious reason should be given more weight than his political reason, but second, from an official of the Republican Party, who tells him exactly the opposite. If we take A and B as the respective statements of the religious and political figures, then the advice of these two authorities can be represented

through the defaults r_3 and r_4 , where r_3 is $A \rightarrow n_2 \prec n_1$ and r_4 is $B \rightarrow n_1 \prec n_2$. Nixon is now faced with his initial conflicting reasons, as well as further conflicting reasons as to how that initial conflict should be resolved. Finally, though, let us suppose that he seeks further counsel, perhaps from his wife, Pat, who tells him that the advice of the religious figure is to be preferred to that of the party official. If we take C as Pat's statement, her advice can be represented through the default r_5 , where r_5 is $C \rightarrow n_4 \prec n_3$.

The variable priority default theory that provides the background for Nixon's reasoning is $\Delta_3 = \langle \mathcal{W}, \mathcal{D} \rangle$, where \mathcal{W} contains the propositions A , B , C , Q , and R —according to which Nixon is a Quaker and a Republican, and his religious and political advisors, as well as his wife, said what they did—and where \mathcal{D} now contains r_1 , r_2 , r_3 , r_4 , and r_5 .¹¹ As the reader can verify, this theory allows the unique proper scenario $\mathcal{S}_1 = \{r_1, r_3, r_5\}$, supporting the conclusions that $r_4 < r_3$, that $r_2 < r_1$, and that P . The scenario corresponds to the course of action in which Nixon is moved by Pat's advice to take the advice of the religious figure more seriously than that of the political official, and so accepts the religious figure's recommendation that his religious reason is to be preferred to his political reason, and therefore accepts pacifism on the basis of his religious reason.

4 Precedent

We have focused so far on default logic and its relation to the theory of reasons, including the way in which it allows for reasoning about the priorities among reasons. Let us turn now to the topic of precedent in the common law. I will quickly review my own proposal that the concept should be understood in terms of an ordering relation on legal reasons.

¹¹The set \mathcal{W} must also contain appropriate instances of the asymmetry schema, but since these can be generated automatically from the defaults contained in that theory, I will not mention them explicitly.

We start with a few basic concepts. Legal cases will be represented in terms of *factors*, where a factor is a legally significant fact or pattern of facts.¹² Cases in different areas of the law will be characterized by different sets of factors, of course. In the domain of trade secrets law, for example, where the factor-based analysis has been developed most extensively, a case will typically concern the issue of whether the defendant has gained an unfair competitive advantage over the plaintiff through the misappropriation of a trade secret; and here the factors involved might turn on, say, questions concerning whether the plaintiff took measures to protect the trade secret, whether a confidential relationship existed 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.

The analysis set out here relies on several simplifications, which would have to be relaxed in a more complete theory. We will assume, first, that all factors have polarities, favoring one side or the other; second, that the factor representation is separable, so that factors have whatever force they do regardless of the presence of other factors; and third, that the reasoning under consideration involves only a single step, proceeding at once from the factors present in a case to a decision for the plaintiff or defendant, rather than moving through a series of intermediate legal concepts. Finally, it must be emphasized 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.

Formally, then, let us begin by postulating a set F of legal factors. A *fact situation* X , of the sort presented in a legal case, can then be defined as some particular subset of these factors: $X \subseteq F$. We will let $F^\pi = \{f_1^\pi, \dots, f_n^\pi\}$ represent the set of factors favoring

¹²The analysis of legal cases in terms of factors, initially taken only as points along legally significant dimensions, was first introduced by Rissland and Ashley (1987); see Ashley (1990) for a canonical treatment.

the plaintiff and $F^\delta = \{f_1^\delta, \dots, f_m^\delta\}$ the set of factors favoring the defendant. Given our assumption that each factor favors one side of the other, we can suppose that the entire set of legal factors is exhausted by those favoring the plaintiff together with those favoring the defendant: $F = F^\pi \cup F^\delta$. Where s is a side—either π or δ —we let \bar{s} represent the opposite side: $\bar{\pi} = \delta$ and $\bar{\delta} = \pi$. And where X is a fact situation, we let X^s represent the factors from X that support the side s : $X^\pi = X \cap F^\pi$ and $X^\delta = X \cap F^\delta$.

A *precedent case* will be represented as a fact situation together with an outcome and a rule justifying that outcome. Such a case, then, can be defined as a triple of the form $c = \langle X, r, s \rangle$, where X is a fact situation containing the legal factors presented by the case, r is the rule of the case, and s is its outcome, a decision for a particular side.

We will suppose that the rule r contained in a precedent case is a special *factor default* rule of the form $Y \rightarrow s$, where Y is a *factor reason* favoring the side s —that is, some nonempty set of factors uniformly favoring s as an outcome, a nonempty subset of F^s . Factor reasons are to be interpreted conjunctively, so what a default of the form $Y \rightarrow s$ means is that the factors from Y , taken together, support a decision in favor of s . To illustrate: $\{f_1^\pi, f_2^\pi\}$ is a factor reason favoring π , and so $\{f_1^\pi, f_2^\pi\} \rightarrow \pi$ is a factor default, according to which the presence of f_1^π and f_2^π , taken together, support a decision for the plaintiff. Note that the set $\{f_1^\pi, f_1^\delta\}$ is not a factor reason, since the factors it contains do not uniformly favor one side, and that $\{f_1^\pi, f_2^\pi\} \rightarrow \delta$ is not a factor default, since the factor reason that forms its premise favors the side π , but its conclusion is δ .

We impose two coherence constraints relating the rule of a precedent case $c = \langle X, r, s \rangle$ to the facts of the case as well as its outcome. The first is that the rule of the case must actually apply to the fact situation: $Premise(r) \subseteq X$. The second is that the conclusion of that rule must match the outcome of the case itself: $Conclusion(r) = s$.

These various concepts can be illustrated through the concrete case $c_1 = \langle X_1, r_1, s_1 \rangle$, where $X_1 = \{f_1^\pi, f_2^\pi, f_1^\delta, f_2^\delta\}$, with two factors each favoring the plaintiff and the defendant, where r_1 is the rule $\{f_1^\pi\} \rightarrow \pi$, and where the outcome s_1 is π , a decision for the plaintiff. Evidently, the case satisfies our two coherence constraints: the rule of the case is applicable to the facts, and the conclusion of this rule matches the outcome of the case. This particular precedent, then, represents a case in which the court decided for the plaintiff by applying or introducing a rule according to which the presence of the factors f_1^π leads, by default, to a decision for the plaintiff.

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

To motivate my proposed notion of precedential constraint—according to which the concept is to be understood in terms of an ordering on reasons—let us return to the case $c_1 = \langle X_1, r_1, s_1 \rangle$ and ask what information is actually carried by this case; what is the court telling us with its decision? Well, two things, at least. First of all, by appealing to the rule r_1 as justification, the court is telling us that the reason for its decision—that is, $Premise(r_1)$, or $\{f_1^\pi\}$ —is sufficient to justify a decision in favor of the plaintiff. But second, with its decision for the plaintiff, the court is also telling us that this reason must be stronger than the strongest reason presented by the case in favor of the defendant.

To put this precisely, let us first stipulate that, if X and Y are factor reasons favoring the same side, then Y is at least as strong as X for that side—or alternatively, X is at least as weak as Y —whenever $X \subseteq Y$. Returning to our example, then, where $X_1 = \{f_1^\pi, f_2^\pi, f_1^\delta, f_2^\delta\}$, it is clear that the strongest reason present for the defendant is $X_1^\delta = \{f_1^\delta, f_2^\delta\}$, containing

all those factors from the original fact situation that favor the defendant. Since the court has decided for the plaintiff on the grounds of the reason $Premise(r_1)$, even in the face of the conflicting X_1^δ , it seems to follow, as a consequence of the court's decision, that the reason $Premise(r_1)$ for the plaintiff should be assigned a higher priority than the reason X_1^δ for the defendant—that is, that $\{f_1^\pi\}$ should be assigned a higher priority than $\{f_1^\delta, f_2^\delta\}$. If we let $<_{c_1}$ represent the priority relation on factor reasons that is derived from the particular case c_1 , then this consequence of the court's decision can be put more formally as the claim that $X_1^\delta <_{c_1} Premise(r_1)$, or equivalently, that $\{f_1^\delta, f_2^\delta\} <_{c_1} \{f_1^\pi\}$.

As far as the priority ordering goes, then, the earlier court is telling us at least that $X_1^\delta <_{c_1} Premise(r_1)$, but is it telling us anything else? Perhaps not explicitly, but implicitly, yes. For if the reason $Premise(r_1)$ for the plaintiff is preferred to the reason X_1^δ for the defendant, then surely any reason for the plaintiff that is at least as strong as $Premise(r_1)$ must likewise be preferred to X_1^δ , and just as surely, $Premise(r_1)$ must be preferred to any reason for the defendant that is at least as weak as X_1^δ . It therefore follows from the earlier court's decision in c_1 , not only that $X_1^\delta <_{c_1} Premise(r_1)$, but that $W <_{c_1} Z$ whenever $W \subseteq X_1^\delta$ and $Premise(r_1) \subseteq Z$. To illustrate: from the court's explicit decision that $\{f_1^\delta, f_2^\delta\} <_{c_1} \{f_1^\pi\}$, we can conclude also that $\{f_1^\delta\} <_{c_1} \{f_1^\pi, f_3^\pi\}$, for example.

Generalizing from this example, we can now define the *priority relation* $<_c$ derived from the single case $c = \langle X, r, s \rangle$ by stipulating, where W and Z are factor reasons, that

$$W <_c Z \text{ if and only if } W \subseteq X^{\overline{s}} \text{ and } Premise(r) \subseteq Z.$$

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

$W <_{\Gamma} Z$ if and only if $W <_c Z$ for some case c from Γ .

And we can then define a case base as *inconsistent* if it provides conflicting information about the priority relation among reasons—telling us, for some pair of reasons X and Y , that each has a higher priority than the other, or that

$$X <_{\Gamma} Y \text{ and } Y <_{\Gamma} X.$$

A case base can be defined as *consistent* as long as it is not inconsistent.

Given this notion of consistency, we can at last turn to the *requirement of precedent* itself. The intuition is simply that, in deciding a case, a court is required by precedent only to preserve the consistency of the background case base. Suppose, more exactly, that a court is confronted with a new fact situation X against the background of a consistent case base Γ . Then what precedent requires is that the court base its decision on some rule r leading to an outcome s such that the case base resulting from supplementing Γ with this new decision—that is, the case base $\Gamma \cup \{\langle X, r, s \rangle\}$ —remains consistent.

This idea can be illustrated by assuming as background the case base $\Gamma_1 = \{c_1\}$, containing only the previous 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 is $\{f_1^{\pi}\} \rightarrow \pi$, and where s_1 is π . Suppose that, against this background, the court confronts the fresh situation $X_2 = \{f_1^{\pi}, f_1^{\delta}, f_3^{\delta}\}$ and considers finding for the defendant on the basis of f_1^{δ} , leading to the decision $c_2 = \langle X_2, r_2, s_2 \rangle$, where X_2 is as above, where $r_2 = \{f_1^{\delta}\} \rightarrow \delta$, and where $s_2 = \delta$. This decision would then violate the requirement of precedent. Why? Because the new case c_2 would support the priority relation $\{f_1^{\pi}\} <_{c_2} \{f_1^{\delta}\}$, telling us that the reason $\{f_1^{\delta}, f_2^{\delta}\}$ for the defendant outweighs the reason $\{f_1^{\pi}\}$ for the plaintiff. But Γ_1 already contains the case c_1 , from which we can derive the priority relation $\{f_2^{\delta}\} <_{c_1} \{f_1^{\pi}\}$, telling us exactly the opposite. As a result, the augmented case base $\Gamma_1 \cup \{c_2\}$ would be inconsistent.

5 Constrained natural reasoning

Having seen how default logic can be interpreted as a model of natural reasoning, and also how the requirement of precedent can be defined in terms of a priority ordering on factor reasons, we are now in a position to understand how reasoning with precedent can be understood as constrained natural reasoning—how, more exactly, a problem presented to a court against the background of a case base can be coded as a default theory, in such a way that the defaults representing the reasons provided by precedent cases override those reflecting the court’s own values.

It is often thought, by those who like the picture of common law rules as defeasible, that courts introduce, with their decisions, defeasible rules according to which certain collections of factors favor one side or another by default, though of course, these rules are subject to later exceptions. But this picture does not sit well with the view of factors as having polarities, favoring one side or another. If the factor f_1^π favors the side π , then there is no need for any court to introduce the rule $\{f_1^\pi\} \rightarrow \pi$, for example. The rule tells us that the presence of f_1^π counts as a reason for π , but this is just what it means to say that f_1^π favors π —the rule is, in a sense, already built into the favoring relation.

Generalizing, then, let us postulate a set \mathcal{F} containing, for each factor reason Y favoring the side s , a factor default of the form $Y \rightarrow s$, expressing the fact that the presence of the factors from Y counts as a reason for s . For convenience, we will also define a weak strength ordering among factor default rules favoring the same side, mirroring our earlier strength ordering on the factor reasons that form their premises: where r and r' are factor default rules, then $r \leq r'$ —that is, r' is at least as strong as r —just in case $Premise(r) \subseteq Premise(r')$.

But if courts do not introduce new factor default rules with their decisions—if these rules are already present, implicit in the meaning of factors—then what effect do the decisions

of courts have? The simplest answer is that these decisions impose an ordering relation on factor default rules, telling us which of these rules are to be assigned higher priority than others. This idea has been explored by a number of writers, as well as in my own previous work, but it does not give us what we want.¹³ It provides, at best, an account of defeasible reasoning with legal information alone, not a theory of the way in which this legal information constrains an agent’s natural reasoning. There is no description of a reasoning agent’s own preferences, or values, or of the way in which precedent impacts the decisions that an agent would otherwise reach.

To fill this gap, we begin by introducing special defaults to represent the agent’s own values, as they are reflected in the agent’s preferences among factor reasons and their corresponding default rules. Where r and r' are factor default rules favoring opposing sides, then, let us define a *value default* as a rule of the form $\top \rightarrow n \prec n'$; the force of a default like this is that, according to the agent’s own values, the rule r' is assigned a higher priority than the rule r .¹⁴ For the sake of simplicity, we will ignore the reasoning, surely very complex, through which an agent arrives at his or her own ordering on factor defaults, as well as any specific reasons justifying components of this ordering; we will likewise suppose that there is no priority ordering among value defaults, unrealistically treating all of the agent’s values on a par.

In order to see how reasoning with precedent can be understood as constrained natural reasoning, it will be useful to consider, first, how a problem can be coded as a default theory

¹³See Prakken and Sartor (1998) for a wide-ranging and detailed development of the idea; my own treatment, based on the current account of precedent, is found in my (2011).

¹⁴I refer to these rules as “value defaults” only for terminological clarity, using this phrase in a technical sense, without intending to advance any philosophical claims about the relation between reasons and values. In the same way, when I speak of an agent’s values, also in a technical sense, I mean to refer only to that agent’s reasons for ranking one factor reason over another, and so for ranking one factor default over another.

without any constraints at all, reflecting only the reasoning agent's own values. Suppose, then, that a fact situation X is presented for adjudication to an agent whose values are represented by a set \mathcal{V} of value defaults. We can take $\mathcal{D}_{\mathcal{V}} = \mathcal{F} \cup \mathcal{V}$ as the set of defaults guiding the agent's reasoning—that is, the entire set \mathcal{F} of factor defaults, together with the set \mathcal{V} of value defaults, reflecting the agent's own reasons for ranking some factor defaults as more important than others. We let the hard information \mathcal{W}_X to which these defaults are applied include X —the facts of the situation at hand—together with all instances of the *weak transitivity* schemata

$$\begin{aligned}(n < n' \wedge n' \leq n'') &\supset n < n'', \\ (n \leq n' \wedge n' < n'') &\supset n < n'',\end{aligned}$$

in which the variables are replaced with names of defaults from \mathcal{F} . The point of these schemata is simply to guarantee that, whatever priority ordering on factor defaults the agent finally settles on, it must respect the weak strength ordering defined earlier on rules favoring the same side. The *problem presented by the fact situation X to an agent with values \mathcal{V}* can then be represented by the variable priority default theory $\Delta_{X,\mathcal{V}} = \langle \mathcal{W}_X, \mathcal{D}_{\mathcal{V}} \rangle$, where $\mathcal{D}_{\mathcal{V}} = \mathcal{F} \cup \mathcal{V}$ and \mathcal{W}_X includes X together with instances of weak transitivity for all defaults from \mathcal{F} .

To illustrate, suppose an agent is presented with the situation $X_2 = \{f_1^\pi, f_1^\delta, f_3^\delta\}$, considered earlier. Out of the entire set \mathcal{F} of factor defaults, only four are applicable to this situation:

$$\begin{aligned}r_1 &= \{f_1^\pi\} \rightarrow \pi, \\ r_2 &= \{f_1^\delta\} \rightarrow \delta, \\ r_3 &= \{f_3^\delta\} \rightarrow \delta, \\ r_4 &= \{f_1^\delta, f_3^\delta\} \rightarrow \delta.\end{aligned}$$

And let us imagine that the agent’s particular set \mathcal{V}_1 of value defaults contains only the single rule

$$r_5 = \top \rightarrow n_1 \prec n_2,$$

according to which the factor default r_2 is to be prioritized over r_1 ; intuitively, this represents the agent’s view that $\{f_1^\delta\}$ is a more important reason in favor of δ than $\{f_1^\pi\}$ is in favor of π . The problem presented by this fact situation to an agent with these values, then, is represented by the default theory $\Delta_{X_2, \mathcal{V}_1} = \langle \mathcal{W}_{X_2}, \mathcal{D}_{\mathcal{V}_1} \rangle$, where the set $\mathcal{D}_{\mathcal{V}_1}$ contains all factor defaults together with the single value default r_5 , and where \mathcal{W}_{X_2} contains the factors from X_2 together with appropriate instances of weak transitivity.

It is easy to verify that this default theory allows the single proper scenario $\mathcal{S}_1 = \{r_2, r_3, r_4, r_5\}$, supporting the conclusion that $r_1 < r_2$, in accord with the reasoning agent’s values, and that the situation X_2 should be decided for the defendant. Note that r_1 , supporting a decision for the plaintiff, is not allowed, since, according to the agent’s values, this rule is defeated by the stronger r_2 , supporting the opposite conclusion. Note also that the scenario contains three separate rules— r_2 , r_3 , and r_4 —supporting a decision for the defendant. My interpretation of this is that the agent is free to choose one of the broader rules r_2 or r_3 , or the narrower rule r_4 , to justify his or her decision; this choice is determined—if it is determined at all—by information that is not represented in the current framework.

Now, to show how natural reasoning—reasoning based on default theories of the kind just developed, reflecting only the agent’s own values—can be constrained by legal reasons, I introduce two new classes of default rules. The first is the class of *case defaults*, representing the reasons provided by previous cases. Where $c = \langle X, r, s \rangle$ is a case, with r as the rule supporting the winning side s , and with $r' = X^{\bar{s}} \rightarrow \bar{s}$ as the rule based on the strongest reason present for the opposite side, this case will be said to *generate* the case default $c \rightarrow n' \prec n$.

The meaning of this case default is simply that the case itself, or the fact that it was decided as it was, provides a reason for assigning a higher priority to r than to r' —a reason for assigning a higher priority to the rule for the winning side than to the rule based on the strongest reason present for the losing side.

Value defaults and case defaults can conflict, of course. Given two factor defaults r and r' favoring different sides, an agent's value default, reflecting the agent's own values, might rank r' above r , while a case default ranks r above r' . In a situation like that, which of the two conflicting factor defaults should the agent apply? The answer is that an agent reasoning under the constraints of precedent—a *court*, as we will say—is required to favor the case default over its own value default. This information is carried, finally, by a third class of default rules. Where r'' is a value default and r''' a case default, a *precedent default* can be defined as a rule of the form $\top \rightarrow n'' \prec n'''$, according to which the case default is to be assigned higher priority than the value default. It is worth noting where these new defaults fall in the hierarchy. Factor defaults provide reasons for deciding a case for some particular side, plaintiff or defendant. Both value defaults and case defaults provide reasons for assigning some factor defaults a higher priority than others—and so, for ranking some reasons for deciding a case for some particular side above others. Precedent defaults, finally, provide reasons for assigning a higher priority to case defaults than to value defaults—and so, for ranking some reasons for ranking reasons for deciding a case for some particular side above other such reasons.

Let us now consider how these various ideas can be pulled together into a default theory representing the problem presented by a fact situation for a court reasoning, not only on the basis of its own values, but under the constraints of precedent. We will take \mathcal{F} , once again, as the entire set of factor defaults, and \mathcal{V} as a set of value defaults representing the court's

own values, or priorities among factor defaults. Where Γ is a case base, we will let \mathcal{C}_Γ be the set containing each case default generated by a case from Γ , and we will let $\mathcal{P}_{\mathcal{V},\Gamma}$ be the set of precedent defaults ranking every case default from \mathcal{C}_Γ above every value default from \mathcal{V} . The entire set $\mathcal{D}_{\mathcal{V},\Gamma}$ of defaults guiding the agent's reasoning can be defined as the union of these factor, value, case, and precedent defaults—that is, $\mathcal{D}_{\mathcal{V},\Gamma} = \mathcal{F} \cup \mathcal{V} \cup \mathcal{C}_\Gamma \cup \mathcal{P}_{\mathcal{V},\Gamma}$. The hard information $\mathcal{W}_{X,\Gamma}$ to which these defaults are applied will include both X and Γ —the facts of the situation at hand, together with the set of cases providing the context within which this situation is considered—as well as appropriate instances of the weak transitivity schemata mentioned earlier. The *problem presented by the fact situation X to a court with values \mathcal{V} under precedential constraints derived from a case base Γ* can then be represented by the variable priority default theory $\Delta_{X,\mathcal{V},\Gamma} = \langle \mathcal{W}_{X,\Gamma}, \mathcal{D}_{\mathcal{V},\Gamma} \rangle$, where $\mathcal{D}_{\mathcal{V},\Gamma} = \mathcal{F} \cup \mathcal{V} \cup \mathcal{C}_\Gamma \cup \mathcal{P}_{\mathcal{V},\Gamma}$ and where $\mathcal{W}_{X,\Gamma}$ includes X and Γ together with instances of weak transitivity for all defaults from \mathcal{F} .

A court reasoning with a theory of this kind can legitimately be described as engaged in a process of constrained natural reasoning: the court is reasoning in the natural way, on the basis of its own values, except when those values conflict with decisions reached in previous cases, in which case the court must then defer to the previous decisions. It is not difficult to verify that a court reasoning with such a default theory will satisfy the requirement of precedent, defined in the previous section, in the sense that a decision based on any factor default belonging to any proper scenario based on the theory will be consistent with the background case base. Put more precisely, we have the following observation:

Where Γ is a consistent case base, X a fact situation, and \mathcal{V} a set of values, let $\Delta_{X,\mathcal{V},\Gamma}$ represent the decision problem presented by the fact situation X to a court with values \mathcal{V} under precedential constraints derived from a case base Γ .

Then a factor default rule r supporting the outcome s belongs to some proper scenario based on $\Delta_{X,\mathcal{V},\Gamma}$ just in case $\Gamma \cup \{\langle X, r, s \rangle\}$ is consistent.

This observation—that the results of constrained natural reasoning satisfy the requirements of precedent—can be illustrated by continuing with our earlier example, in which the situation $X_2 = \{f_1^\pi, f_1^\delta, f_3^\delta\}$ is presented to a court whose values are represented by the set $\mathcal{V}_1 = \{r_5\}$, where the value default r_5 , displayed earlier, provides a reason for prioritizing the factor default r_2 over r_1 . This time, however, we will suppose that the situation is evaluated under precedential constraints derived from the background case base $\Gamma_1 = \{c_1\}$, considered at the end of Section 4, containing the single 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 , again, is $\{f_1^\pi\} \rightarrow \pi$, and where s_1 is π . We then have $\mathcal{C}_{\Gamma_1} = \{r_6\}$ as the set containing the single case default derived from the single background case, where

$$r_6 = c_1 \rightarrow n_7 \prec n_1,$$

and r_7 is $\{f_1^\delta, f_2^\delta\} \rightarrow \delta$. We have $\mathcal{P}_{\mathcal{V}_1, \Gamma_1} = \{r_8\}$ as the set containing the single precedent default necessary to rank the single case default from \mathcal{C}_{Γ_1} above the single value default from \mathcal{V}_1 , where

$$r_8 = \top \rightarrow n_5 \prec n_6.$$

And we thus have $\mathcal{D}_{\mathcal{V}_1, \Gamma_1} = \mathcal{F} \cup \mathcal{V}_1 \cup \mathcal{C}_{\Gamma_1} \cup \mathcal{P}_{\mathcal{V}_1, \Gamma_1}$ as the entire set of factor, value, case, and precedent defaults guiding the court's reasoning. The hard information $\mathcal{W}_{X_2, \Gamma_1}$ to which these defaults are applied will include the set X_2 of case facts, the set Γ_1 of cases that form the background context, as well as all instances of weak transitivity for factor defaults. The resulting default theory $\Delta_{X_2, \mathcal{V}_1, \Gamma_1} = \langle \mathcal{W}_{X_2, \Gamma_1}, \mathcal{D}_{\mathcal{V}_1, \Gamma_1} \rangle$ therefore represents the problem presented to the court.

It is easy to see that the default theory $\Delta_{X_2, \nu_1, \Gamma_1}$ supports the two extensions $\mathcal{S}_2 = \{r_1, r_6, r_8\}$ and $\mathcal{S}_3 = \{r_3, r_4, r_6, r_8\}$. It follows from our observation, then, that a decision based on any of the factor defaults contained in either of these two extensions—that is, a decision based on any of r_1 , r_3 , or r_4 —is consistent with the background case base, and so satisfies the requirements of precedent in the sense defined earlier. The court’s reasoning is constrained in such a way that it is bound to reach a decision satisfying the requirements of precedent.

It is useful to compare the scenarios supported by the constrained default theory $\Delta_{X_2, \nu_1, \Gamma_1}$ with $\mathcal{S}_1 = \{r_2, r_3, r_4, r_5\}$, supported by the earlier theory Δ_{X_2, ν_1} , in which the court reasons in an unconstrained way, entirely on the basis of its own values, about the same situation. Both \mathcal{S}_2 and \mathcal{S}_3 , supported by the constrained theory, contain the precedent default r_8 , according to which the case default r_6 has higher priority than the value default r_5 . Both \mathcal{S}_2 and \mathcal{S}_3 , therefore, contain r_6 rather than r_5 , which is itself defeated by r_6 . It follows from r_6 that r_1 , according to which the plaintiff is favored on the basis of f_1^π , has a higher priority than r_2 , according to which the defendant is favored on the basis of f_1^δ . In contrast to \mathcal{S}_1 , therefore, neither \mathcal{S}_2 nor \mathcal{S}_3 can contain r_2 . The scenario \mathcal{S}_2 represents an outcome in which the court decides for the plaintiff on the basis of r_1 itself. The scenario \mathcal{S}_3 represents an outcome in which the court decides for the defendant, but there is no longer the option of justifying this decision on the basis of r_2 . Instead, a decision for the plaintiff must now be based on either r_3 or r_4 .

6 Conclusion

My goal in this paper has been to set out a new proposal according to which reasoning with precedent is treated as constrained natural reasoning—like natural reasoning, except that

the weights, or priorities that a reasoning agent would normally assign to reasons might have to be adjusted in accord with reasons provided by a background case base. The proposal is developed against the background of a model of natural reasoning as default reasoning, with reasons explicated in terms of default rules.

Since the details of the proposal are somewhat complex, a review may be helpful: Cases before a court are represented as sets of factors, each favoring one side or another. Factor reasons are defined as collections of reasons uniformly favoring one particular side, either the plaintiff or the defendant, with the force of factor reasons captured through factor default rules. Each case before a court, then, presents a number of conflicting factor reasons, triggering conflicting factor defaults. Different individuals might assign different weights to these factor reasons, of course, and so to their corresponding factor defaults. The weights that an individual assigns to factor reasons are represented through that individual's value defaults, which define a priority ordering of factor defaults. An individual deciding a case entirely on the basis of his or her own ordering among factor defaults, reflecting only his or her own values, is engaging in purely natural reasoning.

In addition to an individual's own values, there are also, however, other reasons for assigning greater weight to some factor reasons than to others, and so for prioritizing some factor defaults over others. In the modern common law, previously decided cases provide further reasons for favoring certain factor defaults over others; the force of these reasons is represented through case defaults. These case defaults may provide reasons conflicting with those provided by an individual's own values—an individual might disagree with the decisions reached in some previous case. But according to the proposal set out here, considerations of precedent can be thought of as providing individuals with further, higher-level reasons for favoring reasons derived from earlier cases over those deriving from their own values; these

reasons are captured by precedent defaults, prioritizing case defaults over value defaults.

The central result of the paper is that an individual taking these higher-level reasons into account but otherwise engaging in natural reasoning will be bound, in any particular case, to reach a decision satisfying the requirements of precedent. The complexity of the paper results from the need to define the necessary concepts carefully enough that this result can be stated clearly, and verified.

The approach sketched here suggests a number of further problems, of which I will mention only two. First, I have simplified by ignoring the reasoning through which an individual's values support an ordering on factor defaults, collapsing all of this reasoning into a single step. But of course, this simplification would have to be relaxed in developing a richer theory, and when it is, we would then have to consider more sophisticated ways in which accommodating the case defaults derived from a background case base might interfere with the reasoning through which an individual arrives at his or her own ordering on reasons. Second, I have concentrated on the case in which individuals, in their reasoning, are constrained by precedent, so that they are required to assign case defaults a higher priority than their own value defaults. A different, and perhaps more difficult, question is how we might adapt the present framework to model the reasoning of individuals for whom precedents are persuasive, but not authoritative.

These problems, and others, remain unresolved. But in spite of the problems, and in spite of the complexity of the present proposal, I hope that the ideas mapped out here help to show that the development of a precise theory of reasons and their weight might have applications in a number of different areas.

References

- Alexander, Larry (1989). Constrained by precedent. *Southern California Law Review*, 63:1–64.
- Alexander, Larry and Sherwin, Emily (2001). *The Rule of Rules: Morality, Rules, and the Dilemmas of Law*. Duke University Press.
- Alexander, Larry and Sherwin, Emily (2009). *Demystifying Legal Reasoning*. Cambridge University Press.
- Ashley, Kevin (1990). *Modeling Legal Argument: Reasoning with Cases and Hypotheticals*. The MIT Press.
- 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.
- Lamond, Grant (2005). Do precedents create rules? *Legal Theory*, 11:1–26.
- Pollock, John (1995). *Cognitive Carpentry: A Blueprint for How to Build a Person*. The MIT Press.
- Prakken, Henry and Sartor, Giovanni (1998). Modelling reasoning with precedents in a formal dialogue game. *Artificial Intelligence and Law*, 6:231–287.
- Reiter, Raymond (1980). A logic for default reasoning. *Artificial Intelligence*, 13:81–132.
- Rissland, Edwina and Ashley, Kevin (1987). A case-based system for trade secrets law. In *Proceedings of the First International Conference on AI and Law (ICAIL-87)*, pages 60–66. The Association for Computing Machinery Press.

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

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