Reasoning with Dimensions and Magnitudes

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

ABSTRACT

In response to problems raised by Bench-Capon [4], this paper shows how two models of precedential constraint can be broadened to include legal information represented through dimensions, as well as standard factors.

CCS CONCEPTS

• **Computing methodologies** → **Knowledge representation and reasoning**; Lexical semantics; • **Applied computing** → *Law*;

ACM Reference format:
John Horty. 2017. Reasoning with Dimensions and Magnitudes. In *Proceedings of International Conference on Artificial Intelligence and Law, London, UK, June 12–15, 2017 (ICAIL'17),* 10 pages.
DOI: http://dx.doi.org/10.1145/3086512.XYZ

1 INTRODUCTION

The field of artificial intelligence and law is characterized by a tradition of reconciliation between case-based and rule-based, or logical, methods of reasoning. Tentative efforts in this direction can traced back to Loui et al. [16], but substantial progress was first achieved by Prakken and Sartor [18], who showed that many of the patterns of legal reasoning and legal argument first studied in the case-based framework of HYPO [3]—as well as in successor systems, such as CABARET [21] and CATO [1]—could also be modeled in the framework of a defeasible logic with variable priorities. This unification of ideas from case-based reasoning with ideas from the rule-based framework of defeasible logic is one of the great success stories of the entire field.

But problems remain—or at least barriers to reconciliation—deriving from the differences, both in case-based knowledge representation and in the reasoning it supports, between "factors" and "dimensions." Usage in the field is not entirely uniform, but let us say, for present purposes, that a *factor* is a legally significant proposition, which may or may not hold in a given situation, but which, when it does hold, always favors the same side in a dispute. A *dimension*, by contrast, is an ordered set of legally significant *values*, where the ordering among values reflects the extent to which the fact that the dimension takes on that particular value favors one side or the other.

The contrast between factors and dimensions can be illustrated with examples from the field of trade secrets law, the original application domain of the HYPO and CATO systems, which explores the conditions under which a defendant can be said to have gained an unfair competitive advantage over a plaintiff through the misappropriation of a trade secret. One relevant consideration is whether the defendant has, or has not, signed a non-disclosure agreement. This consideration can naturally be represent as a factor—say, the proposition that a non-disclosure agreement was signed—since this proposition either holds or does not, but always supports the plaintiff if it does hold, since it indicates the presence of a genuine secret. Another consideration concerns the extent to which the information alleged to be a secret has already been disclosed to outsiders. This consideration is best represented as a dimension, with the possible values along that dimension—the number of outsiders to whom the information was disclosed—arranged in such a way that disclosure to more outsiders progressively strengthens the case for the defendant, since it provides stronger support for the idea that the information in question was not in fact a secret.

A third consideration concerns measures taken to protect the information purported to be a secret—again best represented as a dimension, with protective measures as values, and these values ordered in such a way that stronger measures provide stronger support for the plaintiff's claim that the information was indeed a secret. Imagine that the information in question is data stored on a disk, and consider four possible values along the protective measures dimension: (1) the plaintiff has taken no protective measures, (2) the plaintiff has encrypted the disk, (3) the plaintiff has locked the disk in a safe, (4) the plaintiff has both encrypted the disk and locked it in a safe. These values might naturally be ordered so that the second and third provide stronger support for the plaintiff than the first, but are incomparable to each other, and the fourth provides stronger support for the plaintiff than all the others.

The last example highlights three points about dimensions. It shows, first, that the values along a dimension need not correspond to a numerical range, but could be entirely qualitative, and second, that the ordering among these values need not be linear. Third, the example illustrates the fact that the polarity of some particular value along a dimension can be ambiguous. Consider a case in which the protective measures dimension takes the second value listed above: the disk was encrypted, but not locked away. It is easy to imagine the plaintiff arguing that this value supports the conclusion that the information was indeed a secret, since, after all, it was encrypted. It is also easy to imaging the defendant arguing that the same value supports the conclusion that the information

ICAIL'17, London, UK
 978-1-4503-3522-XYZ…$15.00
DOI: http://dx.doi.org/10.1145/3086512.XYZ

was not a secret, since it was merely encrypted, and not locked away as well.

Given this distinction between factors and dimensions, then, what problem does it pose for Prakken and Sartor's reconstruction of case-based reasoning in a rule-based setting? Simply this: On one hand, many major case-based systems, with the notable exception of CATO, support reasoning based on dimensions, not just factors, and most researchers in the field, including the authors of CATO, believe that full dimensional resources are necessary for an adequate representation of legal information.[1] But on the other hand, the logical reconstruction of ideas from case-based reasoning offered by Prakken and Sartor takes only factors into account—relying on rules whose premises are conjunctions of factors alone, without analysis of the connection between these factors and the underlying dimensional information.

The point can be illustrated with one of Prakken and Sartor's own hypothetical examples concerning the issue whether an individual who has spent time in another country has changed fiscal domicile with respect to income tax. Among the considerations bearing on this issue is the duration of the individual's stay abroad, where greater duration provides stronger support for change of domicile. It seems most natural to represent duration through a dimension that can take on a variety of values—the individual might have stayed in the other country for a week, a month, six months, a year, five years, and so on. But in fact, Prakken and Sartor bypass the full range of available values and deal with this dimensional information, instead, only through the introduction of a pair of factors—"long-duration" and "short-duration"—where the first favors change of fiscal domicile and the second favors no change.

Even if the problem of relating factors to underlying dimensional information is real, however, it may seem like a small problem. This seems to be what Prakken and Sartor themselves thought.[2] But in a brief and, I feel, somewhat neglected paper, Bench-Capon [4] sets out a number of arguments that raise real concerns about the possibility of handling dimensional information in the kind of rule-based systems used by Prakken and Sartor. On Bench-Capon's view, there are two crucial problems: first, as I have emphasized, no particular value along a dimensional scale necessarily favors one side or the other, in the way that factors do, and second, that if several dimensions are present, strength along one dimension can be traded off for strength along another.

As far as I know, the concerns raised by Bench-Capon have not been resolved. My goal in this paper is to do just that—to propose one way, at least, in which dimensional reasoning can be modeled in a rule-based system. In carrying out this project, I will not be working directly with Prakken and Sartor's logic, which is designed to model legal argument, but instead in a rule-based framework of my own, designed to characterize the concept of legal constraint. Still, if my proposal is successful, it can be adapted to Prakken and Sartor's logic, or to any other framework in which the reasons underlying common law decisions are carried by defeasible rules.

I begin, in the next section, by describing a standard representation of legal cases based on factors alone, and then reviewing—and slightly reformulating—two models of precedential constraint developed within this standard setting. The first is the "result model" of constraint, supporting only a fortiori reasoning and so considered, by many, as too weak to be plausible. The second is what I call the "reason model," supporting a somewhat stronger notion of constraint, since it allows the reasons behind a court's decisions to be taken into account. I show in Section 3 both how the standard representation can be modified to allow for dimensions rather than factors and how the two models of constraint, result and reason, can then be adapted to this new dimensional setting. Finally, in Section 4, I describe one way of interpreting standard information within the dimensional setting and explore relations, in this setting, between the two models of constraint.

2 THE STANDARD SETTING

2.1 Factors, rules, and cases

Let us begin by postulating a set F of *standard factors*—the phrase is meant to distinguish these factors from a different kind of factor to be introduced later on, which will be related to dimensional information. We simplify by imagining 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. As a result, we can suppose that the entire set of standard factors is exhausted by the set $F^\pi = \{f_1^\pi, \ldots, f_n^\pi\}$ of those favoring the plaintiff together with the set $F^\delta = \{f_1^\delta, \ldots, f_m^\delta\}$ of those favoring the defendant: $F = F^\pi \cup F^\delta$. As this notation suggests, we take π and δ to represent the two sides in a dispute, plaintiff and defendant, and where s is a side, we let we let $\overline{s}$ represent the opposite side: $\overline{\pi} = \delta$ and $\overline{\delta} = \pi$.

Within the standard setting, a situation confronting the court—that is, a *standard fact situation*—can be defined simply as some particular subset X of the standard factors: $X \subseteq F$. And where X is a standard fact situation, we let X^s represent the standard factors from X that support the side s: $X^\pi = X \cap F^\pi$ and $X^\delta = X \cap F^\delta$.

Rules will be defined in terms of reasons, where a *standard reason favoring the side s* is some set of factors favoring that side. To illustrate: $\{f_1^\pi, f_2^\pi\}$ is a standard reason favoring the plaintiff, while $\{f_1^\delta\}$ is a standard reason favoring the defendant; but the set $\{f_1^\pi, f_1^\delta\}$ is not a reason, since the factors it contains do not uniformly favor one side or another. Reasons of this kind are to be interpreted conjunctively, so that, for example, the reason $\{f_1^\pi, f_2^\pi\}$ represents the conjunction of the propositions represented by the factors f_1^π and f_2^π.

The idea that a factor holds in a particular situation, or that the situation satisfies that factor, can be defined very simply in the standard case, and then lifted from factors to reasons, or sets of factors, by stipulating that a situation satisfies a set of factors whenever it satisfies each factor from that set.

Definition 2.1. Where X is a standard fact situation and f_n^s is a standard factor, X satisfies f_n^s—written, $X \models f_n^s$—if and only if f_n^s belongs to X.

Definition 2.2. Where X is a fact situation and W is a matching reason, X satisfies W—written, $X \models W$—if and only if X satisfies each factor contained in W.

[1] See both Bench-Capon and Rissland [7] and Rissland and Ashley [20] for arguments supporting the importance of dimensions in legal knowledge representation.
[2] Speaking of dimensions, as well as hypotheticals, they write that "there are no theoretical objections to extending our analysis with these features" [18, p. 279].

We can then define a relation of entailment between reasons, but stipulating that one reason entails another whenever any situation that satisfies the first of these reasons also satisfies the second.

Definition 2.3. Where W and Z are matching reasons, W entails Z—written, $W \Vdash Z$—if and only if $X \models Z$ whenever $X \models W$, for any matching fact situation X.

These definitions call for two comments. First, both Definitions 2.2 and 2.3 contain the requirements that reasons, or reasons and fact situations, must be matching. These requirements can be ignored for now, since we currently have before us only standard reasons and standard fact situations, which we can assume to match. Second, logical entailment among reasons corresponds to strength as a reason—that is, where W and Z are reasons supporting the same side, it is natural to suppose that W is at least as strong a reason as Z for that side whenever $W \Vdash Z$.

Given the notion of a standard reason introduced here, a *standard rule* r can be defined as a statement of the form $W \rightarrow s$, where W is a standard reason supporting the side s. We define two functions—*Premise* and *Conclusion*—picking out the premise and the conclusion of a rule, so that, in the case of this particular rule r, for example, we would have $Premise(r) = W$ and $Conclusion(r) = s$. A rule like this is to be interpreted as defeasible, telling us that its premise entails its conclusion, not as a matter of necessity, but only by default. What the rule $W \rightarrow s$ means, then, is that, if the factors from W hold in some fact situation, then as a default, the court ought to reach a decision in favor of the side s—or perhaps more intuitively, that the factors from W, taken together, provide the court with a reason for deciding for s.

A *standard precedent case* is defined as a standard fact situation together with an outcome and a standard rule through which that outcome is reached or justified. Such a case, then, is a triple of the form $c = \langle X, r, s\rangle$, where X is a fact situation containing the standard factors presented by the case, r is the rule of the case, and s is its outcome. We define three functions—*Facts*, *Rule*, and *Outcome*—mapping cases into their component parts, so that, in the case c above, for example, we have $Facts(c) = X$, $Rule(c) = r$, and $Outcome(c) = s$. The concept of a case is subject to two coherence conditions: first, that the rule of the case must actually apply to the underlying fact situation, or equivalently, that the fact situation satisfies the reason that forms the premise of that rule, and second, that the conclusion of the case rule must match the outcome of the case itself.

These various concepts and constraints 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. 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 factor f_1^π leads, by default, to a decision for the plaintiff.

Finally, a *standard case base* can be defined as a set Γ of standard cases—a set of fact situations presented to various courts, together with their outcomes and the rules justifying these outcomes.

2.2 Two models of constraint

Now, how does an existing case base like this constrain decisions in future cases? We begin by reviewing two models of precedential constraint developed in previous work.

The first, drawn from my [12], is the *result model*, according to which an existing case base constrains a later court only when that court is presented with an a fortiori fact situation—a situation that is at least as strong for the winning side of some precedent case as the fact situation of that precedent case itself.[3]

Obviously, this model relies on some ordering through which different fact situations can be compared in strength for one side or another. The ordering I propose, in the standard setting, is one according to which a fact situation Y presents a case for the side s that is at least as strong as that presented by the fact situation X whenever Y contains all the factors from X that support s, and X contains all the factors from Y that support $\bar{s}$, the opposite side. If we let $\leq^s$ represent the strength ordering for the side s, this idea can then be defined formally as follows:

Definition 2.4. Let X and Y be standard fact situations. Then Y is at least as strong as X for the side s—written, $X \leq^s Y$—if and only if $X^s \subseteq Y^s$ and $Y^{\bar{s}} \subseteq X^{\bar{s}}$.

This definition can be illustrated by considering the previous fact situation $X_1 = \{f_1^\pi, f_2^\pi, f_1^\delta, f_2^\delta\}$ along with the new fact situation $X_2 = \{f_1^\pi, f_2^\pi, f_1^\delta\}$. We then have $X_1 \leq^\pi X_2$, since X_2 contains all the factors from X_1 that support π, and X_1 contains all the factors from X_2 that support δ; and we can see, likewise, that $X_2 \leq^\delta X_1$.

With this strength ordering $\leq^s$ in place, it is straightforward to define the concept of a fortiori constraint at work in the result model:

Definition 2.5. Let Γ be a case base and X a matching fact situation confronting the court. Then the result model of constraint requires the court to reach a decision in X for the side s if and only if there is some case c from Γ such that $Outcome(c) = s$ and $Facts(c) \leq^s X$.

To continue with the same example, suppose the background case base is $\Gamma_1 = \{c_1\}$, containing only the previous case c_1, and that the court is currently confronting the situation X_2. Then the result model of constraint requires a decision in this situation for the plaintiff, since $Facts(c_1) = X_1$ and $Outcome(c_1) = \pi$, and since, as we have seen, $X_1 \leq^\pi X_2$.

The result model presents a picture of precedential constraint that depends only on the comparative strength for a side of the current fact situation relative to the facts of some precedent case, regardless of the explicit rule formulated by the court to justify its decision in that precedent case. There is a long history behind the idea that a court's own efforts at justifying its decision should carry less weight than the decision itself. This history goes back at least to Goodhart's [11] thesis that the *ratio decidendi* of a case is determined only by the decision in that case together with its material facts, and through Goodhart, to earlier work by the American legal realists.[4]

[3] The phrase "result model" is due to Alexander [2].

[4] An interesting discussion of the realist influence on Goodhart is found in Duxbury [9, pp. 80–90].

But, though the history behind this idea may be long, the idea itself represents a minority opinion. Most writers—including Eisenberg [10], Raz [19] and Simpson [22]—believe that the reasons offered by the court must be taken seriously as the basis for its decision. This perspective is captured in the *reason model* of constraint, with roots in the work of Grant Lamond [15], first set out precisely in [13], and developed in the context of artificial intelligence and law by Bench-Capon and myself in [14]. The central feature of this model is that it makes explicit what is, I feel, generally only implicit in case law: a priority ordering representing the strength of the reasons underlying judicial decisions.

To motivate this model of constraint, let us return to the case $c_1 = \langle X_1, r_1, s_1 \rangle$—where $X_1 = \{f_1^\pi, f_2^\pi, f_1^\delta, f_2^\delta\}$, where r_1 is $\{f_1^\pi\} \rightarrow \pi$, and where s_1 is π—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 on the basis of the reason $Premise(r_1)$, the court is also telling us that this reason for the plaintiff—and indeed any reason for the plaintiff that is at least as strong as this one—carries greater weight, or has higher priority, than any reason for the defendant that holds in X_1, the fact situation of the case.

We can recall that a reason W favoring the defendant holds in X_1 whenever $X_1 \models W$, and that a reason Z for the plaintiff is at least as strong as $Premise(r_1)$ whenever $Z \Vdash Premise(r_1)$. If we let $<_{c_1}$ represent the priority relation on reasons that is derived from the particular case c_1, therefore, the force of the court's decision in this case is that, if W is a reason favoring the defendant and Z is a reason favoring the plaintiff, then $W <_{c_1} Z$ whenever $X_1 \models W$ and $Z \Vdash Premise(r_1)$. To illustrate: since $\{f_1^\delta\}$ and $\{f_1^\pi, f_3^\pi\}$ are reasons favoring the defendant and the plaintiff respectively, and since we have both $X_1 \models \{f_1^\delta\}$ and $\{f_1^\pi, f_3^\pi\} \Vdash Premise(r_1)$, it follows that $\{f_1^\delta\} <_{c_1} \{f_1^\pi, f_3^\pi\}$—the court's decision in the case c_1 tells us that the reason $\{f_1^\pi, f_3^\pi\}$ favoring the plaintiff is assigned a higher priority than the reason $\{f_1^\delta\}$ favoring the defendant.

This line of argument leads to the following definition of the priority ordering among reasons that can be derived from a single case:

Definition 2.6. Let $c = \langle X, r, s \rangle$ be a case, and let W and Z be matching reasons favoring the sides $\bar{s}$ and s respectively. Then the relation $<_c$ representing the priority ordering on reasons derived from the case c is defined by stipulating that $W <_c Z$ if and only if $X \models W$ and $Z \Vdash Premise(r)$.

A reader who is familiar with earlier versions of this definition—set out in [13] or [14], for example—might note that I have here replaced the earlier condition that $W \subseteq X$ with the current $X \models W$, and the earlier condition that $Premise(r) \subseteq Z$ with the current $Z \Vdash Premise(r)$. As it turns out, the new conditions are equivalent to the earlier conditions in the standard setting, but allow the definition to generalize to the dimensional setting.

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

Definition 2.7. Let Γ be a case base, and let W and Z be matching reasons. Then the relation $<_\Gamma$ representing the priority ordering on reasons derived from the case base Γ is defined by stipulating that $W <_\Gamma Z$ if and only if $W <_c Z$ for some case c from Γ.

And we can then define a case base as consistent as long as it does not tell us, for some pair of reasons, that each has a higher priority than the other.

Definition 2.8. Let Γ be a case base with $<_\Gamma$ its derived priority ordering. Then Γ is inconsistent if and only if there are matching reasons W and Z such that $W <_\Gamma Z$ and $Z <_\Gamma W$, and consistent otherwise.

Given this notion of consistency, we can turn to the reason model of constraint itself. The intuition is simply that, in deciding a case, a court is required to preserve the consistency of the case base. More exactly, supposing that, against the background of a case base Γ, a court is confronted with a new fact situation X, what the reason model requires is simply that the court reach a decision in the situation X that is consistent with Γ.

Definition 2.9. Let Γ be a case base and X a matching fact situation confronting the court. Then the reason model of constraint requires the court to base its decision on some rule r supporting an outcome s such that the new case base $\Gamma \cup \{\langle X, r, s \rangle\}$ is consistent.

This definition can be illustrated by assuming once more that the background case base is $\Gamma_1 = \{c_1\}$, containing as its single member the familiar case c_1. Suppose that, against this background, the court confronts the fresh situation $X_3 = \{f_1^\pi, f_1^\delta\}$ and considers finding for the defendant on the basis of $\{f_1^\delta\}$, leading to the decision $c_3 = \langle X_3, r_3, s_3 \rangle$, where X_3 is as above, where r_3 is $\{f_1^\delta\} \rightarrow \delta$, and where s_3 is δ. Since $X_3 \models \{f_1^\pi\}$ and $\{f_1^\delta\} \Vdash Premise(r_3)$, we would then have $\{f_1^\pi\} <_{c_3} \{f_1^\delta\}$, according to which the reason $\{f_1^\delta\}$ for the defendant would have to have a higher priority than 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 $\{f_1^\delta\} <_{c_1} \{f_1^\pi\}$, telling us exactly the opposite. As a result, the augmented case base $\Gamma_1 \cup \{c_3\}$ would be inconsistent, so that the reason model would prevent the court from deciding the situation X_3 for the defendant on the basis of the rule r_3.

What are the relations between these two models of constraint, result and reason? It is easy to see that constraint in accord with the result model entails constraint in accord with the reason model.

Observation 1. Let Γ be a consistent case base and X a new fact situation confronting the court, and suppose the result model requires the court to reach a decision for the side s in the situation X. Then the reason model likewise requires the court to reach a decision for the side s.

But at least in the standard setting, the converse entailment fails, as the example just considered shows. Supposing that the court is confronting the new fact situation $X_3 = \{f_1^\pi, f_1^\delta\}$ against the background of the case base $\Gamma_1 = \{c_1\}$, the result model of constraint does not require a decision for the plaintiff, since X_3 is not at least as strong for the plaintiff as some case already decided for the plaintiff—in particular, we do not have $X_1 \leq^{\pi} X_3$. But as we have just seen, the reason model of constraint does require a decision for the plaintiff in this situation, since a decision for the defendant on the only grounds available, the reason $\{f_1^{\delta}\}$, would render the background case base inconsistent. The reason model allows the court to narrow the rule supporting its judgment, thus broadening the scope of its decision and constraining more future cases.

3 THE DIMENSIONAL SETTING

3.1 Dimensions and magnitudes

We now move from the standard setting, with a case representation based on standard factors, to the dimensional setting. While a factor is a legally significant proposition, which either holds or does not, but always favors the same side when it does hold, a dimension, we recall, is an ordered set of values, with the ordering corresponding to the extent to which these values favor one side or the other. The importance of dimensions was illustrated earlier with Prakken and Sartor's change of fiscal domicile example, where the issue under dispute is whether a period of residence in a foreign country counts as a change of fiscal domicile, and where one dimension to consider is the duration of that period, with various lengths of time as values and longer lengths of time favoring change of domicile more strongly. Of course, there may be more than one dimension to consider in a given dispute. In the present example, another relevant dimension might be the proportion of the individual's income derived from organizations based in the foreign country, with particular percentages as values and larger percentages favoring change of domicile.

To represent information like this, we start by postulating a set $D = \{d_1, d_2, \ldots, d_n\}$ of dimensions relevant to some area of dispute. For each dimension, we assume an ordered set of values, ranging from those favoring the side s to those favoring the side $\bar{s}$. Where p and q are values along some fixed dimension, we take the statement

$$p \leq^{s} q$$

to mean that the assignment of the value q to this dimension favors the side s at least as strongly as the assignment of p. This ordering on dimension values is assumed to satisfy the transitivity and antisymmetry conditions

$$p \leq^{s} q \text{ and } q \leq^{s} r \text{ implies } p \leq^{s} r,$$
$$p \leq^{s} q \text{ and } q \leq^{s} q \text{ implies } p = q,$$

as well as a duality condition

$$p \leq^{s} q \text{ if and only if } q \leq^{\bar{s}} p,$$

according to which q favors the side s at least as much as p just in case p favors the opposite side $\bar{s}$ at least as much as q.

This notation can be illustrated with the fiscal domicile example if we imagine that the plaintiff is the individual's native country, which is arguing against change of domicile in order to tax the individual's income, and that the defendant is the individual, who is arguing for change of domicile in order to pay, we can suppose, the lower tax rates available in a foreign country. Here, two possible values along the dimension representing the period of residence abroad are six months and eighteen months. If these values are represented simply as 6 and 18, we have $6 \leq^{\delta} 18$, since the longer period favors the defendant's argument for change of domicile; duality then tells us that $18 \leq^{\pi} 6$, since the shorter period favors the plaintiff's argument against change.

Where p is a value along the dimension d, the pair $\langle d,p \rangle$ is a *value assignment*, according to which the dimension d takes on the value p. In contrast to a standard fact situation, defined earlier as a set of standard factors, a *dimensional fact situation*

$$X = \{\langle d,p \rangle : d \in D\}$$

can be defined as a set of values assignments, one for each dimension, subject to the condition that if $\langle d,p \rangle$ and $\langle d,p' \rangle$ both belong to X, then $p = p'$. A dimensional fact situation, in other words, is a function mapping each dimension to a value along that dimension. We take $X(d)$ as the value assigned to the dimension d in the fact situation X, where this idea is defined in the usual way:

$$X(d) = p \text{ if and only if } \langle d,p \rangle \in X.$$

The central conceptual problem presented by the dimensional setting is that, while standard fact situations are constructed out of standard factors, which always favor one side or the other, dimensional fact situations are constructed out of value assignments, which need not favor any particular side. We address this problem by introducing a different class of factors, like standard factors in possessing a definite polarity, but keyed to the value assignments found in dimensional fact situations.

To motivate the proposal, imagine that an individual has been living in a foreign country for two and a half years, and the question of fiscal domicile hinges on whether a period of that length counts as a long duration. How could a court reach a decision in this situation, and how could it justify its decision? My suggestion is that the court might focus on some particular value along the dimension scale—a value that seems to be salient—and then both make and justify its decision by comparing the value of the dimension in the current fact situation to that salient value. Suppose, for instance, that the value of one year seems, to the court, like a sufficient and salient length of time to count as a long duration. The court could then register its decision in the current situation by ruling for change of fiscal domicile, and so in favor of the defendant, on the grounds that the period of residence in the foreign country lasted at least a year.

This proposition—to spell it out, that the actual period of residence abroad favors the defendant at least as much as a period of one year—is a kind of factor: it either holds or does not hold in any fact situation, and always favors the same side, the defendant, when it does hold. Generalizing from our example, then, where p is some value along the dimension d, we now introduce the concept of a *magnitude factor favoring the side s*, as a statement of the form

$$M^{s}_{d,p}$$

carrying the meaning: the actual value assigned to the dimension d favors the side s as least as strongly as the value p. If we take d_1 as the dimension representing length of time abroad, the magnitude factor at work in our example can now be expressed as $M^{\delta}_{d_1,12}$, the proposition that the actual value assigned to d_1 in the situation at hand favors the defendant at least as much as a value of twelve months, or one year.

Once these magnitude factors have been introduced, we follow the pattern from the standard setting by defining a *magnitude reason favoring the side s* as a set of magnitude factors favoring that side. A factor collection of the form $\{M^{\pi}_{d_1,p}, M^{\pi}_{d_2,q}\}$, then, would be a magnitude reason favoring the plaintiff, carrying the conjunctive meaning that the actual value assigned to d_1 favors the plaintiff as least as strongly as p and the actual value assigned to d_2 favors the plaintiff as least as strongly as q.

Now, what about satisfaction and entailment? In the standard setting, where fact situations were simply sets of standard factors, a fact situation could be said to satisfy a factor whenever that factor belonged to the situation. But this idea, set out in Definition 2.1, cannot carry over to the dimensional setting, since a fact situation is now defined as a set of value assignments and a magnitude factor is something else entirely—a statement of the form $M^{s}_{d,p}$, carrying the meaning, once again, that the value assigned to dimension d favors the side s at least as strongly as the value p. Here, since the value assigned to the dimension d in some situation X is simply $X(d)$, and since this value favors the side s at least as strongly as the value p whenever $p \preceq^{s} X(d)$, the conditions under which a dimensional fact situation satisfies a magnitude factor can be defined as follows:

Definition 3.1. Where X is a dimensional fact situation and $M^{s}_{d,p}$ is a magnitude factor, X satisfies $M^{s}_{d,p}$—written, $X \models M^{s}_{d,p}$—if and only if $p \preceq^{s} X(d)$.

To illustrate, imagine that the dimensional fact situation X contains the value assignment $\langle d_1, 30\rangle$, indicating that the individual has been in the foreign country for a period of two and a half years, or thirty months, so that $X(d_1) = 30$. Since a period of thirty months favors the conclusion of an absence of long duration, and so the defendant, more than a period of twelve months, we therefore have $12 \preceq^{\delta} X(d_1)$, which now tells us that $X \models M^{\delta}_{d_1,12}$.

Once the previous treatment of factor satisfaction from Definition 2.1 has been replaced with this new treatment, from Definition 3.1, our earlier definitions of reason satisfaction and entailment can be carried over without change to the dimensional setting. We can say, in accord with Definition 2.2, that a dimensional fact situation satisfies a magnitude reason whenever that fact situation satisfies each magnitude factor belonging to that reason, and in accord with Definition 2.3, that one magnitude reason entails another whenever every dimensional fact situation that satisfies the first also satisfies the second. In interpreting these definitions, of course, we must attend to the matching requirements, which guarantee that they apply only within settings, not across settings.

If W is a magnitude reason favoring the side s, we can now define $W \rightarrow s$ as a *magnitude rule*, where this rule is interpreted defeasibly, just like a standard rule, and where the functions *Premise* and *Conclusion* picking out the premise and conclusion of this rule are defined as before. And following the pattern of the standard setting, we can define a *dimensional case* as a triple $c = \langle X, r, s\rangle$, where X is a dimensional fact situation, r is a magnitude rule justifying a particular outcome, and s is the case outcome itself. As before, we have three functions—*Facts*, *Rule*, and *Outcome*—mapping cases into their component parts. And we require, as coherence condition on the concept of a case, both that the rule of the case should apply to its fact situation and that the conclusion of the case rule should match the case outcome.

These ideas can be illustrated through the concrete example $c_4 = \langle X_4, r_4, s_4\rangle$, where $X_4 = \{\langle d_1, 30\rangle, \langle d_2, 60\rangle\}$ is the underlying dimensional fact situation, where r_4 is the magnitude rule $\{M^{\delta}_{d_1,12}\} \rightarrow \delta$, and where s_4 is δ, a decision for the plaintiff. If we take d_1 and d_2 as the dimensions representing length of time in a foreign country and proportion of income earned from organizations based in that country, then X_4 represents a fact situation in which an individual spent two and a half years in a foreign country, and during that period earned sixty percent of his or her income from organizations based in that country. The case, then, is one in which, faced with this situation, the court ruled for change of fiscal domicile, and so in favor of the defendant, on the grounds that the individual spent at least a year in the foreign country.

Finally, and as in the standard setting, we define a *dimensional case base* Γ as a set of dimensional cases.

3.2 Constraint

How can our two models of constraint, result and reason, be adapted to the dimensional setting?

The result model, we recall, was meant to capture a fortiori reasoning—according to which a court is constrained to decide a situation for a particular side whenever that situation is at least as strong for that side as the fact situation from some case that has already been decided for that side—and so depends on an ordering through which different fact situations can be compared in strength for one side or another. In the standard setting, with standard fact situations built from standard factors, this ordering was set out in Definition 2.4, but of course, that definition is no longer applicable in the dimensional setting. Fortunately, it is plain how the new definition should go: the dimensional fact situation Y should now be defined to be at least as strong for a side as the dimensional fact situation X whenever, for every dimension, the value assigned by Y to that dimension favors that side at least as much as the value assigned by X. Continuing to use $\preceq^{s}$ to represent strength for a side s, this new definition can be stated formally as follows:

Definition 3.2. Let X and Y be dimensional fact situations. Then Y is at least as strong as X for the side s—written, $X \preceq^{s} Y$—if and only if $X(d) \preceq^{s} Y(d)$ for each dimension d from D.

And once this new concept of strength for a side is in place, our previous specification of the result model in terms of strength for a side, set out in Definition 2.5, can be carried over without change.

The result model, in the dimensional setting, can be illustrated by taking as background the dimensional case base $\Gamma_2 = \{c_4\}$, containing only the earlier dimensional case c_4, and imagining that the court is confronting the new situation $X_5 = \{\langle d_1, 36\rangle, \langle d_2, 65\rangle\}$, representing a state of affairs in which an individual spent three years in a foreign country while earning sixty-five percent of his or her income from organizations based there. Comparing this fresh situation to the earlier $X_4 = \{\langle d_1, 30\rangle, \langle d_2, 60\rangle\}$, our new Definition 3.2 tells us that $X_4 \preceq^{\delta} X_5$, since $30 \preceq^{\delta} 36$ along the dimension d_1 and $60 \preceq^{\delta} 65$ along the dimension d_2. Definition 2.5 then tells us that, according to the result model, X_5 must be decided for the defendant, since this situation is at least as strong for the defendant as X_4, the fact situation from the previous case c_4 that was already decided for the defendant.

The result model, then, can be adapted in a straightforward way to the dimensional setting, but what about the reason model? Here we can see the point of our reformulation of the reason model in terms of the logical ideas of reason satisfaction and reason entailment—for it now turns out that, subject only to matching restrictions, the treatment of the reason model set out earlier carries over without change to the dimensional setting.

To spell it out: According to Definition 2.9, the reason model requires a court, faced with a fresh fact situation and working against the background of an existing case base, to reach a decision that maintains consistency of the case base. A case base is consistent, according to Definition 2.8, as long as there are no two reasons each of which is prioritized over the other on the basis of the priority ordering derived from that case base, where this idea is set out in Definition 2.7, which itself relies on the central concept, set out in Definition 2.6, of the priority ordering on reasons derived from a single case. This latter definition draws on the ideas of reason satisfaction and reason entailment from Definitions 2.2 and 2.3, which themselves bottom out, in the standard setting, in the treatment of standard factor satisfaction from Definition 2.1 and, in the dimensional setting, in the new treatment of magnitude factor satisfaction from Definition 3.1. The entire structure of the reason model is thus identical in the standard and dimensional settings, differing only at the very bottom, with different definitions of satisfaction for standard and magnitude factors.

Still, even though the dimensional reason model simply parallels the standard version, it is worth discussing a few examples in order to understand the shape of this model in the more complex dimensional setting.

Example 1. Continuing to work in the fiscal domicile domain, we begin with the simplifying assumption that there is only one dimension of concern: length of time abroad. Suppose that $\Gamma_3 = \{c_6\}$, with $c_6 = \langle X_6, r_6, s_6\rangle$, where $X_6 = \{\langle d_1, 30\rangle\}$, where r_6 is $\{M^{\delta}_{d_1,12}\} \rightarrow \delta$, and where s_6 is δ. The single case in this case base, then, is one in which the defendant has spent two and a half years abroad and the court ruled for change of fiscal domicile, and so in favor of the defendant, on the grounds that the period abroad lasted at least a year. Now, against this background, imagine that the new situation $X_7 = \{\langle d_1, 18\rangle\}$, representing a defendant who has spent a year and a half, or eighteen months, abroad comes before a different court. The new court, we can suppose, has stricter standards for change of fiscal domicile and would prefer to rule against change in this situation, and so for the plaintiff, on the grounds that the defendant has spent less than two years abroad—that is, that the actual period abroad favors the plaintiff at least as much as a period of two years, or twenty-four months. The resulting decision would be represented by the new case $c_7 = \langle X_7, r_7, s_7\rangle$, where X_7 is as above, where r_7 is $\{M^{\pi}_{d_1,24}\} \rightarrow \pi$, and where s_7 is π.

Can the new court rule as it prefers? It can, according to the reason model, since the resulting case base $\Gamma_3 \cup \{c_7\}$ is consistent.

Example 2. As in the previous example, we take as background the case base $\Gamma_3 = \{c_6\}$, with c_6 as before, and again imagine that a new court is confronted with the fact situation $X_7 = \{\langle d_1, 18\rangle\}$. This time, however, the new court embraces standards for change of fiscal domicile very much stricter than those of the original court, and would prefer to rule against change, and so for the plaintiff, not on the grounds that the defendant spent less than two years abroad, but on the grounds that the defendant spent less than three years, or thirty-six months, abroad. The resulting decision would be represented by the case $c_8 = \langle X_8, r_8, s_8\rangle$, where $X_8 = X_7$, where r_8 is $\{M^{\pi}_{d_1,36}\} \rightarrow \pi$, and where s_8 is π.

Again we must ask whether the new court can rule as it prefers, and the answer this time is that it cannot, since the resulting case base $\Gamma_3 \cup \{c_8\}$ would be inconsistent. This fact can be seen informally by noting that the new case rule r_8 would indeed apply to the previous situation X_6, and conflict with the decision reached regarding that situation by the c_6 court. For a more formal verification, we need only note that $X_6 \models Premise(r_8)$ and of course that $Premise(r_6) \Vdash Premise(r_6)$, from which it follows by Definition 2.6 that $Premise(r_8) <_{c_6} Premise(r_6)$, and also that $X_8 \models Premise(r_6)$ and $Premise(r_8) \Vdash Premise(r_8)$, from which it likewise follows that $Premise(r_6) <_{c_8} Premise(r_8)$.

Example 3. We now consider a fiscal domicile example in which there is more than one dimension of concern. Take as background the previous case base $\Gamma_2 = \{c_4\}$, containing $c_4 = \langle X_4, r_4, s_4\rangle$, where $X_4 = \{\langle d_1, 30\rangle, \langle d_2, 60\rangle\}$, where r_4 is $\{M^{\delta}_{d_1,12}\} \rightarrow \delta$, and where s_4 is δ. This case, once again, represents a situation in which the defendant spent two and a half years in a foreign country while earning sixty percent of his or her income there, and in which the court decided for the defendant on the grounds that the period abroad lasted at least a year. Now imagine that, against this background, a different court confronts the new fact situation $X_9 = \{\langle d_1, 36\rangle, \langle d_2, 10\rangle\}$, representing an individual who easily satisfies the rule of the c_4 court by spending three years in a foreign country, but during that period earned only ten percent of his or her income from organizations based there .

In considering this situation, the new court might be struck by the remarkably low proportion of income earned from foreign organizations, suppose that this possibility escaped notice of the c_4 court when it formulated its rule based solely on length of stay abroad, and therefore rule against change of domicile, and so in favor of the plaintiff, on the grounds that less than twenty-five percent of income was earned from foreign organizations. This decision would be represented by the case $c_9 = \langle X_9, r_9, s_9\rangle$, where X_9 is as above, where r_9 is $\{M^{\pi}_{d_2,25}\} \rightarrow \pi$, and where s_9 is π. We leave it to the reader to verify that this decision is allowed by the rule model of constraint, since the resulting case base $\Gamma_2 \cup \{c_9\}$ is consistent. What this example shows is that strength along one dimension can be overcome by weakness along another—there are tradeoffs among dimensions.

Observation 1 continues to hold in the dimensional setting, from which it follows that result constraint entails reason constraint. As we saw earlier, the converse does not hold in the standard setting: a court can be constrained by the reason model to reach a decision for a particular side even though that decision is not forced by the result model. Surprisingly, though, the converse of Observation 1 does hold in the full dimensional setting: not only does result constraint entail reason constraint, but reason constraint entails result constraint.

Observation 2. Let Γ be a consistent dimensional case base and X a new dimensional fact situation confronting the court, and suppose the reason model requires the court to reach a decision for the side s in the situation X. Then the result model likewise requires the court to reach a decision for the side s.

Does this mean that there is no difference, in the dimensional setting, between the reason and result models of constraint, that the reason model simply collapses into the result model? The answer is that the two models do collapse when the full range of magnitude reasons is allowed, but that they can come apart, and the reason model once again offers a stronger notion of constraint, when the range of magnitude reasons under consideration is restricted. We return to this topic in the next section.

4 INTERPRETING STANDARD INFORMATION

4.1 An interpretation

Having explored the notion of precedential constraint in two general settings, standard and dimensional, we now describe one way in which information from the standard setting can be interpreted into the dimensional setting.

It is sometimes suggested that standard factors can be thought of as points along a dimension.[5] The interpretation set out here is based, however, on an understanding of standard factors themselves as dimensions. Considered as dimensions, the standard factors are seen as taking on the boolean values 1 and 0, indicating presence or absence, with these values ordered so that, for a standard factor favoring the side s, its presence favors s at least as much as its absence, and its absence favors the opposite side $\overline{s}$ at least as much as its presence. This idea is encoded using our value ordering notation by stipulating that, where f is a standard factor and p and q are its possible boolean values, and where $\leq$ is the usual ordering on numbers, then

$$p \leq^s q \text{ if and only if } \begin{cases} p \leq q & \text{and } f \in F^s \\ q \leq p & \text{and } f \in F^{\overline{s}}. \end{cases}$$

The reader can verify that the $<^s$ ordering defined in this way satisfies our earlier conditions of transitivity, antisymmetry, and duality.

We now define a *dimensionalization* function $\mathcal{D}$, mapping items from the standard setting into their dimensional counterparts, in five steps. First, where X is a standard fact situation, we take

$$\mathcal{D}(X) = \{\langle f, 1\rangle : f \in X\} \cup \{\langle f, 0\rangle : f \in F - X\},$$

as its corresponding dimensional fact situation, where, of course, the value assignments $\langle f, 1\rangle$ and $\langle f, 0\rangle$ indicate the presence or absence of the standard factor f in X. There is a slight wrinkle when it comes to interpreting standard reasons in the dimensional setting, since standard reasons are objects of the same type as standard fact situations—sets of standard factors—and so would likewise be mapped by $\mathcal{D}$ into dimensional fact situations, rather than magnitude reasons. As our second step, we therefore introduce a separate function $\mathcal{D}'$ mapping standard reasons into their dimensional counterparts in such a way that, where W is a standard reason favoring the side s, its dimensionalization is

$$\mathcal{D}'(W) = \{M^s_{f,1} : f \in W\},$$

a magnitude reason that is satisfied by a dimensional fact situation just in case, for each standard factor present in W, that dimensional fact situation contains a value assignment indicating the presence of that factor.

Third, where r is a standard rule supporting the outcome s, its dimensionalization $\mathcal{D}(r)$ is a magnitude rule of the form

$$\mathcal{D}'(Premise(r)) \rightarrow s,$$

supporting the same outcome as the original, and taking as its premise the dimensionalization of the standard reason that forms the premise of the original rule. Fourth, where $c = \langle X, r, s\rangle$ is a standard case, its dimensionalization

$$\mathcal{D}(c) = \langle \mathcal{D}(X), \mathcal{D}(r), s\rangle$$

is the dimensional case containing the dimensionalization of the fact situation and rule from the original case, and the same outcome. Fifth, and finally, where Γ is a standard case base, its dimensionalization is

$$\mathcal{D}(\Gamma) = \{\mathcal{D}(c) : c \in \Gamma\}$$

containing dimensionalizations of each case belonging to the original.

We can see these definitions at work by calculating the dimensionalization of the case base $\Gamma_1 = \{c_1\}$, considered earlier, containing the single case $c_1 = \langle X_1, r_1, s_1\rangle$, where $X_1 = \{f^\pi_1, f^\pi_2, f^\delta_1, f^\delta_2\}$, where r_1 is $\{f^\pi_1\} \rightarrow \pi$, and where s_1 is π. Descending through the steps in our definition, we have $\mathcal{D}(\Gamma_1) = \{\mathcal{D}(c_1)\}$, with $\mathcal{D}(c_1) = \langle \mathcal{D}(X_1), \mathcal{D}(r_1), s_1\rangle$, where $\mathcal{D}(X_1)$ is the dimensional fact situation $\{\langle f^\pi_1, 1\rangle, \langle f^\pi_2, 1\rangle, \langle f^\delta_1, 1\rangle, \langle f^\delta_2, 1\rangle\}$ and where $\mathcal{D}(r_1)$ is the magnitude rule $\{M^\pi_{f^\pi_1,1}\} \rightarrow \pi$.

The question now arises: to what extent are the constraint relations defined in the standard setting preserved under the mapping described here from standard to dimensional information? Or more exactly: given a standard case base Γ and fact situation X, is it the case that a court working against the background of Γ is constrained to decide X for s just in case, moving to the dimensional setting, a court working against the background of $\mathcal{D}(\Gamma)$ is constrained to decide $\mathcal{D}(X)$ for s? Since we are working with two models of constraint, result and reason, we need to ask the question separately for each model.

Beginning with the result model, it turns out that, here, the notion of constraint defined in the standard setting carries over without change to the dimensional setting.

Observation 3. Let Γ be a standard case base and X a standard fact situation confronting the court. Then the result model of constraint requires the court working against the background of Γ to reach a decision in X for the side s if and only if, moving to the dimensional setting, the result model of constraint requires the court working against the background of the dimensional case base $\mathcal{D}(\Gamma)$ to reach a decision in the dimensional fact situation $\mathcal{D}(X)$ for the side s.

[5] See, for example, Bench-Capon [6].

Things are different, though, when we turn to the reason model: here, the concept of constraint defined in the standard setting fails to survive interpretation into the dimensional setting. We can see this by reconsidering our earlier example in which a court faces the standard fact situation $X_3 = \{f_1^\pi, f_1^\delta\}$ against the background of the standard case base $\Gamma_1 = \{c_1\}$. As we saw earlier, the reason model of constraint requires the court to decide this situation for the plaintiff, since a decision for the defendant would prioritize $\{f_1^\delta\}$ over $\{f_1^\pi\}$, but the opposite priority ordering is already supported by the background case base. If the example is interpreted in the dimensional setting, however—that is, supposing the court faces the dimensional fact situation $\mathcal{D}(X_3) = \{\langle f_1^\pi, 1\rangle, \langle f_2^\pi, 0\rangle, \langle f_1^\delta, 1\rangle, \langle f_2^\delta, 0\rangle\}$ against the background of the dimensional case base $\mathcal{D}(\Gamma_1)$—then the reason model no longer requires a decision for the plaintiff. This fact follows at once from Observation 2, according to which, in the dimensional setting, the reason model requires a decision for a particular side only if the result model also requires a decision for that side. But the result model does not require a decision for the plaintiff in this case, since we do not have $\mathcal{D}(X_1) \leq^\pi \mathcal{D}(X_3)$. Why not? Because there is at least one dimension, the dimension f_2^π, whose value in the fact situation $\mathcal{D}(X_3)$ does not favor the plaintiff as strongly as its value in $\mathcal{D}(X_1)$; in particular, $\mathcal{D}(X_1)(f_2^\pi)$—that is, the value assigned to f_2^π in the fact situation $\mathcal{D}(X_1)$—is 1 while $\mathcal{D}(X_3)(f_2^\pi)$ is 0, and since f_1^π favors the plaintiff, it follows from our ordering definition that we do not have $\mathcal{D}(X_1)(f_2^\pi) \leq^\pi \mathcal{D}(X_3)(f_2^\pi)$.

4.2 Restricting the reasons

Since, by Observation 2, the reason model collapses into the result model in the dimensional setting, it is natural to wonder whether the reason model can be modified so that it allows, in the dimensional setting, a pattern of constraint that aligns with that of the standard reason model. In fact, it can, by restricting the reason model so that appeal to inappropriate reasons and rules is blocked .

In order to do this in the most general way, we proceed in three steps. First, we generalize the reason model of constraint from Definition 2.9 to allow for restrictions on the rules available to justify decisions. Although this new version of the reason model differs from the original only in a single word—the requirement that the rule appealed to should be *appropriate*—it is formulated explicitly here for the sake of clarity:

Definition 4.1. Let Γ be a case base and X a matching fact situation confronting the court. Then the reason model of constraint requires the court to base its decision on some appropriate rule r supporting an outcome s such that the new case base $\Gamma \cup \{\langle X, r, s\rangle\}$ is consistent.

Now, what does it mean for a rule to be appropriate? That varies, and can be a matter of contention—in many domains, this would be a research question. For present purposes, however, where we are considering the interpretation of standard information into the dimensional setting, we can stipulate that the appropriate reasons and rules are those that result from the dimensionalization of standard reasons and rules. Second, then, let us define a *dimensionalized standard reason* as a magnitude reason of the form $\mathcal{D}'(W)$ for some standard reason W, and a *dimensionalized standard rule* as a magnitude rule with a dimensionalized standard reason as its premise—that is, a magnitude rule of the form $\mathcal{D}(r)$ for some standard rule r. Finally, if we stipulate for present purposes that all and only the dimensionalized standard rules are appropriate, the pattern of constraint defined by the reason model in the standard setting is, at last, preserved under our mapping from standard to dimensional information, as we can see from the following observation:

Observation 4. Let Γ be a consistent standard case base and X a standard fact situation confronting the court. Then there is some standard rule r supporting the outcome s such that the standard case base $\Gamma \cup \{\langle X, r, s\rangle\}$ is consistent if and only if, moving to the dimensional setting, there is some dimensionalized standard rule r' supporting the same outcome such that the dimensional case base $\mathcal{D}(\Gamma) \cup \{\langle \mathcal{D}(X), r', s\rangle\}$ is consistent.

The upshot is this: In the dimensional setting, the difference between the result and reason models of constraint can be captured by working with a single definition of constraint, set out in Definition 4.1, but classifying different reasons and rules as appropriate. The result model emerges as the special case in which all reasons and rules are classified as appropriate, while restrictions on the set of appropriate reasons and rules lead to patterns of constraint characteristic of the reason model.

5 CONCLUSION

My goal has been to show—in response to concerns raised by Bench-Capon—how two simple models of precedential constraint, the result and reason models, can be broadened from the standard setting to the dimensional setting, allowing a richer representation of legal information, not just through sets of standard factors, but through dimensions that can take on various values. The path followed in this paper is, in many ways, straightforward, but it led to a surprise. The surprise is that the two models of constraint, which are distinct in the standard setting, collapse in the richer dimensional setting. As a response to this collapse, in order to restore a contrast between the result and reason models of constraint, we generalized the reason model in a way that allows for limiting consideration only to certain preferred, or appropriate, reasons and rules.

This generalization of the reason model may appear to be artificial, or ad hoc, especially given the somewhat technical problem through which it was motivated. In fact, though, I feel that, by introducing the question of what reasons and rules should be classified as appropriate, and why, the highlights some of the deeper issues underlying the notion of precedential constraint. To support this point, I close by describing a different and less technical setting in which a contrast between the result and reason models hinges on the question whether a particular rule is appropriate.

Returning to the change of fiscal domain scenario, suppose an individual has lived in a foreign country for a year and ten days and a court rules that this period counts as a long duration, and so for the defendant, on the grounds that the individual has been away for over a year. According to the result model, of course, the reason the court gives to support its judgment has no bearing on constraint; this particular judgment constrains a later case only if the facts of that case are at least as strong for the defendant as the facts of the present case—that is, only if an individual has spent at least a year and ten days in a foreign country. And what Observation 2 tells us is that, surprisingly, the initial version of the reason model, set out in Definition 2.9, yields the same outcome: any future situation in which an individual has lived in a foreign country for less than a year and ten days is distinguishable, even if the period of foreign residence is longer than a year, and so falls within the scope of the original rule.

But now, imagine actually trying to distinguish in such a case. Suppose that, against the background of the previous decision, a case arises concerning an individual who has lived abroad for a year and five days . According to the initial version of the reason model, the court could now consistently rule that this period of time should not count as a period of long duration, and so for the plaintiff, on the grounds that it is, say, no longer than a year and six days. I think that we would all find such a judgment to be objectionable, not on the grounds that it is formally inconsistent with the background case base—it is, again, not inconsistent—but on the grounds that the rule set out by the court to justify its decision is so contrived, that it appears to be crafted simply to justify a desired outcome in a particular case, rather than reflecting a coherent principle, and that it promises to gerrymander the range of potential cases, rather than reflecting any natural division.

What the new Definition 4.1 version of the reason model allows us to do is reject such a rule as inappropriate, thus preventing the reason model from collapsing into the result model in the dimensional setting. Suppose, for example, that we decide that only rules based on yearly increments are appropriate—one year, two years, three years, and so on. Under that assumption, the initial court's decision does in fact constrain the new situation, in which an individual spent a year and five days abroad, not because there is no rule on the basis of which this situation can be consistently distinguished, but because no such rule is appropriate.

In a way, it should come as no surprise that we should find ourself considering such higher-order features of rules. Some of the most interesting work in artificial intelligence and law has focused on the evaluation of rules with respect to various higher-order considerations, most particularly the values advanced by those rules—this theme was first sounded by Bench-Capon himself [5], and echoed in the work of, among others, Bench-Capon and Sartor [8] and Prakken [17]. But what is surprising, at least to me, is that, once we move to the dimensional setting, we should be driven to consider higher-order features of rules simply in order to maintain a distinction between the result model and the reason model of constraint.

REFERENCES

[1] Vincent Aleven and Kevin Ashley. 1997. Evaluating a learning environment for case-based argumentation skills. In *Proceedings of the Sixth International Conference on Artificial Intelligence and Law (ICAIL-97)*. The Association for Computing Machinery Press, 170–179.

[2] Larry Alexander. 1989. Constrained by precedent. *Southern California Law Review* 63 (1989), 1–64.

[3] Kevin Ashley. 1990. *Modeling Legal Argument: Reasoning with Cases and Hypotheticals*. The MIT Press.

[4] Trevor Bench-Capon. 1999. Some observations on modelling case based reasoning with formal argument models. In *Proceedings of the Sixth International Conference on Artificial Intelligence and Law (ICAIL-99)*. The Association for Computing Machinery Press, 36–42.

[5] Trevor Bench-Capon. 2002. The missing link revisited: the role of teleology in representing legal argument. *Artificial Intelligence and Law* 10 (2002), 79–94.

[6] Trevor Bench-Capon. 2012. Representing Popov v. Hayashi with dimensions and factors. *Artificial Intelligence and Law* 20 (2012), 15–35.

[7] Trevor Bench-Capon and Edwina Rissland. 2001. Back to the future: dimensions revisited. In *The Fourteenth Annual Conference on Legal Knowledge and Information Systems (JURIX-2001)*. IOS Press, 41–52.

[8] Trevor Bench-Capon and Giovanni Sartor. 2003. A model of legal reasoning with cases incorporating theories and values. *Artificial Intelligence* 150 (2003), 97–143.

[9] Neil Duxbury. 2005. *The Nature and Authority of Precedent*. Cambridge University Press.

[10] Melvin Eisenberg. 1988. *The Nature of the Common Law*. Harvard University Press.

[11] Arthur Goodhart. 1930. Determining the *ratio decidendi* of a case. *Yale Law Journal* 40 (1930), 161–183.

[12] John Horty. 2004. The result model of precedent. *Legal Theory* 10 (2004), 19–31.

[13] John Horty. 2011. Rules and reasons in the theory of precedent. *Legal Theory* 17 (2011), 1–33.

[14] John Horty and Trevor Bench-Capon. 2012. A factor-based definition of precedential constraint. *Artificial Intelligence and Law* 20 (2012), 181–214.

[15] Grant Lamond. 2005. Do precedents create rules? *Legal Theory* 11 (2005), 1–26.

[16] Ronald Loui, Jeff Norman, Jon Olson, and Andrew Merrill. 1993. A design for reasoning with policies, precedents, and rationales. In *Proceedings of the Fourth International Conference on Artificial Intelligence and Law (ICAIL-93)*. The Association for Computing Machinery Press, 201–211.

[17] Henry Prakken. 2002. An exercise in formalising teleological case-based reasoning. *Artificial Intelligence and Law* 10 (2002), 113–133.

[18] Henry Prakken and Giovanni Sartor. 1998. Modelling reasoning with precedents in a formal dialogue game. *Artificial Intelligence and Law* 6 (1998), 231–287.

[19] Joseph Raz. 1979. *The Authority of Law*. Oxford University Press.

[20] Edwina Rissland and Kevin Ashley. 2002. A note on dimensions and factors. *Artificial Intelligence and Law* 10 (2002), 65–77.

[21] Edwina Rissland and David Skalak. 1989. Interpreting statutory predicates. In *Proceedings of the Second International Conference on Artificial Intelligence and Law (ICAIL-89)*. The Association for Computing Machinery Press, 46–53.

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