Artificial Intelligence and Law
https://doi.org/10.1007/s10506-023-09384-5

ORIGINAL RESEARCH

Toward representing interpretation in factor-based models of precedent

Adam Rigoni[1]

Accepted: 20 November 2023

Abstract
This article discusses the desirability and feasibility of modeling precedents with multiple interpretations within factor-based models of precedential constraint. The main idea is that allowing multiple reasonable interpretations of cases and modeling precedential constraint as a function of what all reasonable interpretations compel may be advantageous. The article explains the potential benefits of extending the models in this way with a focus on incorporating a theory of vertical precedent in U.S. federal appellate courts. It also considers the costs of extending the models in this way, such as the significant increase in the functional size of the case base and the need to provide some kind of ordering on interpretations to select a "best" interpretation. Finally, the article suggests partially incorporating multiple interpretations of dimensions as a realistic starting point for incorporating interpretations generally, and shows how doing so can help address difficulties with dimensions. The conclusion remarks on the use of interpretations to deal with inconsistent precedents.

Keywords Legal factors · Precedential constraint · Dimensions · Interpretation · Case-based reasoning

1 Introduction

This article discusses the potential role of interpretation in factor-based models of common law judicial reasoning (henceforth "factor-models"). It stresses the importance of "where" in the common law judicial reasoning process models place particular inferences and decisions and how this relates to precedential constraint. Let us stipulate, although it is an oversimplification, that common law judicial reasoning

✉ Adam Rigoni
arigoni@asu.edu

[1] Barrett Honors College, Arizona State University, Tempe, AZ, USA

Published online: 12 January 2024

in a case has three steps: (1) fact-finding for the current case, (2) interpretation of past cases, and (3) the application of (2) to (1).[1] There is a cluster of factor-models of precedent (even though many employ dimensions as distinct from factors) in recent AI and Law literature (Horty and Bench-Capon 2012; Atkinson et al. 2013; Al-Abdulkarim et al. 2015; Rigoni 2015; Mullins 2022). All of these models only represent step (3). They begin with a case base wherein every case is already interpreted: the relevant rules, factors, priority orderings, etc. are established for each of those cases. They also start with the current case divided into facts/points on dimensions/factors that can trigger those rules, be prioritized by those orderings, etc.

More recent work (Ashley and Brüninghaus 2009; Branting 2020; Bench-Capon and Atkinson 2022; Collenette et al. 2023), as well as some past work (Ashley 1991), incorporates these factor based approaches, or variants thereof within larger frameworks that model other aspects of the reasoning process, including the ascription of factors to a the current case, and the use of factor-based determinations to establish issues that are then used to reach an outcome. My suggestions here are in line with and motivated by this work, but I focus on the value of allowing cases to have multiple interpretations within the model. The main idea is that allowing multiple reasonable interpretations of cases and modeling precedential constraint as a function of what all reasonable interpretations compel is advantageous. Although it focuses on factor models of precedential constraint, the main points are applicable more broadly. To be clear, my claim is not that the relevant reasoning processes *are actually* interpretative instead of fact-finding or applying a legal rule to a particular fact. Rather, my goal is to explain the potential benefits and costs for modeling the relevant processes as interpretative as compared to modeling them in other ways.

For example, consider factor-models such as (Horty and Bench Capon 2012; Rigoni 2018). On these models, roughly, a case is a triple $<F, R, O>$ where F is the set of factors in the case, R is the rule of the case, and O is the outcome. O must be either {plaintiff} or {defendant}, depending on the outcome of the case. The factors in the case, F, are reasons for or against deciding the case for the plaintiff or the defendant. We will use superscripts to denote the valence of the factor and subscripts to number them, hence F^p_1, F^p_2 are two different factors favoring the plaintiff and F^d_1, F^d_2 are two different pro-defendant factors. R is a conditional where the consequent is the outcome of the case and the antecedent is a set of the factors that are both present in the case and that favor the outcome. Call the antecedent of a rule, R, *Premise (R)*. For example, if F^p_1 and F^p_2 are present in a case that was decided for the plaintiff, then the rule from the case could only be either $\{F^p_1, F^p_2\} \rightarrow$ Plaintiff, or $\{F^p_1\} \rightarrow$ Plaintiff, or $\{F^p_2\} \rightarrow$ Plaintiff. Rules cannot involve any factors opposed to the outcome. Rules in cases work to generate a weighing of sets of factors or an ordering of sets factors according strength: the rule, R, from a case establishes that that the set of factors in *Premise (R)* outweigh any set of opposing factors present in

[1] We might add that the application of interpretation to facts occurs not only in decide the ultimate outcome of a case, but to decide each issue within the case, which then combine to produce the outcome. This would follow the four part approach found in (Bench-Capon 2023, sec 4; Collenette et al. 2023, sec. 2.1.4), where cases move from evidence to facts, facts to factors, factors to determinations of various issues, and then from those determinations to an final outcome.

the case. *Premise (R)* then provides a stronger reason for the outcome than any of the opposing sets provide for the contrary outcome.

To illustrate, suppose we have a case, $<F, R, O>$, where F: $\{F_1^p, F_2^p F_1^d, F_2^d\}$, R: $\{F_1^p\} \rightarrow$ Plaintiff, and O: Plaintiff. Then this case establishes that $\{F_1^p\}$ outweighs $\left\{F_1^d, F_2^d\right\}, \left\{F_2^d\right\}, \{F_1^d\}$. Using $>$ to denote the weighing of factors, we can say $\{F_1^p\} > \left\{F_1^d, F_2^d\right\}, \left\{F_1^p\right\} > \left\{F_2^d\right\}$, and $\left\{F_1^p\right\} > \{F_1^d\}$. On these models a set of factors is always stronger than its proper subsets and weaker than its proper supersets, hence $\left\{F_1^d, F_2^d\right\} > \left\{F_2^d\right\}$, and $\left\{F_1^p F_2^d\right\} > \{F_1^d\}$. This can be stated more perspicuously if we define a function "*LF*" which extracts from a case the set of all losing factors, i.e. factors present in that case that do not favor the outcome. We can define the connection between rules and the ordering of factors as follows:

Rules and Factor Ordering: If c is a case with rule, R, then *Premise (R)* $> LF(c)$

Given what we know about the ordering of factors, *Premise (R)* $> LF(c)$ entails that *Premise (R)* outweighs each member of the power set of losing factors, hence $\forall x(Premise\,(R) > x : x \in \mathbf{P}(LF(c)))$.

This ordering on sets of factors is how precedent constrains on these models. Judges in current cases are constrained insofar as their decision must be consistent with the weighing established in past cases. For example, suppose we have the past case, $<\{F_1^p, F_2^p F_1^d, F_2^d\}, \{F_1^p\} \rightarrow$ Plaintiff, Plaintiff$>$, and the factors in the current case are $\{F_1^p, F_2^d\}$. Can the current judge decide for the defendant? If she did, the rule she used would have to be $\{F_2^d\} \rightarrow$ Defendant. That would entail that $\left\{F_2^d\right\} > \left\{F_1^p\right\}$. But we just saw that the rule in our past case requires that $\left\{F_1^p\right\} > \left\{F_2^d\right\}$. Therefore, the judge cannot decide the current case for the defendant.

The models are able to capture the common law judicial practice of distinguishing, wherein a court seems to be bound by a past rule, but some aspect of the current case, a novel reason, marks it for different treatment. Often this involves the current court holding that some aspect of the current case is an exception to the rule from the past case. We will say that a rule, R, from a past case is triggered in a current case with factors, F, whenever the factors from the R's premise/antecedent are all present, i.e. when $Premise\,(R) \subseteq F$. A rule, R, is binding when it is triggered and $Premise\,(R)$ is stronger than all the opposing factors present in the current case according to the weighing established by past cases. A rule can be triggered but fail to bind in two ways: $Premise\,(R)$ is weaker than some set of opposing factors present in the current case, or $Premise\,(R)$ has not yet been weighed against some set of the opposing factors present in the current case.

To illustrate, consider two past cases, c_1, c_2. Let c_1 be $<\{F_1^p F_1^d\}, R_1 : \{F_1^p\} \rightarrow$ Plaintiff, Plaintiff$>$ and c_2 be $<\{F_1^p, F_1^d, F_2^d\}, R_2 : \{F_1^d, F_2^d\} \rightarrow$ Defendant, Defendant$>$. Suppose the current case has these factors: $\{F_1^p, F_1^d, F_2^d\}$. Here the rule from c_1, R_1, will be triggered but it will not bind because $Premise\ (R_1) \not> \left\{F_1^d, F_2^d\right\}$. In fact, the rule from c_2, R_2, tells us that $\left\{F_1^p\right\} < \left\{F_1^d, F_2^d\right\}$, hence $Premise\ (R_1) < \left\{F_1^d, F_2^d\right\}$. Here R_2 binds, because it is triggered and $Premise\ (R_2) > \left\{F_1^p\right\}$. The judge is constrained to hold for

the defendant. This is case where the rule of one past case is trumped by the binding rule of another past case.

Now suppose the current case has these factors: $\{F_1^p, F_2^p\ F_1^d, F_2^d\}$.Again, both R_1 and R_2 are triggered, but now neither one binds. R_1 is still trumped by R_2, but the weighing does not have either *Premise* $(R_2) > \{F_1^p, F_2^p\}$ or *Premise* $(R_2) < \{F_1^p, F_2^p\}$. The current court has discretion to determine the weighing between *Premise* (R_2) and $\{F_1^p, F_2^p\}$. The court could decide for the defendant, which would strengthen *Premise* (R_2) in future cases. Or it could decide for the plaintiff, determining that the presence of F_2^p creates a novel reason (F_1^p, F_2^p) justifying a different outcome in this case. If decided for the plaintiff, the current case would then have the rule, $\{F_1^p, F_2^p\} \rightarrow$ Plaintiff, and establish $\{F_1^p, F_2^p\} > \{F_1^d, F_2^d\}$ in the ordering of factors. The current court crafts a new rule that trumps R_2 on the basis of the novel reason present in the case. This is distinguishing on the factor-based models.

From jurisprudential perspective we can consider whether to think of distinguishing as a matter of interpretation, i.e. as interpreting (or perhaps re-interpreting) a past case to involve a rule with an exception that exempts the current case, or as a matter of application, i.e. as deciding not to apply the rule from a past case because of the presence of a novel reason and introducing a new rule involving that novel reason. Horty (2015) showed these two approaches are formally equivalent, but there is a slight caveat. The language used in his "Raz/Simpson conditions," as well as in Raz and Simpson (1961, 2009, p. 189), is that of modifying the rule of the past case. We must be careful not to take that literally, because that would allow a lower court in a hierarchy to create rules for courts above its level, which is typically forbidden, and because standard citation practice is to cite the distinguishing case for the relevant exception to the rule, whereas you would simply cite the original case if the rule was actually modified in that case (Rigoni 2021, pp. 215–6). The idea then is not that the rule is literally modified in the past case, but that the past case is re-interpreted as containing a rule with the relevant exception.

Consider the U.S. Supreme Court case of *United States* v. *Robinson*, which held that police have "categorical" authority to search the contents of arrestee's pockets.[2] In a later decision, *United States* v. *Wurie* the First Circuit Court of Appeals rejected a "literal reading of the *Robinson* decision" and held that the government could not search a cellphone found in an arrestee's pocket, finding (or creating) a cell phone exception to the *Robinson* rule.[3] The Supreme Court not only upheld the First Circuit holding, but favorably quoted it in another opinion wherein it reversed a California court that held that *Robinson* permitted a search of a pocketed cellphone.[4]

This could be cast as a fairly standard instance of distinguishing on factor-based models, with the First Circuit finding a novel reason (cellphones, which contain reams of personal information, simply didn't exist in the 1970s) that justified trumping the applicable rule from *Robinson*. Per Horty (2015), you could equivalently

[2] United States v Robinson 414 U.S. 218, 225 (1973). The example is discussed as part of an example of "narrowing" by Re (2016, p. 954).

[3] United States v. Wurie (2013), 728 F.3d, 3–7.

[4] See Riley v. California, 573 U.S. 403 (2014), in which the Supreme Court consolidated *United States* v. *Wurie* and *Riley* v. *California.* See discussion in Re (2016, p.955).

understand this as an instance of interpreting the rule from *Robinson* as being inapplicable to cellular phones.

However, the language of the First Circuit's opinion only supports the interpretative understanding. They do not write that *Robinson*'s rule applies but is then trumped by a novel reason or fact, and they certainly do not claim to be modifying the rule from *Robinson*. Rather, they argue that properly understood, *Robinson* is not applicable in the context of cellular phones. This interpretative framing of the argument is probably familiar to every lawyer that has written a brief—unless the case is extremely on point, the first argumentative move is to argue that the cases with adverse outcomes should be understood as inapplicable.

We might think that these approaches amount to six of one and half dozen of the other. After all, what does it matter whether we call it interpretation or application? But it does matter. If the *Wurie* court is understood as just interpreting *Robinson*, then that part of their reasoning is in step (2) and hence should not be modeled by the cluster of theories under discussion. If we stray from the language of the opinion, however, we can model it as part of step (3) as distinguishing. The outcome may be the same, but the explanation and justification for it will not be.[5] What we have is a trade-off between fidelity to the language of the opinion and the range of reasoning explained by the model. In the next section, I discuss more detailed examples of this issue.

2 A bramble bush: distinguishing, narrowing, and overruling in vertical precedent in U.S. federal courts

Broughton (2019) provides a number of improvements for the factor-based model to account for the details of the hierarchical structure of the U.S. federal court system. Of particular interest, he provides a evidence for the claim that lower courts in the U.S. judicial hierarchy are not permitted to distinguish higher court precedent. He writes,

The U.S. doctrine of horizontal precedent and the U.S. doctrine of vertical precedent are significantly different, even once the possibility of overruling is set aside. In particular, while a U.S. court faced with an applicable horizontal precedent that it is unwilling or unable to overrule is often permitted to either apply the precedent or distinguish it, a court faced with an applicable vertical precedent is required to apply the rule and reach the outcome it prescribes (Broughton 2019, p. 262).

He provides a number of supporting statements showing that lower courts cannot refuse to apply an applicable precedent, nor may they modify such a precedent. This leads him to advocate for treating higher court precedent as strict rules, i.e. rules that bind whenever they are triggered.

[5] This might matter a great deal. In the next section we will see that, if Broughton (2019) is correct, then an explanation in terms of distinguishing may characterize the court as having made an illegitimate decision while an explanation in terms of interpretation characterizes the court as having acted legitimately.

Now this claim about vertical precedent may come as somewhat of a surprise, since lower court cases that say they are distinguishing higher court precedent are easy to find. Broughton is well aware of this, and explains that distinguishing can have two different senses: a weaker or thin sense according to which "one distinguishes a precedent simply by showing that the associated rule does not apply to the present dispute," and a stronger or thicker sense, according to which distinguishing requires that precedent must be first recognized as applicable and then modified to carve out an exception for the current case (Broughton 2019, p. 256). These senses of distinguishing correspond to "tame" and "strong" distinguishing found in Raz (2009, p.185), which are also discussed in the context of AI and Law in Carey (2013). In the factor models, thin "distinguishing" is just showing that the relevant rule is not triggered by the facts of the current case, while the thick sense of distinguishing is what the models just call "distinguishing." It involves acknowledging that the rule is triggered but identifying a novel reason that allows the court to reach a different outcome.

Re (2016, p. 228–9) offers a helpful taxonomy of ways of using precedent,[6] under the assumption that a past case has a "best reading."[7] For him, following consists in applying a precedent when it is best read as applying, while extending a precedent consists in applying a precedent to a case when under the best reading it does not apply. Distinguishing consists in not applying a precedent where it is best read as not applying (Broughton's thin sense of the term), and narrowing consists in interpreting a precedent as not applying when it is best read as applying. Narrowing also seems to closely fit Broughton's thick sense of distinguishing, except the court need not (and usually does not) admit that the past case is "best read" as applying. The taxonomy is summarized in the table below.

Following: precedent is best interpreted as applying and precedent is so applied
Extending: precedent is best interpreted as not applying, and precedent is still applied
Distinguishing: precedent is best read as not applying, and precedent is not applied
Narrowing: precedent is best read as applying, and precedent is still not applied

The notion of narrowing is especially relevant here. Like the factor-based models, Re assumes that there is a definitive best reading of a past case. However, in his view courts, even lower courts, sometimes do legitimately decide not to apply that best reading when the precedent can be reasonably interpreted as having a narrower holding (a more specific rule) that would not govern the present case. To be eligible for narrowing, the precedent must be "ambiguous" insofar as it admits of a number of reasonable interpretations. Narrowing for Re requires at

[6] The idea that U.S. judges will intentionally read past precedent in ways favorable to their preferred result is hardly new. It is one of the points emphasized (perhaps too much so) by the Legal Realists, such as Holmes (1897), and later by legal pragmatists (Posner 2008) and in social science approaches to the U.S. common law (Cross and Tiller 1998).

[7] He notes that the assumption of a best reading is controversial and subject to qualification (Re 2016, p. 928).

least a tentative belief that the precedent was incorrect, potentially in result but at least with respect to the rule created, and confidence that the narrowed precedent does not contradict background legal principles (Re 2016, p. 935). There are further additional requirements for lower courts to narrow higher court opinions, which depend on one's theory of *stare decisis* (Re 2016, p. 936). The discussion of *Robinson* and *Wurie* above is taken from the examples Re provides of this practice (Re 2016, pp. 951–71), and that discussion shows that narrowing seems very similar to thick distinguishing.

Re explains partial overruling as differing from narrowing insofar as partial overruling does not require an ambiguous precedent (Re 2016, p. 935). When engaging in partial overruling, the court acknowledges that the past precedent applies and then overrules a part of that precedent. A court can partially overrule even if a precedent is unquestionably applicable. Partial overruling requires a confident belief that the precedent case was wrongly decided (either in result or with respect to the specific rule) and the presence (perhaps) of a special justification for over-ruling.

We now have two concepts, partial overruling and narrowing, which seem to overlap with the notion of distinguishing in the factor models. However, partial overruling is verboten for lower courts dealing with higher court precedent, while narrowing is not. Complicating matters further, Broughton acknowledges,

> when the holding of a precedent case is genuinely indeterminate, it seems plausible that the doctrine of precedent grants even lower courts discretion to "interpret" the case in whatever way seems best, within the range of indeterminacy left open by the higher court. And the lower court may, if application of the precedent rule to the present dispute seems unwise, read the precedent more narrowly than it might have, so that the rule does not reach the present dispute…[The court doing so has not] distinguished the precedent in the thick sense either, since it has not necessarily narrowed the scope of the rule handed down by the higher court. Instead, it has resolved an ambiguity in that court's formulation of the rule (Broughton 2019, p. 278).

It's not quite clear if this process is simply Re's narrowing in different terms. Broughton might be thinking of "genuinely indeterminate cases" as just those cases that lack a "best interpretation," which makes this process distinct from Re's narrowing. Or, he might be thinking that "genuinely indeterminate cases" are all the cases that have multiple different reasonable interpretations, which is all the ambiguity that narrowing requires. Broughton argues that this interpretative process, however understood, must be exogenous to the models under discussion because those models assume the rule of the case is always clear (Broughton 2019, p. 278). The case law seems to compound the confusion, as Re points out that courts often portray decisions to narrow as decisions not to extend, because they present the narrowed reading as if it were the best (or only) reading (Re 2016, n. 154). This can be seen in the *Wurie* example above, where the court simply discusses the proper understanding of *Robinson* in the context of related decisions on searches incidence to arrest.

To begin to untangle this bramble bush, we should observe that, functionally, distinguishing in the factor-based models is similar to narrowing. This becomes clear from considering some examples of narrowing:

- In *Simmons*[8] the Missouri Supreme Court heard a capital punishment case that seemed to be governed by the Supreme Court decision in *Stanford* v. *Kentucky*.[9] Essentially, the Missouri Court read *Stanford* as announcing a "contingent rule—a legal requirement whose continued application depends on the continued presence of the reasons for its creation in the first place" (Re 2016, p. 952).[10] Re notes that the "Missouri Supreme Court's view finds some support in a traditional conception of stare a decisis" (Re 2016, n. 141). *Simmons* was upheld by the Supreme Court despite dissent protests that the decision overruled *Stanford*.[11]
- Re discusses the aforementioned *Wurie* (728 F.3d at 11–12) in terms very familiar to formal modelers: "Concluding that *Robinson*'s stated rule could be read as defeasible, or subject to implied exceptions, the First Circuit declined to apply the Court's stated holding according to its own terms" (Re 2016, p. 954).[12]

The factor-based models' notion of distinguishing seems to significantly overlap with Re's notion of narrowing, even though the former does not involve multiple interpretations. This might explain why models that do permit lower courts to distinguish still seem to work well when applied to case bases with lower and higher court precedent, such as the U.S. trade secret database.[13] The models are capturing (perhaps imperfectly) what is described by cases and commentators as an interpretive process.

This suggests that Re's narrowing is just the factor-models distinguishing under a different name. Unfortunately, there are a few problems with identifying distinguishing in the models with narrowing. First, as Broughton (2019) has pointed out, this seems to ignore the language of the case law and commentary, which is fairly uniform in treating the process as interpretive. Narrowing depends on ambiguity in the past opinion. To illustrate, recall our example of distinguishing. We had the factors $\{F_1^p, F_2^p\ F_1^d, F_2^d\}$ in the current case with rules R_1: $\{F_1^p\} \rightarrow$ Plaintiff and R_2: $\{F_1^d, F_2^d\} \rightarrow$ Defendant, and the weighing $\left\{F_1^p\right\} > \{F_1^d\}$ and $\left\{F_1^d, F_2^d\right\} > \left\{F_1^p\right\}$ in the background. Let's further suppose that the current court is at a lower level of the U.S. federal judicial hierarchy than the courts for the background cases. On the factor-based models without Broughton (2019)'s modifications, the current court is free to distinguish R_2 in this context by introducing the new rule $\left\{F_1^p, F_2^p\right\} \rightarrow$ Plaintiff. On Re's view, you could narrow R_2 so that it would not apply in the current case; for example, you could narrow it to something like $\{F_1^d, F_2^d, F_3^d\} \rightarrow$ Defendant. where F_3^d is something incompatible with F_2^p. But it can only be narrowed in that way if

[8] State *ex rel.* Simmons v. Roper, 112 S.W.3d 397 (Mo. 2003).

[9] 492 U.S. 361 (1989).

[10] See also *Simmons*, 112 S.W.3d at 406, which comments on "flexible and dynamic" interpretation as a "fundamental premise" of *Stanford*.

[11] State *ex rel*. Simmons vs. Roper, 543 U.S. 551 (2005).

[12] But even here, notice the hedge between interpretation and application. "Implied exceptions" means the rule somehow contains the exceptions, but that's not required for the rule to be defeasible.

[13] See Ashley and Brüninghaus ((2009), Al-abdulkarim et al. (2016a). Of course, this isn't the only possible explanation.

the language of the past case allows that. The language in c_2 might be such that R_2 can only be understood as a strict rule, which would mean it cannot be narrowed and will bind. This would force the current court to decide for the defendant.[14]

This difference in focus on the language of the past opinion may also give rise to differences in how current judges may modify past rules, although much depends on the particular factor-model selected. Suppose we have a past case best understood as introducing the rule that "no dogs are allowed on the bus" and the current case involves a service dog. Is the current judge able to free to introduce an exception for "service dogs"? For both the factor models and Re, the answer depends on the past opinion. For some factor models,[15] precedents may only be represented with factors that were actually present in the past case. On these models, the only question is whether the past case involved a service dog (and if this was clear from the opinion). If so, the judge cannot make the exception; if not, she is free to do so. For Re, the issue of fit with the past opinion makes the matter more complicated. For example, the past opinion might state, "no dogs, not even service dogs, are allowed on busses," even if the dog in question was not a service dog. For the aforementioned factor-models, this mention of service dogs is irrelevant because it goes beyond the facts of the case (and there may be good jurisprudential reasons for thinking this way). However, there is no reasonable way to interpret that language as allowing service dogs on the bus, so narrowing will not be permitted in this instance but distinguishing would be possible. Other factor-based models (call them permissive factor models) permit the use of hypothetical factors in representing cases.[16] These models would permit treating the past case as having considered service dogs, even if the past case did not actually involve service dog. Hence, they would not permit distinguishing in this situation in line with the standards for narrowing, since the service dog would not be a novel factor in the current case.

However, Re and the permissive factor models will still come apart with respect to the type of changes they permit the current court to make. The factor models only allow adding an exception to the past rule for a novel reason in the current case. Essentially, there is a general rule plus a more specific exception: dogs are not allowed on the bus, except service dogs are. Distinguishing leaves no open ground;

[14] Under Broughton (2019)'s modifications R_1 and R_2 they are strict rules, not in virtue of the particular language used in the cases, but in virtue of their coming from a court higher up in the judicial hierarchy. That means any of the high court rules that are triggered and not trumped by other background rules will bind. Under Broughton's modifications, R_2 will bind here and the court must find for the defendant.

[15] E.g. (Horty 2011).

[16] (Rigoni 2015, 2018) are models of this kind. A different, but still permissive approach found can be found in the original HYPO (Rissland and Ashley 1987) as well as in more recent work (Bench-Capon and Atkinson 2022), which would be treat the case as introducing a dimension with respect to permissibility of different types of dogs. While (Rigoni 2018) represents dimensions as sets of factors, some of which are not present in the case, these other models represent dimensions as orderings of individual facts from which one can derive the presence of a factor. These models could represent the past court's statement that "no dogs, not even service dogs, are allowed on buses" as establishing that the fact (rather than factor) that the dog is a service dog does not permit the derivation of a factor that favors the dog's owner. I still consider these models permissive because they permit hypothetical facts (the past case doesn't involve a service dog) to generate binding precedent.

for cases that fell within the original rule, they will all either fall into the general rule (no dogs) or the exception (service dog). Narrowing countenances a wider range of modifications of the past rule. The current court could narrow the "no dogs on the bus" rule from the past case to "no untrained dogs on the bus," but have the holding of the current case be "service dogs are permitted." Here the exception at issue in the case is narrower than the modification to the rule. Cases involving dogs that are trained but not service dogs, which were covered by the original rule, are now left open.[17] Narrowing permits all the same modifications as distinguishing and more, so narrowing does not entail distinguishing.

To see this formally, let the plaintiff be a bus company suing for damages that they attribute to the dog owned by the defendant, who bought a ticket and was riding the bus with his dog. Let F_1^p be a factor for the presence of the owner's dog on the bus and F_1^d be the factor for the rider having bought a ticket. Then let c_1 be as it was previously, i.e. $<\{F_1^p\, F_1^d\}, \{F_1^p\} \rightarrow$ Plaintiff, Plaintiff$>$. This case introduces the "no dogs on the bus" rule. Let F_2^d be the factor for the dog being a licensed service dog, and let the current case have this set of factors: $\{F_1^p, F_1^d, F_2^d\}$. On the factor model approach, the current court can distinguish R_1 by introducing the new rule R_2: $\{F_1^d, F_2^d\} \rightarrow$ Defendant. If it does so, R_1 will still be triggered and potentially bind whenever F_1^p is present in a future case. In future cases involving service dogs (and riders who bought a ticket), R_1 will be trumped by R_2, as these cases fall within the exception created by F_2^d. For example, in a future case with a dog that is trained but not a service dog, R_1 will be triggered although possibly distinguishable. On the other hand, narrowing would involve the current case re-interpreting R_1, changing it to $R_{1'}$ $\left\{F_1^p, F_2^p\right\} \rightarrow$ Plaintiff, while also introducing the rule $\left\{F_1^d, F_2^d\right\} \rightarrow$ Defendant. Horty (2015) shows that these approaches are equivalent if $F_2^p = \neg F_2^d$. In our example, if F_2^p is the factor for the dog *not* being a service dog, then narrowing will be equivalent to distinguishing. In a future case with a dog that is trained but not a service dog, $R_{1'}$ will still be triggered. But narrowing does not require that $F_2^p = \neg F_2^d$ in these instances. F_2^p could be the factor for the dog not being trained, for example, so it would entail the dog not being a service dog, but not be equivalent to it. Then in a future case with a dog that is trained but not a service dog, $R_{1'}$ as re-interpreted, will not be triggered at all and the judge is entirely unconstrained in formulating a rule to govern the case.

[17] Re considers further kinds of rule modification when he discusses Scalia's claim that the Supreme Court distinguished a previous case on an "accurate-in-fact (but inconsequential in principle distinction)," which sounds very much like distinguishing based on a difference in fact without any corresponding difference in reasons or factors (Re 2016, p. 935 (quoting Arizona State Legislature v. Arizona Independent Redistricting Committee, 576 U.S. 857 (2015) (Scalia, J., dissenting))). This would be illegitimate as distinguishing on any factor model, but Re says it would "qualify as legitimate narrowing" if it was based on a reasonable reading of the past precedent and violated no background principles. I do not follow Re on this point. Scalia seems to be criticizing the court for drawing such a distinction, not legitimating such behavior. He argues that *if* such a distinction was a legitimate basis for distinguishing, then the current case could also be so distinguished. However, he thinks the court is wrong to draw such distinctions. See Arizona State Legislature v. Arizona Independent Redistricting Committee, 576 U.S. 857-8 (2015) (Scalia, J., dissenting).

Thus, narrowing is not quite distinguishing in the factor models. Cases that can be distinguished but not narrowed may be rare, unless we stipulate, following Broughton (2019), that lower courts cannot distinguish higher court rulings. Thick distinguishing might seem more like partial overruling. The language associated with partial overruling is a better fit, as it acknowledges that the precedent rule applies to the current case. However, thick distinguishing is not equivalent to partial overruling. Partial overruling, just like overruling, does not require a novel reason (or relevant difference) in the facts of the current case, while distinguishing requires a novel reason (or relevant difference). Partial overruling requires a special justification, but the special justifications for overruling are reasons outside the specific facts of the current case. For example, the quality of reasoning within the precedent opinion and the degree to which the precedent case is relied upon are special justifications for overruling,[18] but neither of these have anything to do with the facts of the current case. Hence partial overruling will be legitimate in some cases where distinguishing is not (where there is no relevant difference but there is a special justification), and partial overruling will be illegitimate in cases where distinguishing would be (where there is a relevant difference but no special justification).

From the perspective of the factor-models, you could accommodate partial over-ruling, but it won't be a substitute for distinguishing. To get partial over-ruling, you need to model cases as having multiple holdings, as (Rigoni 2015; Bench-Capon and Atkinson 2021) suggest. Partial over-ruling is then when a court over-rules some but not all of the holdings in a case (Re 2016, p. 929). This differs from distinguishing, which involves the distinguished rule being trumped by a new rule. Over-ruling involves removing the past rule entirely. As mentioned above, distinguishing requires a novel reason in the facts of the current, which is not required for over-ruling. Suppose the past case involves two holdings, one of which uses R_1: $\{F_1^p\} \rightarrow$ Plaintiff, and this rule established the weighing $F_1^p > F_1^d$. Suppose the current case has this set of factors: $\{F_1^p, F_1^d\}$. R_1 cannot be distinguished in this context, but it could be over-ruled, which removes R_1 from the background for future cases, and then the current court could use R_2, $\left\{F_1^d\right\} \rightarrow$ Defendant to reach its holding. So long as the rule for the other holding from the past case is not over-ruled, this will be partial (not total) over-ruling.

3 Incorporating interpretation for narrowing

From the perspective of a modeler working with factors, a few options are available to try to incorporate narrowing. One could put a premium on the language used in opinions and treat what is described as interpretation as interpretation, meaning narrowing is excluded from the model. Broughton advocates for a form of this by making it impossible for lower courts to distinguish higher court cases within the model, but allowing courts to distinguish cases from courts at or below their place in

[18] See Janus v. American Federation of State, County, and Municipal Employees (2018), 138 S.Ct. 2472-4 (discussing special justifications for overruling).

the hierarchy. He acknowledges that lower courts resolve ambiguous interpretations of higher court opinions, but stresses that this is explicitly interpretative and hence should be outside the model (Broughton 2019, Section 6).

Broughton's model retains thick distinguishing for horizontal precedents, but one could go even further and simply eliminate thick distinguishing from the model. It's not clear that courts declining to follow horizontal precedent do anything that is neither narrowing nor (complete or partial) over-ruling. Re, for example, thinks that extending, following, narrowing (both horizontally[19] and from above and below), and complete and partial over-ruling (only horizontally and from above) exhaust the ways US federal courts use precedent. Although thick distinguishing differs from both partial over-ruling and narrowing, it's not clear that a model would need distinguishing if it already had both partial over-ruling and narrowing. Distinguishing would only be needed to capture cases where it is very clear the court (i) understood the past precedent as providing an applicable rule (so language of the opinion will not fit with narrowing), (ii) was not using a special reason to over-rule that precedent, and (iii) declined to follow that precedent on the basis of a novel reason in the facts of the current case. In practice such cases may be few and far between.

Jettisoning distinguishing provides (plausibly) a good fit with the language of opinions and a greatly simplified model of precedent. But we pay a price for this in terms of explanatory power. The more reasoning that gets put into the interpretative step of the process, the less reasoning the model itself explains. And case interpretation, which is itself a controversial and specialized mode of language processing, is very difficult to model. I will sketch an outline of a model for interpretation shortly, but working the out the details will take some time.

These are good reasons to further develop models of interpretation. But they are also reasons to consider retaining distinguishing in the models even if that requires deviating from the language of the caselaw, which leads to the second option for the modelers. Narrowing itself already requires deviating from (or going beyond) the language of the caselaw—few (if any) judges will ever admit in an opinion that they know of a "better" reading of a past case than the one they are providing. If we tolerate a bit more deviance, we can model a good deal of narrowing in the standard factor models– recall the substantial overlap between narrowing and thick distinguishing. Depending on the goals of the model, this may strike a favorable trade-off between adherence to the language of the cases and explanatory power. Still, we must be clear that we are making a trade-off in taking this route.[20]

I promised a sketch of a model of interpretation to accommodate narrowing within a factor-based framework. Here is the idea: currently, the factor models begin with each past case having an exogenously provided interpretation, which is a single formalization in terms of factors, dimensions, rules, etc. For narrowing, we need to allow each case to have multiple such interpretations (formalizations). What used

[19] See Re (2014).

[20] Ignoring the purposes of models can cause much confusion. It leads to "toy" models being presented as accurate depictions, see Pfleiderer (2020).

to be case becomes an interpretation of a case, and cases are now sets of interpretations. To be clear, this, in effect, exponentially increases the size of the case base.[21] The exogenous process of interpretation needs to provide a range of reasonable interpretations, and identify one as the best or most reasonable. Lower courts may not distinguish in the thick sense; all higher court rules are strict so if one applies and is not trumped by another higher court rule then it binds the lower court. But all courts can pick any one of the reasonable interpretations of a past case. Narrowing occurs when the best interpretation applies in the current case, but the court picks an alternative interpretation that does not apply.[22] Extending occurs when the best interpretation does not apply and the court picks an alternative interpretation that does apply. Following is when the best interpretation does apply and the court selects this one. If a court is not allowed to distinguish the case, then it is bound by precedent on this view when all the reasonable interpretations produce the same result. If the court may distinguish, then it is bound when all reasonable interpretations agree in result and none of them can be distinguished.

This approach is, I think, worth elaborating in future work, especially if the production of interpretations can be automated,[23] but it has a number of shortcomings. Most obviously, it seems very difficult to generate a set of reasonable interpretations and to pick the best of these. The model also needs to manage a possibly very large set of interpretations.[24] Further, it has to reckon with the precedential status of interpretations of past cases. When a court narrows or extends a precedent, that reading of the precedent becomes precedential. If the court narrows the rule from a case from "No dogs on the bus" to "No non-service dogs on the bus," then that narrowed rule is precedent for lower courts. Carey (2013) calls these attribution holdings, and provides a general account. There the interpretation is a propositional variable or the negation of one, so much work would be needed to incorporate the more nuanced formalizations of cases found in the factor-based models. In addition, it's not clear how to manage the reflexivity of interpretation: should we allow all past

[21] This approach is discussed in the context of analogical approaches, like (Stevens 2018), that allow for the current case to influence the interpretation of the past case directly in Rigoni (2021). In those analogical approaches, one past case may have different interpretations depending on the current case, so the casebase itself is relative to the current case. The approach here follows Re in treating interpretation as independent of the current case.

[22] It's straightforward to extend this for reasoning with parts of precedents, holdings for different issues, etc.; I ignore that complexity here.

[23] See the work on automatically classifying cases according to factors in Ashley and Brüninghaus (2009), Branting et al. (2021), Gray et al. (2022, Mumford et al. (2022). I suspect it would be beneficial to break opinions up into issues and generate interpretations for each of those portions of the opinion, but that may raise problems of coherence—the combination of best interpretations for each issue may not be as coherent as other combinations of interpretations.

[24] A subtlety: Re reads "the best interpretation" subjectively, i.e. narrowing only happens when the judge or judges self-consciously apply a reading that they know is not the best (Re, 2019, n.37). Nonetheless, Re points to objective evidence to draw inferences about the judge's subjective states. It is not clear that this is better than understanding "best interpretation" objectively, which would allow narrowing to capture instances of confirmation bias, wherein judges unconsciously adopt inferior readings because they support their favored outcome. Since the rule of the past case is effectively narrowed whether the judge knows it or not, the objective characterization seems superior in this context.

interpretations to themselves be reinterpreted?[25] Can narrowings be narrowed? Can they be extended? These are the sort of problems that will need to be ironed out if we are to both model a good deal of U.S. Federal Court precedent and hew tightly to the language in opinions.

The point is not that any one of these approaches: jettisoning distinguishing and leaving narrowing outside the model, including distinguishing and using it to capture some of narrowing, or incorporating interpretations, is *the* right way to model the underlying process. Each approach prioritizes different desiderata. Jettisoning distinguishing prioritizes fit with the language used in opinions and computational tractability. Including distinguishing deemphasizes the language of opinions but retains computational tractability with increased explanatory power. Incorporating interpretations promises to combine good fit with the case language with wide explanatory power, but greatly (perhaps intractably) increases the complexity of the model. Of course, the different approaches can be adopted to varying degrees, such as excluding distinguishing for certain levels of judicial hierarchy. Ultimately, the right model depends on one's purposes and the availability of supplemental models. The tripartite division of judicial reasoning into fact-finding, interpretation, and rule-application is itself motivated by the practical concern to simplify the process rather than firm convictions that every aspect of the reasoning process *really is* interpretation, or application, or fact-finding. Hence, it seems proper to draw on pragmatic considerations when putting aspects of the judicial reasoning process into these categories, and modelers ought to make clear what considerations are motivating our divisions of the reasoning process.

This section has shown the potential benefit of incorporating interpretations on a large scale in factor-based models. However, it also suggests a pragmatic approach in determining when and to what degree to incorporate interpretations within the models. Pragmatically, it may be best to start by incorporating interpretations on a smaller scale, using them in more limited contexts where we can gain some explanatory power or fit with case language without paying too high of a cost. The next section discusses of how interpretation might be incorporated in this way to deal with dimensions in factor-models.

4 Dimensions in factor-models

Recent work has offered a number of proposals for treating dimensions within factor-models. A dimension is an aspect of a case that takes a value from a range of values, where those values can be ordered as more or less favorable for a particular party or conclusion. Taking an example from Horty (2017), consider a court trying to determine whether a taxpayer can claim a change of domicile to avoid paying

[25] In the same vein, Casey (2013, p.364) notes, "It also appears that in cases where an earlier opinion posits some holding h and a later opinion overrules the earlier opinion by positing $\sim h$, it may be possible to formulate a large number of hypothetical holdings that are implied by h (and may be extremely similar to h), but that are not contradicted by $\sim h$ and thus are not deemed to have been overruled.".

tax. Suppose two relevant dimensions are the length of time the taxpayer lived abroad and the percentage of their income that was earned abroad. The longer the time spent abroad, the stronger the case is for a change of domicile, and likewise for the higher the percentage of income earned abroad. The reverse holds for the case against change of domicile: longer (shorter) times and higher (lower) percentages weaken (strengthen) the case for that conclusion. Let us stipulate that the defendant wants a finding of a change of residence.

Horty (2017, 2019, 2021), Al-abdulkarim et al (2016b), Bench-Capon and Atkinson (2017, 2021, 2022), Rigoni (2018), and Prakken (2021) all discuss approaches to modeling dimensions in factor-models. Rigoni (2018) treats values on a dimension simply as factors, meaning that each value is a reason either for or against a particular conclusion (or party[26]). The remainder all avoid modeling values on a dimension as factors by either allowing some values to count neither for nor against a particular conclusion or by stipulating that each value creates a reason for and against a particular conclusion. These models are all motivated by a difficulty discussed in Bench-Capon (1999) and summarized by Prakken thus: "that in practice it is often hard to specify for some or even all value assignments to a dimension in a case whether it is for or against [a particular outcome]" (2021, p. 560). Instead of the binary approach of categorizing values as pro or con a given conclusion, the values are treated as more or less supportive of a particular conclusion. On many of these models, such as Horty (2019) and Prakken (2021)'s alternative dimensional based approach, we cannot say that a specific value on a dimension, such as 90% of the defendant's income being earned abroad, is simply pro-change-of-residence (i.e., pro-defendant). Rather, that 90% point is understood as more pro-change-of-residence than lower percentages and less pro-change-of- residence than higher percentages. Prakken, anticipating some of the problems discussed below, suggests, "allowing but not requiring [the model] to indicate that a particular value assignment favours a particular outcome," i.e., allowing but not requiring values on a dimension to be factors (2021, p. 580). Similarly, Horty (2021) seems to build in a method of treating some dimensions as de facto factors.

The difficulties with modeling dimensions can be illustrated by an example.[27] Suppose we have a case (Case 1) holding that there is a change of residence when the defendant spent 18 months living abroad and earned 60% of her income abroad. This case could mean either of these three things, assuming income and length abroad are dimensions that relate to distinct factors[28]:

(A) 60% income earned abroad is a reason for change that outweighs 18 months abroad (a reason against change).
(B) 18 months abroad is a reason for change that outweighs 60% of income earned abroad (a reason against change).

[26] The model assumes a pro-plaintiff or pro-defendant holding for each issue, and then labels the factors favoring a holding according to which party the holding favors. I ignore this complication.

[27] A similar example is given in Prakken (2021, p. 577).

[28] Dropping this assumption is critical to the Bench-Capon and Atkinson approach from 2017 onward (Bench-Capon and Atkinson 2017, 2021, 2022), which is discussed later in this section.

(C) 18 months abroad is a reason for change and so is 60% of income earned abroad.

Bench-Capon (1999) uses the language of dimensions "trading off" with one another to describe what is happening with the two dimensions here, while Bench-Capon and Atkinson (2017) use the language of "balancing dimensions." I prefer describing the dimensions as "entangled": you cannot tell which dimension is doing what, but you know what the result of both is.[29]

Approaches that reduce dimensions to factors, like Rigoni (2018), require disentangling dimensions. In fact, the reason model underlying Horty (2019) requires disentangling dimensions to some extent, as Prakken (2021) shows. The reason model requires that a ratio contains only pro-plaintiff factors or dimensions. Hence, if the ratio in the Case 1 included both the duration of the stay abroad and the percentage of income earned abroad, both of those would have to be pro-change-of-residence factors and the case could only have meaning (C). If only the 18-month duration was in the ratio, then the case would be ambiguous between (B) and (C); if only the 60% of income was in the ratio, then the case would be ambiguous between (A) and (B). Reductive approaches require making a choice about how to interpret that past case to disentangle the factors, hence they are "arguably harder to apply in practice" (Prakken 2021, p. 581).

The reductive approaches leave more work for exogenous interpretation, while the alternatives do some of that work themselves. By allowing the dimensions to remain entangled, the non-reductive approaches do not require as much interpretative work. This gives them a significant practical advantage over the reductive approaches. Suppose the current case is Case 2, in which the defendant has spent 19 months abroad and earned 65% of his income aboard. Although the non-reductive approaches differ considerably, they will all find that Case 1 forces a finding of change of residence in Case 2 without requiring a decision between (A), (B), or (C) as interpretations of Case 1.[30]

In addition, the non-reductive approaches might be superior when the number of dimensions changes between the first and second case. Suppose our first case is Case 1 (18 months abroad, 60% of income aboard, finding of change of residence) and the current case doesn't have information on the percentage of income earned aboard, it just has that the defendant live abroad for 17 months. If you disentangle the dimensions, you *can* get Case 1 to bind the current case: on reading (A) 18 months is a reason against change, so 17 months opposes change even more strongly, hence the result in a case with just that factor must be no change of residence.[31] But, depending on how the opinion in Case 1 is written, that can seem like the wrong result. It might not be clear just how the 18 months is supposed to be understood. Here it is

[29] "Entanglement" in a somewhat different sense was used in (Roth and Verheij 2004).

[30] The rest of this section makes use of a number of hypothetical cases. Please see Table 1 for a summary of each hypothetical case.

[31] Complications may arise from whether to treat the lack of information on percentage of income earned as a factor. For this example, I'm assuming the only factor in the current case is the 17 months stay.

not just that the reductive approach is harder to apply, but that the best strategy is to leave the dimensions entangled.

However, in other instances disentangling seems like the right strategy.[32] Suppose our first case is Case 3, with a 50-month duration, 0% income earned abroad, and holding for a change of domicile. Suppose our second case is Case 4, with 50 months abroad and no information on income earned abroad. Case 3 intuitively means that 50 months or more is sufficient for change of domicile, regardless of income earned. But without disentangling the two dimensions, we cannot get that reading, and Case 3 cannot bind Case 4.

Similarly, suppose the past case is Case 5 (18 months abroad plus a holding for a change of domicile) and the current case (Case 6) has an 18-month duration and 2% income earned abroad. Suppose the court decides against a change of domicile. A natural reading of Case 6 is that it distinguishes Case 5 on the basis of a reason against a change of domicile (2% income earned abroad). But you cannot give that reading if you cannot characterize 2% of income earned abroad as reason against a change of domicile, which requires disentangling it from the 18-month duration.

The prevailing non-reductive approaches to these difficulties are Horty (2019, 2021)'s magnitude factors and Bench-Capon and Atkinson (2017, 2021, 2022)'s two-stage approach. Considering Horty's approach first, his magnitude factors state that the actual value of dimension in the case favors a side at least as much as reference value. Magnitude factors favor a particular side, but the underlying values on a dimension in the case create magnitude factors favoring each of the opposing sides. For example, a stay of 12 months will generate a pro-defendant magnitude factor of "a stay of 12 months or more" and a pro-plaintiff magnitude factor of "a stay of 12 months or less." The court may also decide to use a lower reference value, so a 12-month stay could generate a pro-defendant magnitude factor of "a stay of 6 months or more" and a corresponding opposite magnitude factor. Notably, a value on a dimension could function as a "de facto factor" by creating both a pro-defendant magnitude factor and a pro-plaintiff magnitude factor such that one has priority over (or outweighs) the other.

In Horty's most recent approach in Horty (2021), the ratio can only contain a set of magnitude factors such that each one favors the prevailing party and the set outweighs the magnitude factors for the other party. Hence, dimensions can be disentangled: if the ratio only contains one magnitude factor, then that magnitude factor must outweigh its corresponding opposite magnitude factor, and that value on the dimension is a de facto factor. Horty denies "that there is a switching point [where values switch from favoring one side to favoring another], supposing instead that reasons favoring each side of a dispute can be found all along any given dimension, varying only in strengths" (2021, p. 284). However, a single fact that produces opposite magnitude factors such that one of them (call it pro-*s* magnitude factor)

[32] Prakken (2021) points out some similar issues with his approach when dealing with dimensions not mentioned in the ratio of the case. He notes that you can avoid those by essentially treating every dimension of the past case as part of the ratio and only allowing distinguishing in favor of side *s* if the value on some dimension in the current case is more favorable to *s* than the value on that dimension in the past case. That is, you can avoid the problem by using a results model approach with dimensions. However, that approach still will not work when the number of dimensions in each case varies.

outweighs the other (con-*s* magnitude factor) seems very close to a factor, and a point on dimension above or below which all values produce these kinds of pairs of magnitude factors seems very close to a switching point.

Formally, Horty (2021, p. 274) states, "where p is some value along the dimension d, we define a magnitude factor favoring the side s as a statement of the form $M^{s}_{d,p}$, meaning the actual value assigned to the dimension d favors the side s at least as strongly as the reference value p." Consider Case 1, which involves four magnitude factors. First, there is $M^{def}_{d1,18}$, meaning the value assigned to dimension d_1 favors the defendant at least as strongly as 18 months abroad. Then there is $M^{pla}_{d1,18}$, which is the magnitude factor meaning the value assigned to dimension d_1 favors the plaintiff at least as strongly as 18 months abroad. Next there is $M^{def}_{d2,60\%}$ meaning the value assigned to dimension d_2 favors the defendant at least as strongly as 60% of income earned abroad. Finally, there is $M^{pla}_{d2,60\%}$, meaning the value assigned to dimension d_2 favors the plaintiff at least as strongly as 60% of income earned abroad. So, each value on the two dimensions (18 months and 60% of income) generates two magnitude factors, one for each side in the case.

The ratio in a case establishes a weighing for *some* of the magnitude factors in the case, in the same way ratios do that in non-dimensional factor models. The factors in the antecedent of the ratio, i.e. *Premise* (r), all favor the prevailing side and that set outweighs all the opposing factors. Consider these potential ratios for Case 1, where the defendant prevails[33]:

Rule 1: $\{M^{def}_{d1,18}\} \rightarrow$ Defendant

Rule 2: $\{M^{def}_{d2,60\%}\} \rightarrow$ Defendant

Rule 3: $\{M^{def}_{d1,18}, M^{def}_{d2,60\%}\} \rightarrow$ Defendant

Rule 1 entails that $\left\{M^{def}_{d1,18}\right\} > \{M^{pla}_{d1,18}, M^{pla}_{d2,60\%}\}$, and as in the non-dimensional factor models, if a set of factors, S_1, outweighs 1 another set, S_2, then S_1 outweighs every subset of S_2. Hence $(S_1 > S_2) \rightarrow (S_1 > S\colon\ S \in \boldsymbol{P}(S_2))$. Hence $\left\{M^{def}_{d1,18}\right\} > \left\{M^{pla}_{d1,18}\right\} M^{def}_{d1,18} > \{M^{pla}_{d2,60\%}\}$. This ratio seems to entail that the fact that the stay is at least 18 months acts like a traditional pro-defendant factor (is a de-facto factor), since it introduces a pair of magnitude factors such that the pro-defendant one is always stronger than the pro-plaintiff one: $\left\{M_{d1,18}\right\}^{def} > \left\{M^{pla}_{d1,18}\right\}$. Rule 1 won't tell us anything about the comparative weight between $M^{def}_{d2,60\%}$ and $M^{pla}_{d2,60\%}$. Rule 2 will tell us that $\left\{M^{def}_{d2,60\%}\right\} > \{M^{pla}_{d2,60\%,} M^{pla}_{d1,18}\}$ and hence $\left\{M^{def}_{d2,60\%}\right\} > \left\{M^{pla}_{d2,60\%,}\right\}$, making 60% or great percentage of income earned abroad a de-facto factor. Rule 2 leaves the weight between $\left\{M^{def}_{d1,18}\right\} and \left\{M^{pla}_{d1,18}\right\}$ undetermined. Rule

[33] Many more options are available. Without constraint from background cases, the court can set a reference value of the factor in the ratio as lower than the value that occurs in the case. That is, the court could use $M^{def}_{d1,12}$ in these ratios, because the value of 18 months would support the defendant as least as much as the value of 12 months.

3 tells us nothing about the weightings of either pair of magnitude factors, it only tells us that $\{M^{def}_{d1,18}, M^{def}_{d2,60\%}\} > \{M^{pla}_{d1,18}, M^{pla}_{d2,60\%}\}$. On Rule 3 the factors remain entangled. Constraint is achieved by requiring that the ratio of the current case is consistent with the ordering of factors/magnitude factors from past cases. Distinguishing is permitted when there is a novel reason, as in the non-dimensional versions.

Horty's (2021) approach can easily get Case 1 to bind Case 2, once we select a ratio for Case 1. Let the ratio for Case 1 be Rule 1. Then $\left\{M^{def}_{d1,18}\right\} > \{M^{pla}_{d1,18}, M^{pla}_{d2,60\%}\}$. But Case 2 satisfies $M^{def}_{d1,18}$ because a 19-month stay is at least as favorable to the defendant as an 18 month stay. So, Rule 1 is triggered. Further, there is no novel reason to permit distinguishing Rule 1. A 19-month stay is at most as favorable to the plaintiff as an 18-month stay, because longer stays are worse for the plaintiff, and 65% income earned abroad is at most as favorable to the plaintiff as 60%, because higher percentages as worse for the plaintiff. Rule 1 shows that $M^{def}_{d1,18}$ is stronger than a set of factors that is at least as strong in favor of the plaintiff as the set of all pro-plaintiff magnitude factors present in the case. Hence the court is bound to follow Rule 1.[34]

Horty's (2021) approach assumes complete information on dimensions in both the current and past cases, and there are difficulties in extending it cases with incomplete dimensions. It can get Case 3 to bind Case 4, if it gives Case 3 the ratio $\{M^{def}_{d1,50}\} \rightarrow$ Defendant, which means $\left\{M^{def}_{d1,50}\right\} > \{M^{pla}_{d1,50}, M^{pla}_{d2,0\%}\}$, where $M^{def}_{d1,50}$ is the magnitude factor favoring the defendant at least as much as a 50 month stay, $M^{pla}_{d1,50}$ is the magnitude factor favoring the plaintiff at least as much as a 50 month stay, and $M^{pla}_{d2,0\%}$ is the magnitude factor favoring the plaintiff at least as much as a 0% income earned abroad. Case 4 will trigger this rule and the only relevant pro-plaintiff factor will be $M^{pla}_{d1,50}$. Since we already know that that $\left\{M^{def}_{d1,50}\right\} > \{M^{pla}_{d1,50}\}$ from the background weighing, Case 3 compels a pro-defendant ruling in Case 4.[35]

The approach can get Case 6 to distinguish Case 5 as well. Let Case 5 have the ratio, $\{M^{def}_{d1,18}\} \rightarrow$ Defendant, where $M^{def}_{d1,18}$ is the pro-defendant magnitude factor for spending at least 18 months abroad. This induces the weighing $\left\{M^{def}_{d1,18}\right\} > \left\{M^{pla}_{d1,18}\right\}$, where $M^{pla}_{d1,18}$ is the pro-plaintiff magnitude factor for having spent 18 months abroad. In Case 6 we have a novel reason: the pro-plaintiff reason associated with a 2% income earned abroad, $M^{pla}_{d2,2\%}$, makes for a set of pro-plaintiff factors has not yet been weighed against the pro-defendant set in this case. This can be used to create the rule $\{M^{pla}_{d1,18}, M^{pla}_{d2,2\%}\} \rightarrow$ Plaintiff, which distinguishes the rule from Case 5 and induces the weighing $\{M^{pla}_{d1,18}, M^{pla}_{d2,2\%}\} > \left\{M^{def}_{d1,18}, M^{def}_{d2,2\%}\right\}$.

However, the approach won't deal with a situation in which Case 7, with 50 months abroad, no income information, and a holding of change of residence, happens before the current case, Case 8, with 50 months abroad and 100% of the

[34] Case 1 using any one of Rule 1, Rule 2, or Rule 3 will force the result in Case 2, a feature discussed *infra* at Sect. 5.

[35] Horty (2021)'s theory can also construct Case 3 in such a way that it won't bind Case 4. See *infra* at Sect. 5.

income earned abroad. The ratio in Case 7 will be the same as in Case 3, i.e. $\{M^{def}_{d1,50}\} \rightarrow$ Defendant with the ordering $\left\{ M^{def}_{d1,50} \right\} > \{ M^{pla}_{d1,50}\}$. It seems like Case 7 should force a holding of a change of residence in Case 8, but it will not. Rather, the 100% of income earned will introduce magnitude factors favoring and opposing a change of residence and the new court will have discretion in how to prioritize them.[36] Case 8 will have four magnitude factors: $M^{def}_{d1,50}, M^{def}_{d2,100\%}, M^{pla}_{d1,50}, M^{pla}_{d2,100\%}$. $\{M^{pla}_{d1,50}, M^{pla}_{d2,100\%}\}$ is a pro-plaintiff set of factors that is not yet weighed against $M^{def}_{d1,50}$, and hence it could be used to distinguished Case 8.

As Cases 7 and 8 show, dimensions remain entangled on Horty's (2021) approach when more than one magnitude factor occurs in the ratio. In those cases, the ratio entails only that those magnitude factors as a set outweigh the set of their opposite magnitude factors (plus any other opposing factors), so any one of those magnitude factors may still be outweighed by its opposite. Technically the ratio still only provides "reasons" favoring the outcome, but the intuition behind that conception of the ratio is largely lost. The ratio may refer to values on a dimension that produce magnitude factors opposing the outcome that are much stronger than the corresponding magnitude factor in the ratio. We gave Case 3 the ratio of $\{M^{def}_{d1,50}\} \rightarrow$ Defendant, but we could have given it this ratio: $\{M^{def}_{d1,50}, M^{def}_{d2,0\%}\} \rightarrow$ Defendant. This ratio involves the magnitude factor in favor of finding a change in domicile that is created by having zero percent of one's income earned abroad, which is plainly weaker than the corresponding magnitude factor in favor of finding no change in domicile for earning no income abroad.

This creates some potential difficulties in the context of distinguishing, which we saw when Case 8 distinguished in favor of the plaintiff on the basis of a fact that seemed clearly to favor the defendant. In pre-dimensional versions of the factor-based model (see Horty 2011, 2015) courts distinguished on the basis of a novel factor favoring the relevant outcome, i.e., the ratio in the distinguishing case had to include the novel opposing factor. This nicely captured the idea that distinguishing involved identifying an exception to the previous rule that distinguished the current case. But for any novel value on a dimension, there will always be two novel magnitude factors, one for each side. In theory then the court could distinguish on the basis of a magnitude factor despite that factor having a stronger corresponding magnitude factor for the opposite conclusion.

To see this, suppose the court has distinguished Case 6 from Case 5 as described above. That is, it distinguishes Case 5, which has the ratio $\{M^{def}_{d1,18}\} \rightarrow$ Defendant, on the basis of the pro-plaintiff magnitude factor of 2% of income earned abroad. It does so by introducing a rule that involves both that magnitude factor and the pro-plaintiff magnitude factor for the 18 months stay abroad: $\{M^{pla}_{d1,18}, M^{pla}_{d2,2\%}\} \rightarrow$ Plaintiff.

[36] This example might be taken to suggest that the extrema of dimensions be treated as factors, although it's not clear that that changing from 100 to 95% of income earned abroad, for instance. As discussed *infra* at ** of the Bench-Capon and Atkinson approach addresses this issue in a slightly different way that enables it to treat certain values on a dimension as producing non-entangled factors.

Table 1 Hypothetical cases

Case Number	Change of domicile	Months abroad	% Income earned abroad
1	Yes	18	60%
2	? (yes, if forced by Case 1)	19	65%
3	Yes	50	0
4	? (yes, if forced by Case 3)	50	?
5	Yes	18	?
6	No	18	2
7	Yes	50	?
8	? (yes, if forced by Case 7)	50	100
9	? (no, if forced by Case 6)	?	2

So we have $\{M^{pla}_{d1,18}, M^{pla}_{d2,2\%}\} > \left\{M^{def}_{d1,18}, M^{def}_{d2,2\%}\right\}$ and $\left\{M^{def}_{d1,18}\right\} > \left\{M^{pla}_{d1,18}\right\}$ as our ordering from past cases. Now we get Case 9, which just has the information that the individual earned 2% of her income abroad. Case 9 will involve only $M^{def}_{d2,2\%}, M^{pla}_{d2,2\%}$. Since simply adding the fact of 2% of income earned abroad in Case 6 tilted the scales in favor of no change of domicile in Case 9, it's tempting to think the court here is compelled to find no change of domicile. That is, it's tempting to assume that $\left\{M^{pla}_{d2,2\%}\right\} > \left\{M^{def}_{d2,2\%}\right\}$, but the ordering does not entail this. Case 9 will not trigger the ratio from Case 5 or Case 6, because Case 9 lacks $M^{def}_{d1,18}$. Therefore, the court will have discretion in deciding Case 9. It could rule that $\left\{M^{def}_{d2,2\%}\right\} \rightarrow$ Defendant, meaning that $\left\{M^{def}_{d2,2\%}\right\} > \left\{M^{pla}_{d2,2\%}\right\}$.

To be clear, these problematic cases are not in conflict with anything in Horty (2021), because there it is assumed that every dimension is assigned a value in the past and current case. Rather, these cases illustrate some of the difficulties in extending that account to cases with incomplete dimensions. Recent work (Odekerken et al. 2023) has attempted to model dimensions to contexts with incomplete dimensions, and I will discuss how this compares with my interpretive proposal in Sect. 5.

Turning to the second prevailing approach, consider the two-stage approach to dimensions and factors from Bench-Capon and Atkinson (2017, 2021, 2022). On this approach dimensions are orderings of facts from which one can derive factors, hence the two-stages: we move from dimensions in the first stage to the factors in the second stage. In addition, these models do not assume that every factor is related to a single dimension. Dropping that assumption means our Case 1 from this section has four possible interpretations, not three:

(A) 60% income earned abroad is a reason for change that outweighs 18 months abroad (a reason against change).
(B) 18 months abroad is a reason for change that outweighs 60% of income earned abroad (a reason against change).

(C) 18 months abroad is a reason for change and so is 60% of income earned abroad.
(D) 60% income and 18 months (the combination of values of each dimension) is a reason for change (a single factor).

The new available interpretation, (D), is the one that the two-stage approach would assign to Case 1. The relevant factor would be something like SufficientAbsenceGivenIncome,[37] which is present when both the value on the dimension for income earned abroad is 60% and when the stay abroad is 18 months in length.

The two-stage models represent entangled factors as present if the values of dimensions lie within an area of two-dimensional space. Using just the information from Case 1, this would be the space that lies above both 60% on the income earned dimension and above 18 months on the length of stay dimension. That is, we would say SufficientAbsenceGivenIncome is present whenever income earned abroad is at least 60% and length of stay is at least 18 months. Hence the model will force a finding of change of domicile in Case 2 based on Case 1. SufficientAbsenceGivenIncome *could* also be present for lower values on either or both dimensions, but the Case 1 doesn't provide that information. Adding more cases provides additional constraints on when that factor is or can be present.

Although the dimensions remain entangled for Case 1, the two-stage approach allows for dimensions to become disentangled for certain values. For example, the dimension of income earned abroad might be associated with four separate factors: the opposing factors of InsufficientIncome (for some lower range of values) and SufficientIncome (for some higher range of values), and the opposing entangled factors SufficientAbsenceGivenIncome (for some middle range of values) and InsufficientAbsenceGivenIncome (for the remaining middle range). Likewise, the dimension of length of stay might be associated with the opposing factors of InsufficientLength (for lower values) and SufficientLength (for higher values), and the opposing entangled factors SufficientAbsenceGivenIncome (for some middle range of values) and InsufficientAbsenceGivenIncome (for the remaining middle range). This allows the model to represent Case 3 as presenting two opposing (not entangled factors) of SufficientLength and InsufficientIncome, with the outcome establishing a priority for SufficientLength over InsufficientIncome.[38] Case 4 can be modelled as a case with the pro-change of domicile factor SufficientLength and no other factor, so the pro-change outcome is compelled there.

Case 5 followed by Case 6 raises questions about how to represent unknown values on a dimension. One could model Case 5 as establishing the presence of the factor SufficientLength for values at and above 18 months (remember, Case 1 isn't in the background for Case 5). In Case 6 this factor could be present but opposed by the entangled factor InsufficientAbsenceGivenIncome and this would give a basis

[37] I'm thankful to an anonymous reviewer for explanation on this point.

[38] You could also represent the priority more strongly, with sufficient length being established as one of (Bruninghaus and Ashley 2003)'s "knock-out factors.".

for distinguishing and finding no change of domicile.[39] Case 6 could also be understood as involving SufficientLength and the opposing but not entangled factor of InsufficientIncome, which would still give the court a basis for distinguishing.

In Case 7 we would have the factor SufficientLength established by the 50 months abroad, and in Case 8 we would have SufficientLength (from the 50 months) and SufficientIncome (from 100% income earned abroad). Since both of those factors point the same way, there is no discretion to decide Case 8 differently than Case 7. Finally, the two-stage model can get Case 6 to force a result of no change in domicile in Case 9, using the interpretation of Case 6 as involving SufficientLength and InsufficientIncome. On that interpretation of Case 6, 2% of income earned abroad establishes InsufficientIncome. Hence in Case 9 we have InsufficientIncome, and since this is unopposed in Case 9 the court is compelled to rule against a change in domicile. This last result reminds us of the importance of the interpretation of the past case. This forms the basis for my proposal in the next section, which uses multiple interpretations to model these phenomena.

5 Incorporating interpretation for dimensions

In Sect. 3 we saw that adding interpretations to the models can provide increased explanatory power and better fit with the language of cases, and that insight can be fruitfully applied here. Start with a reductive model like Rigoni (2018), in which cases are modeled using factors, ratios, and dimensions, which are sets of factors with priorities. Then allow that past cases can have multiple interpretations for the values of dimensions, when there are multiple reasonable readings of opinion on that this point. Consider again our Case 1, which held that there is a change of residence when the defendant spent 18 months living abroad and earned 60% of her income abroad. On this proposal that case has three possible reasonable interpretations:

(A) 60% income is a reason for change that outweighs 18 months (a reason against change).
(B) 18 months is a reason for change that outweighs 60% of income (a reason against change).
(C) 18 months is a reason for change and so is 60% of income.

Instead of forcing a choice between (A), (B), and (C), let the model first check to see if all of the reasonable interpretations compel the same result in the current case. If so, then treat the current case as forced without selecting an interpretation for the past case. This allows the model to capture the binding force of past cases

[39] Within the two-stage approach you could also model Case 5 as containing the entangled factor SufficientAbsenceGivenIncome, if you treat "unknown" as a value on the dimension for income earned. In Case 6 we would then have the factor InsufficientAbsenceGivenIncome unopposed. A number of complications arise from treating unknown values in this way, so I ignore it in the body text. But is an option available on the two stage views.

with entangled dimensions, without further requiring that multiple dimensions in past cases always be treated as entangled.

The general idea is that interpretation does not require resolving every ambiguity of the past opinion; only the ambiguities that could give rise to different results in the current case matter. If every reasonable way of resolving the ambiguity leads to the same result, then the current court will know it has to follow the past case. Returning to Case 1, interpretations (A), (B), and (C) will all compel a finding of a change of residence in Case 2. Keeping Case 1 but now considering Case 5, this approach will not compel a result in Case 5 if (A)-(C) are all equally reasonable interpretations. Case 3 will only admit of the reasonable interpretation of 50 months as a factor that outweighs any percent of income earned, which will then compel the right result in Case 4. Without magnitude factors, the 100% income earned abroad in Case 8 can only function as a pro-change factor, so Case 7 will bind Case 8. Distinguishing still requires identifying a factor favoring the losing side of the past case, so the trouble regarding Case 9 given a background of Cases 5 and 6 is avoided.

Allowing multiple interpretations for cases with respect to the value of dimensions seems to increase explanatory power over a reductivist account like Rigoni (2018)'s and Horty's alternative. However, Sect. 3 also showed that adding interpretations can have significant costs. Yet, introducing interpretations in this limited way, as options for the valence of values on a dimension, is much more manageable than introducing multiple interpretations for cases in general. Here we are only looking at a narrow range of ambiguity—instead of every reasonable reading of the whole opinion, we are only looking at reasonable characterizations of dimensions. Further, we are only looking at a specific type of ambiguity: whether each value is for or against the relevant decision. For each dimension in the case, there are only two ways it could go, with the further restriction that at least one of them has to favor the prevailing side. Hence, where x is the number of relevant dimensions, the maximum number of interpretations is $2^x - 1$. Finally, although the exogenous process has to determine which interpretations are reasonable, it does not, in this form, have to select one interpretation as the best.

The approach advocated here is in-line with the proposal for treating incomplete information on dimensions in (Odekerken et al. 2023). That approach considers instances where the current (the "focus case" in their terminology) has incomplete information on one or more dimensions. Applied to our examples, those would be current cases where the number of months abroad or the percentage of income earned abroad is unknown. The approach holds that the current court is bound by a past case if every way of completing the information in the current case would make the past case bind it. For example, on Oderkerken et al. (2023)'s approach, Case 3 will bind Case 4 because every value for percentage of income earned abroad would make Case 4 at least as good of a case for the defendant as Case 3.[40]

[40] Interestingly, on Horty's (2021) approach, the Case 3 won't bind Case 4 if we give Case 3 this ratio: $\left\{M^{def}_{d1,50}, M^{def}_{d2,0\%}\right\} \rightarrow$ Defendant. However, that is the "ratio" for that case under the results model as used by Oderkerken et al. (2023). Horty's approach can get Case 3 to bind Case 4, as we saw, but it depends on the selection of the ratio. Case 3 must bind Case 4 on Oderkerken et al.'s view.

The approach would need to be modified so that past cases rather than current cases are capable of having incomplete information for a dimension, with the idea being that a past case binds if and only if any way of completing the dimensional information within it would bind. With this modification, the Oderkerken et al. (2023) model will allow Case 5 to not bind Case 6, because there are ways of completing the information in Case 5 that make it stronger for the defendant than Case 6. This would also enable Case 7 to bind Case 8, because every possible value for percentage of income earned abroad in Case 7 would make it at least as strong of a case for the defendant as Case 8.

However, the Oderkerken et al. (2023) model would need to be further modified to bring it in line with my proposal, because my proposal retains the reason model. The Oderkerken et al. (2023) model uses the results model of precedent, on which the weighing established by a case is that the set of all factors (or the set of the strongest dimensional values) favoring the winning side outweigh the set of all factors (or the set of the strongest dimensional values) favoring the losing side. The results model is equivalent to the reasons model if we assume that *Premise* (*r*) is the set of all reasons (or set of strongest dimensional values) favoring the victor. The reasons model allows for *Premise* (*r*) to be a subset of the reasons (or a set of weaker dimensional values) in favor of the victor. Considerations in favor of the reasons model in general are discussed in (Horty and Bench-Capon 2012; Horty 2015) and considerations in favor of it in the context of dimensions are discussed in (Rigoni 2018; Horty 2021).

We saw in Sect. 4 that the two-stage approach to dimensions and factors from Bench-Capon and Atkinson (2017, 2021, 2022) can get the same results as applied to the cases as my proposal. Given what I've argued in Sect. 2, I can hardly claim that reasoning with dimensions is somehow *actually interpretative* and hence my proposal more closely models the *real* reasoning. Rather, I think the virtue of the proposal is that it provides a manageable first step towards dealing with the more complicated interpretive issues discussed in Sect. 3. Case law is replete with interpretative language, and it's worth seeing how closely the formal models can adhere to that. Dimensions provide a nicely constrained area to introduce interpretations and see what the complications arise. The proposal here provides the pieces for modeling more sophisticated phenomena, such as holdings about holdings.[41] For example, current courts may select one of the reasonable interpretations as the interpretation of the past case and make this interpretation itself precedential. A court, reasoning after Case 1 was decided, might choose to assign interpretation (A) to Case 1 in an opinion. This could be an instance of Re's extending, if some of the reasonable interpretations don't bind the current court but the court chooses one that is binding. Arguably, later courts (suitably situated in the hierarchy) are bound to read Case 1 the same way.

[41] See Carey (2013).

6 Concluding remarks: inconsistent case bases

I hope to have shown that allowing for multiple interpretations of cases or re-interpretations of cases within factor-based models is a fruitful direction for further research. I will close with one further brief example of when interpretation can make a difference, which is in dealing with inconsistent case bases. Canavotto (2022) offers an elegant approach to deriving constraint from inconsistent case bases within a factor-models $F_1^p\ F_1^d$. She gives an example of two past cases: $c_5 = (X_5, r_5,$ Plaintiff) where $X_5 = \{F_1^p, F_1^d\}$ and $r_5 = \{F_1^p\} \rightarrow$ plaintiff, and $c_6 = (X_6, r_6,$ Defendant), where $X_6 = \{F_1^p, F_1^d, F_2^d\}$ and $r_6 = \{F_1^d\} \rightarrow$ Defendant. The current case is $c_7 = (X_7, ?, ?)$, where $X_7 = \{F_1^p, F_2^p, F_2^d\}$. Canavotto quite reasonably arrives at result for the plaintiff in c_7, roughly because c_7 has got better facts for the plaintiff than both c_5 and c_6.[42] However, this seems like a set of cases begging for re-interpretation to avoid inconsistency, namely, for the ratio in c_6, r_6, to be revised to $\{F_1^d, F_2^d\} \rightarrow$ defendant. This re-interpretation won't change the result in Canavotto reaches in c_7, but it is closer to what a real judge is likely to write in an opinion if it is at all plausible to read c_6 in this revised way. It will also potentially impact future cases differently, since c_6 now has a different ratio and the case base is now consistent. Considering interpretations can then move us closer to modeling what Dworkin called a sentimental lawyers' cherished trope: the law working itself pure (Dworkin 1986, p. 400).

References

Al-abdulkarim L, Atkinson K, Bench-Capon T (2016a) A methodology for designing systems to reason with legal cases using abstract dialectical frameworks. Artif Intell Law 24(1):1–49. https://doi.org/10.1007/s10506-016-9178-1

Al-abdulkarim L, Atkinson K, Bench-Capon TJM (2016b) Statement types in legal argument. In: Legal Knowledge and Information Systems–JURIX 2016: The Twenty-Ninth Annual Conference. IOS Press, (pp. 3–12)

Al-Abdulkarim L, Atkinson K, Bench-Capon TJM (2015) Factors, Issues and Values: Revisiting Reasoning with Cases. In: Proceedings of the 15th International Conference on Artificial Intelligence and Law. (pp. 3–12)

Ashley KD (1991) Reasoning with cases and hypotheticals in hypo. Int J Man-Mach Stud 34:753–796

Ashley KD, Brüninghaus S (2009) Automatically classifying case texts and predicting outcomes. Artif Intell Law 17(2):125–165. https://doi.org/10.1007/s10506-009-9077-9

Atkinson K, Bench-Capon TJM, Prakken H, Wyner A (2013) Argumentation schemes for reasoning about factors with dimensions. In: Ashley KD (ed) Proceedings of 26th International Conference on Legal Knowledge and Information Systems (JURIX 2013). IOS Press, Amsterdam, (pp. 39–48)

Bench-Capon T (1999) Some observations on modelling case based reasoning with formal argument models. In: Proceedings of the 7h International Conference on Artificial intelligence and Law. pp 36–42

Bench-Capon T (2023) Whatever happened to hypotheticals? In: Proceedings of the 19th International Conference on Artificial Intelligence and Law. Association for Computing Machinery, (pp. 387–391)

[42] See (Canavotto 2022, pp. 29–31). Her notion of support entails that c5 supports a pro-plaintiff ruling in c7 but c6 won't support a pro-defendant ruling in c7.

Bench-Capon T, Atkinson K (2022) Argument schemes for factor ascription. In: Toni F, Polberg S, Booth R, et al. (eds) Proceedings of COMMA. IOS Press, (pp. 68–79)
Bench-Capon T, Atkinson K (2021) Precedential constraint: the role of issues. Association for Computing Machinery
Bench-Capon T, Atkinson K (2017) Dimensions and values for legal CBR. In: Wyner A, Casini G (eds) Legal knowledge and information systems. IOS Press, pp 27–32
Branting LK (2020) Explanation in hybrid, two-stage models of legal prediction. In: XAILA Workshop at JURIX. (pp. 1–10)
Branting LK, Pfeifer C, Brown B et al (2021) Scalable and explainable legal prediction. Artif Intell Law 29(2):213–238. https://doi.org/10.1007/s10506-020-09273-1
Broughton GL (2019) Vertical precedents in formal models of precedential constraint. Artif Intell Law 27(3):253–307. https://doi.org/10.1007/s10506-019-09244-1
Bruninghaus S, Ashley K (2003) Predicting outcomes of case based legal arguments. In: Proceedings of 9th International Confernece on Artificial Intelligence Law, (pp. 233–42)
Canavotto I (2022) Precedential constraint derived from inconsistent case bases. Front Artif Intell Appl 362:23–32. https://doi.org/10.3233/FAIA220445
Carey M (2013) Holdings about holdings: Modeling contradictions in judicial precedent. Artif Intell Law 21(3):341–365. https://doi.org/10.1007/s10506-013-9141-3
Collenette J, Atkinson K, Bench-Capon T (2023) Explainable AI tools for legal reasoning about cases: a study on the European court of human rights. Artif Intell 317:103861. https://doi.org/10.1016/j.artint.2023.103861
Cross FB, Tiller EH (1998) Judicial partisanship and obedience to legal doctrine: Whistleblowing on the Federal Courts of Appeals, Yale Law J. 107
Dworkin R (1986) Law's empire. Harvard University Press, Cambridge
Gray M, Savelka J, Oliver W, Ashley KD (2022) Toward automatically identifying legally relevant factors. In: Francesconi E et al (eds) Legal knowledge and information systems. IOS Press, pp 53–62
Horty J (2021) Modifying the reason model. Artif Intell Law 29(2):271–285. https://doi.org/10.1007/s10506-020-09275-z
Horty JF (2017) Reasoning with dimensions and magnitudes. In: Proceedings of the 15th International Conference on Artificial Intelligence and Law. ACM Press, (pp. 109–118)
Horty JF (2015) Constraint and freedom in the common law. Philos Impr 15(25):1–27
Horty JF (2019) Reasoning with dimensions and magnitudes. Artif Intell Law 27(3):309–345
Horty JF (2011) Rules and reasons in the theory of precedent. Leg Theory 17(01):1–33
Horty JF, Bench-Capon TJM (2012) A factor-based definition of precedential constraint. Artif Intell Law 20(2):181–214
Mullins R (2022) Two factor-based models of precedential constraint: a comparison and proposal. Artif Intell Law. https://doi.org/10.1007/s10506-022-09335-6
Mumford J, Atkinson K, Bench-Capon T (2022) Reasoning with legal cases: a hybrid ADF-ML approach. Legal knowledge and information systems. IOS Press, pp 93–102
Odekerken D, Bex F, Prakken H (2023) Justification, stability and relevance for case-based reasoning with incomplete focus cases. In: Proceedings of the 19th International Conference on Artificial Intelligence and Law. (pp. 177–186)
Pfleiderer P (2020) Chameleons: the misuse of theoretical models in finance and economics. Economica 87(345):81–107. https://doi.org/10.1111/ECCA.12295
Posner RA (2008) How judges think. Harvard University Press, Cambridge
Prakken H (2021) A formal analysis of some factor- and precedent-based accounts of precedential constraint. Artif Intell Law 29(4):559–585
Raz J (2009) The authority of law : essays on law and morality. Oxford University Press, Oxford
Re R (2016) Narrowing supreme court precedent from below. Georgetown Law J 104(1):921–971
Re R (2014) Narrowing precedent in the supreme court. Columbia Law Rev 114:1861–1911
Rigoni A (2015) An improved factor based approach to precedential constraint. Artif Intell Law 23(2):133–160. https://doi.org/10.1007/s10506-015-9166-x
Rigoni A (2021) AA-RM wrestling: comparing analogical approaches and rule models for legal reasoning. Leg Theory 27(3):207–235. https://doi.org/10.1017/S135232522100015X
Rigoni A (2018) Representing dimensions within the reason model of precedent. Artif Intell Law 26(1):1–22. https://doi.org/10.1007/s10506-017-9216-7
Rissland E, Ashley K (1987) HYPO: a precedent-based legal reasoner. Amherst, MA

Roth B, Verheij B (2004) Dialectical arguments and case comparison. Legal knowledge and information systems. Jurix 99–108
Simpson AWB (1961) The ratio decidendi of a case and the doctrine of binding precedent. In: Guest AG (ed) Oxford essys in jurisprudence: a collaborative work. Oxford University Press, London, pp 148–175
Stevens K (2018) Reasoning by precedent– between rules and analogies. Leg Theory 24(3):216–254. https://doi.org/10.1017/S1352325218000113
Wendell Holmes O (1897) The path of the law. Harv Law Rev 10:457

Cases Cited

Arizona State Legislature v. Arizona Independent Redistricting Commission et al
Arizona Independent Redistricting Commission, 576 U.S. 787 (2015)
Janus v. American Federation of State, County, and Municipal Employees, Council 31, 585 U.S. __ (2018) [abbreviated *Janus* v. *AFSCME*]
Riley v. *California*, 573 U.S. 373 (2014)
Stanford v. Kentucky, 492 U.S. 361 (1989)
State Ex rel Simmons v. Roper, 112 S.W. 3d 397 (Mo. 2003)
State Ex rel Simmons v. Roper, 543 U.S. 551 (2005)
United States v. Robinson, 414 U.S. 218 (1973)
United States v. Wurie, 728 F.3d 1 (1st Cir. 2013)