Reasoning with hierarchies of open-textured predicates

Ilaria Canavotto
icanavot@umd.edu
Philosophy Department
University of Maryland
College Park, MD, USA

John Horty
horty@umd.edu
Philosophy Department and
Institute for Advanced Computer Studies
University of Maryland
College Park, MD, USA

ABSTRACT

We develop the reason model of precedential constraint in the context of a hierarchy of intermediate legal concepts, based on the idea that constraint depends, not just on the ultimate decision reached in a case, but on the variety of intermediate decisions that may constitute judicial opinions. After developing this model, we study the relation between constraint in a full hierarchical setting and constraint in a setting in which hierarchies, cases, and case bases are flattened into structures corresponding to those at work in the standard reason model. We show that constraint in the full hierarchical setting may be lost in the flattened setting, but also, and more surprisingly, that new patterns of constraint might appear in the flattened setting that were not present in the full hierarchical setting.

CCS CONCEPTS

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

KEYWORDS

Reason model, issues, abstract factors, hierarchies.

ACM Reference Format:
Ilaria Canavotto and John Horty. 2023. Reasoning with hierarchies of open-textured predicates. In *Nineteenth International Conference on Artificial Intelligence and Law (ICAIL 2023), June 19–23, 2023, Braga, Portugal*. ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3594536.3595148

1 INTRODUCTION

The earliest computational models of case-based legal argument in the Rissland/Ashley tradition were essentially flat, representing legal information in a space of unrelated factors and dimensions [3, 17]. The same holds true of many of the more recent models of precedential constraint, which were constructed on the basis of this initial Rissland/Ashley representation; an early paper on the reason model of constraint, for example, explicitly noted that the account "simplifies by assuming that precedential reasoning involves only a single step, proceeding from the factors present in a case directly to a decision in favor of the plaintiff or defendant, instead of moving through a series of intermediate legal concepts" [12, p. 28].

In the computational study of legal argument, the restriction to a flat representational framework did not last long. Aleven soon showed with CATO how a more subtle understanding of argument moves could be developed on the basis of a hierarchical organization of legal concepts [1, 2]. A similar hierarchical organization figured in Brüninghaus and Ashley's work on legal prediction [8], and in Prakken and Sartor's work on modeling case-based legal argument within a logical framework [16].

Turning from argument analysis and legal prediction to constraint, Bench-Capon and Atkinson [4] have recently recommended an "issue-based" version of the reason model of precedential constraint—where "issues" are high-level concepts in the hierarchy of legal information. More exactly, Bench-Capon and Atkinson think of issues as concepts that are so high in the hierarchy that any further relation between these concepts and an overall legal judgment is purely deductive. Taking the trade-secrets domain as an example, they illustrate the deductive relation between issues and overall judgments by noting that, in this domain, an overall judgment for the plaintiff is determined by the separate conclusions "*both* that [a particular body of] information was a trade secret *and* that it was misappropriated"—so that these two concepts, the concepts of a trade secret and of misappropriation, are the relevant high-level issues [4, p. 17]. Their suggestion, then, is that, rather than evaluating the entire set of factors en masse as supporting an overall judgment for the plaintiff or the defendant, as in the ordinary reason model, it is better to evaluate the factors that bear on these particular legal issues separately, resolve those issues on their own, and then rely on ordinary logic to determine an overall outcome on the basis of their resolution.

We agree with this suggestion, but think there are situations in which it does not go far enough. To understand why, it is useful to consider a fragment of the CATO hierarchy of concepts bearing on the question whether some body of information should be classified as a trade secret, depicted here in Figure 1.[1] As we can see, the factor f_{101}, representing the concept of a trade secret itself, is positioned at the very top of this hierarchy, as an issue—any further relations between f_{101} and an eventual judgment for the plaintiff or defendant are purely deductive. There are, in addition, a number of base-level factors, such as f_1 and f_6, for example, indicating that the information in question was disclosed in negotiations, or that there was an agreement not to disclose this information. But this is not all. Lying between the base-level factors and the top-level

[1] Bench-Capon and Atkinson offer a boolean definition of the concept of a trade secret as information that is both valuable and subject to security measures, where these latter two concepts are the relevant legal issues. This boolean analysis of a trade secret is found earlier in Brüninghaus and Ashley [8], but differs from the treatment originally proposed by Aleven [1, 2], who treats information value and security measures, not as concepts in terms of which a trade secret can be defined in a boolean way, but simply as factors bearing on the question whether information is a trade secret.

ICAIL 2023, June 19–23, 2023, Braga, Portugal

ACM ISBN 979-8-4007-0197-9/23/06.
https://doi.org/10.1145/3594536.3595148

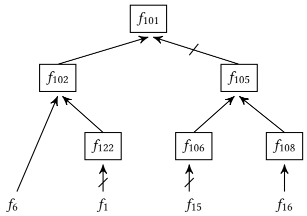

Figure 1: A fragment of the CATO hierarchy

issue, there are also a number of intermediate concepts, or factors, such as f_{102} and f_{105}, indicating that efforts were taken to maintain security, or that the information in question is known or available.

These intermediate factors are not so high in the hierarchy that their relation to an overall judgment is simply deductive, nor so low in the hierarchy that questions concerning their application can be resolved entirely without dispute. Like top-level issues themselves, intermediate factors can be thought of as open-textured predicates, whose applicability in some particular situation must be determined, in a step-by-step process, before considering the application of other concepts further up the hierarchy. And just as Bench-Capon and Atkinson suggest that factors should be considered, not en masse, but only as they bear on particular legal issues, we offer two similar suggestions: First, at any stage of the step-by-step process, only factors bearing on the concepts whose application is in question at that stage are to be considered. And second, decisions about the application of intermediate concepts in particular situations can be taken to constrain decisions about the application of these same concepts in later situations.

These suggestions are not new. Branting [5, 6] has argued both that understanding the intermediate decisions leading up to an overall legal judgment helps to clarify the theory under which a case is decided, and also that these intermediate decisions themselves can constrain later judgments—he refers to intermediate decisions as "precedent constituents."[2] And as mentioned, Prakken and Sartor [16] also work with a multi-step representation of legal arguments, and also suggest, following Branting, that the intermediate steps in such an argument can constrain future decisions. Still, although both Branting and Prakken and Sartor suggest that constraint can be derived from precedent constituents, as well as from entire precedent cases, neither explores this suggestion against the background of a precise model of precedential constraint.

This is what we do. More exactly, our plan in this paper is to develop the reason model of constraint in the context of a hierarchy of intermediate legal concepts. The key idea is that, just as decisions favoring the plaintiff or defendant in the standard reason model lead to a priority ordering over reasons favoring each side, decisions concerning the application of intermediate legal concepts in a hierarchical setting lead to a priority ordering over reasons favoring or opposing application of those concepts, which must likewise be respected by later courts. These intermediate decisions are formed into structures that we characterize as *opinions*, which play the role of rules from the standard reason model in justifying overall outcomes.

[2] The phrase is from Branting [5], which also provides examples in which precedent constituents from one case are appealed to in later cases concerning different issues altogether.

The paper is organized as follows. In Section 2, we introduce the notion of a factor hierarchy that forms the context of our work. In Section 3, we define the concept of an opinion as the generalization of a rule, and define cases as containing opinions rather than rules. Taking advantage of the various precedent constituents belonging to an opinion, Section 4 then adapts the reason model to the hierarchical setting. Finally, Section 5 explores the relation between constraint in the full hierarchical setting and constraint in a setting in which full hierarchies are flattened into structures corresponding to those at work in the standard reason model—we find that, constraint in the full hierarchical setting is often stronger than constraint in the corresponding standard setting, but that, surprisingly, there are also situations in which the opposite holds.

2 FACTORS AND HIERARCHIES

We postulate a set of *factors*: predicates, or concepts—often open-textured—used to characterize a situation. Some of these factors have *contraries*, in the traditional sense that two contrary factors cannot both apply in a particular situation, but at any given point, it may not yet be determined which applies. For illustration, consider the concept of being a trade secret and its contrary, the concept of not being a trade secret. A body of information cannot be classified as both a trade secret and not a trade secret, of course, but it may not yet be determined whether that information is to be classified as a trade secret or not. Where $p, q, r, \ldots$ are factors, we take $p', q', r', \ldots$ as their contraries. We use s, t, and u as variables ranging over factors, with $\bar{s}$, $\bar{t}$, and $\bar{u}$ as contraries. Finally, we assume that the contrary of the contrary of a factor is the factor itself—so that, if s is p, for example, then $\bar{s}$ is p', and if s is p', then $\bar{s}$ is p.

A *factor link* is a statement of the form $s \rightarrow t$ indicating that the presence of the factor s in some situation *directly favors* a decision that t holds as well, or simply that s directly favors t. A *factor hierarchy* is a set $\mathcal{H}$ of factor links.

Within a particular hierarchy, a factor can be either base-level or abstract. The intuition is that base-level factors are concepts about whose application there is no disagreement, while abstract factors are concepts about whose application there may be some dispute. The formal definition is that the factor s is *abstract* if the hierarchy contains a link either of the form $t \rightarrow s$ or of the form $t \rightarrow \bar{s}$, and *base-level*, or at the bottom of the hierarchy, if the hierarchy contains no such link; we let

$$Base_{\mathcal{H}}$$

refer to the set of base-level factors from the hierarchy $\mathcal{H}$. An abstract factor can itself be either intermediate or top-level, where s is *intermediate* if the hierarchy contains some link either of the form $s \rightarrow t$ or of the form $\bar{s} \rightarrow t$, and *top-level*, or at the top of the hierarchy, if the hierarchy contains no such link.

We will assume that base-level factors have no contraries, so that any set of base-level factors is consistent, and that all abstract factors have contraries.[3] The factors that *belong to a hierarchy* are simply those at the head or tail of some link from that hierarchy, along with the contraries of all abstract factors that belong to the hierarchy.

Where s is an abstract factor, we define $s/\bar{s}$ as a *concern*—from an intuitive standpoint, the particular concern indicated by $s/\bar{s}$ is

[3] See Bench-Capon and Atkinson [4, p. 13–14] for a discussion of the issues surrounding factors and their "negations," or contraries.

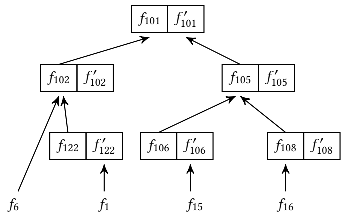

Figure 2: Interpreting the CATO fragment

whether the factor s or its contrary $\bar{s}$ should be applied in some situation. Concerns can be classified as *intermediate* or *top-level*, depending on the classification of the factors involved; a top-level concern is an *issue*. Where s is an abstract factor from the hierarchy $\mathcal{H}$, again

$$F^s_{\mathcal{H}} = \{t : t \rightarrow s \in \mathcal{H}\}$$

is the set of factors from $\mathcal{H}$ directly favoring s, so that $F^{s/\bar{s}}_{\mathcal{H}} = F^s_{\mathcal{H}} \cup F^{\bar{s}}_{\mathcal{H}}$ is the set of factors *bearing on the concern* $s/\bar{s}$—that is, directly favoring either s or $\bar{s}$.

These various ideas are illustrated in Figure 2, which is a transcription into the current setting of the CATO hierarchy from Figure 1. The chief difference between the two figures is that, while the nodes in the original CATO hierarchy represent factors, all but the base-level nodes in the present hierarchy represent concerns—pairs of factors and their contraries. Because factor contraries are now represented explicitly, we can modify the link structure from the original CATO hierarchy, indicating that the presence of one factor favors the contrary of a second, not by a negative link from the first factor to the second, but by an ordinary positive link from the first factor to the contrary of the second. To illustrate: the negative CATO link $f_{105} \nrightarrow f_{101}$, according to which the fact that information was known or available favors the conclusion that it is not a trade secret, can now be replaced by the positive link $f_{105} \rightarrow f'_{101}$, carrying exactly the same meaning.

The hierarchy depicted in Figure 2 contains four base-level factors (f_6, f_1, f_{15}, and f_{16}), five intermediate concerns (f_{102}/f'_{102}, f_{105}/f'_{105}, f_{122}/f'_{122}, f_{106}/f'_{106}, and f_{108}/f'_{108}), and a single top-level issue (f_{101}/f'_{101}). If we let $\mathcal{H}$ represent this hierarchy, the set of factors directly favoring application of f_{101} is $F^{f_{101}}_{\mathcal{H}} = \{f_{102}\}$, while the set of factors bearing on the question whether f_{101} or its contrary $f_{101'}$ should be applied is $F^{f_{101}/f'_{101}}_{\mathcal{H}} = \{f_{102}, f_{105}\}$.

We impose two conditions on factor hierarchies, both implicit in CATO, but which it is worth making explicit. The first is that a factor hierarchy must be acyclic in the following sense:

Definition 1 (Acyclicity). A factor hierarchy is acyclic just in case it does not contain any sequence of links of the form $s_1 \rightarrow s_2$, $s^*_2 \rightarrow s_3, \ldots, s^*_n \rightarrow s_{n+1}$ where, for i from 2 to n, each s^*_i is either s_i or $\overline{s_i}$, and s_{n+1} is either s_1 or $\overline{s_1}$.

Put another way, what acyclicity requires is that the hierarchy cannot present a sequence of concerns $s_1/\bar{s}_1, s_2/\bar{s}_2, \ldots, s_n/\bar{s}_n$ such that some factor from each concern bears on the next, and some factor from the last bears on the first. The reason for this condition is that it guarantees that concerns can be addressed in an appropriate order, so that all concerns bearing on a particular concern can be resolved before we address that concern.[4]

The second condition is that no factor from a hierarchy should favor opposite sides of the same concern. In order to define this condition, we first need to explain what it means for one factor to *favor* another. Of course, we know that the factors directly favoring a particular outcome s in the context of a hierarchy $\mathcal{H}$ are those contained in $F^s_{\mathcal{H}}$, but it seems that other factors might favor s in a more general sense even if they do not belong to $F^s_{\mathcal{H}}$—in the hierarchy from Figure 2, for example, it is natural to suppose that f_{122} favors f_{101} in this more general sense, since f_{102} directly favors f_{101} and f_{122}, directly favors f_{102}. This more general sense of favoring can be defined as follows:

Definition 2 (Favoring). In the context of a hierarchy $\mathcal{H}$, the set of factors favoring an outcome s—written, $Favor^s_{\mathcal{H}}$—is defined in accord with the rules

(1) If $t \in F^s_{\mathcal{H}}$, then $t \in Favor^s_{\mathcal{H}}$.
(2) If $t \in Favor^s_{\mathcal{H}}$ and $u \in F^t_{\mathcal{H}}$, then $u \in Favor^s_{\mathcal{H}}$.
(3) If $t \notin Base_{\mathcal{H}}$ and $t \in Favor^s_{\mathcal{H}}$, then $\bar{t} \in Favor^{\bar{s}}_{\mathcal{H}}$.

Again letting $\mathcal{H}$ be the hierarchy from Figure 2 for illustration, we can see that the factors favoring f_{101} are those belonging to $Favor^{f_{101}}_{\mathcal{H}} = \{f_{102}, f'_{105}, f_{122}, f'_{106}, f'_{108}, f_6, f_{15}\}$ and that the factors favoring f'_{101} are those belonging to $Favor^{f'_{101}}_{\mathcal{H}} = \{f'_{102}, f_{105}, f'_{122}, f_{106}, f_{108}, f_1, f_{16}\}$.[5]

With this notion in hand, we can now state the condition that no factor from a hierarchy should favor opposite sides of the same concern as the requirement that the hierarchy must be factor uniform, where this idea is defined as follows:

Definition 3 (Factor uniformity). A hierarchy $\mathcal{H}$ is factor uniform if and only if, for any abstract factor s belonging to that hierarchy, $Favor^s_{\mathcal{H}} \cap Favor^{\bar{s}}_{\mathcal{H}} = \emptyset$.

This is, in fact, a very strong requirement, which can be violated in many natural situations, where a single consideration can sometimes seem, in different ways, to support each of two sides of some concern.[6] But uniformity is satisfied in the CATO representation, as well as, we suppose, in many legal settings. Again taking $\mathcal{H}$ as the Figure 2 hierarchy, it is easy to see that $Favor^{f_{101}}_{\mathcal{H}} \cap Favor^{f'_{101}}_{\mathcal{H}} = \emptyset$, and that the same property holds for all other abstract factors from that hierarchy. The reason we impose this uniformity condition here is that it facilitates our Section 5 comparison between the reason model of constraint developed in the full hierarchical setting and the same notion of constraint developed in the standard setting of the familiar reason model.

We limit our attention in this paper, then, to factor hierarchies that are both acyclic and factor uniform—these conditions will be

[4] It is tempting, in some ways, to characterize a hierarchy as acyclic just in case it does not contain a cyclic *path*, defined as a link sequence of the form $s_1 \rightarrow s_2 \rightarrow \cdots \rightarrow s_1$. The condition set out in Definition 1 corresponds to the stronger requirement that the hierarchy does not contain a cyclic *generalized path*, as this concept is defined in Horty, Thomason, and Touretzky [14]; this stronger requirement allows us, in the next section, to adapt the notion of degree from [14] to the current framework.
[5] A similar idea can be found in [18]
[6] An example of a natural but non-uniform factor hierarchy is provided in a longer version of this paper; it is non-uniformity of this kind that gives rise to the philosophical debates surrounding *particularism*—see, for example, Dancy [10].

taken as built into the concept of a factor hierarchy. In addition, for simplicity, we will focus here on hierarchies that are also *single-issue*, in the sense that they contain only a single top-level issue. Of course, a single-issue hierarchy can be organized around any particular issue—Figure 2, for example, depicts a hierarchy organized around the issue f_{101}/f'_{101}, whether or not some body of information should be characterized as a trade secret. But, for further simplicity, we will assume that the single issue at stake in the single-issue hierarchies we consider is always π/δ—whether, that is, the factor π should apply, so that the situation under consideration is decided for the plaintiff, or whether δ should apply, so that the situation is decided for the defendant.

For purposes of comparison, we will at times consider factor hierarchies that are not just acyclic, factor uniform, and single-issue, but also *flat* in the sense that they contain only base-level and top-level factors, rather than factors that are truly intermediate. A hierarchy of this kind can be characterized as *standard*, as follows:

Definition 4 (Standard hierarchy). A standard hierarchy is a factor hierarchy containing only the two abstract factors π and δ, as contraries, and in which each link has the form $s \rightarrow \pi$ or $s \rightarrow \delta$.

Apart from π and δ, all other factors from a standard hierarchy $\mathcal{H}$ are base-level, directly favoring one of these two sides, and so belonging to the set $F^{\pi/\delta}_{\mathcal{H}} = F^{\pi}_{\mathcal{H}} \cup F^{\delta}_{\mathcal{H}}$. And since standard hierarchies are flat, it follows from our definition of favoring that $Favor^{\pi}_{\mathcal{H}} = F^{\pi}_{\mathcal{H}}$ and $Favor^{\delta}_{\mathcal{H}} = F^{\delta}_{\mathcal{H}}$, and so from factor uniformity that $F^{\pi}_{\mathcal{H}} \cap F^{\delta}_{\mathcal{H}} = \emptyset$—no factor favors both sides. Standard hierarchies, therefore, correspond to the structures underlying the standard version of the reason model, with the set of factors partitioned into those directly favoring π and those directly favoring δ. The more general factor hierarchies considered here have these standard hierarchies as a special case.

3 FROM RULES TO OPINIONS

We begin by adapting some ideas from the standard version of the reason model to the hierarchical framework. Given a factor hierarchy $\mathcal{H}$, a *fact situation* in the context of that hierarchy will be defined as some subset X of the factors from that hierarchy, and a *base-level* fact situation as a subset $X \subseteq Base_{\mathcal{H}}$ of the base-level factors from that hierarchy. Where s is an abstract factor, a *reason for* s is a subset $U \subseteq F^{s}_{\mathcal{H}}$ of the factors directly favoring s, and a *reason* is a reason for one side or the other; the reason U *holds* in a fact situation X just in case $U \subseteq X$. Where U is a reason for s, a *rule* is a statement of the form $r = U \longrightarrow s$, with $Premise(r) = U$ picking out the premise of this rule and $Conclusion(r) = s$ its conclusion; the rule r is *applicable* in a situation whenever $Premise(r)$ holds in that situation. Finally, a *decision* is defined as a structure of the form $d = \langle X, r, s \rangle$, where X is some fact situation, r is applicable in X, and $Conclusion(r) = s$. Given such a decision $d = \langle X, r, s \rangle$, we define $Facts(d) = X$, $Rule(d) = r$, and $Outcome(d) = s$.

In the standard version of the reason model, with π and δ as the only abstract factors and the base-level factors partitioned into those favoring the plaintiff and those favoring the defendant, a decision of the form specified here is referred to as a case; such a decision justifies either π or δ as an outcome directly in terms of base-level factors—the reasoning moves from base-level factors to a resolution of the top-level issue π/δ all at once, in a single step.

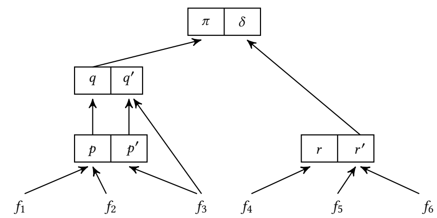

Figure 3: An abstract example

In the hierarchical setting, by contrast, the reasoning that leads from base-level factors to the resolution of a top-level issue may involve a number of steps, requiring the resolution of a variety of intermediate concerns, in an appropriate order, and linked together in an appropriate way. The resulting justification—an opinion—is more complex than a simple rule. We define this notion in three steps: specifying, first, the order in which abstract concerns are to be resolved, then what it means to resolve such a concern, and finally, a non-deterministic algorithm that constructs an opinion by linking these resolutions together.

Beginning, then, with order of resolution, the idea is that concerns must be resolved in accord with the *degree* of the factors they contain, where this concept is defined by stipulating that the degree of the factor s in the context of a hierarchy $\mathcal{H}$—written $Degree_{\mathcal{H}}(s)$—is 0 if s is a base-level factor, and that otherwise, if s is abstract, then

$$Degree_{\mathcal{H}}(s) = 1 + max\{Degree_{\mathcal{H}}(t) : t \in F^{s/\bar{s}}_{\mathcal{H}}\}.$$

It is easy to verify that this notion of degree is well-defined as long as the underlying hierarchy is acyclic, in the sense defined earlier. And once the notion of degree of a factor is defined, it can be lifted to a corresponding notion of degree of a concern, by stipulating that

$$Degree_{\mathcal{H}}(s/\bar{s}) = Degree_{\mathcal{H}}(s),$$

and then used to define the entire set of concerns of some degree n as

$$Degree^{n}_{\mathcal{H}} = \{s/\bar{s} : Degree_{\mathcal{H}}(s/\bar{s}) = n\}.$$

Of course, even at the stage at which we are addressing concerns of degree n, there is no need to address all such concerns. In general, when reasoning about a fact situation X, it is necessary to address only those concerns from the hierarchy that are actually *raised by* this fact situation; these can be collected together in the set

$$Concern_{\mathcal{H}}(X) = \{s/\bar{s} : X \cap F^{s/\bar{s}}_{\mathcal{H}} \neq \emptyset\}.$$

We can then define the *concerns of degree n that are raised by the fact situation X* in the natural way, by intersecting these two sets:

$$Concern^{n}_{\mathcal{H}}(X) = Concern_{\mathcal{H}}(X) \cap Degree^{n}_{\mathcal{H}}.$$

For illustration, we now let $\mathcal{H}$ represent the abstract hierarchy depicted in Figure 3; this hierarchy contains the six base-level factors f_1 through f_6, the three intermediate concerns p/p', q/q', and r/r', and the top-level issue π/δ. Suppose the fact situation under consideration is $X_1 = \{f_1, f_3, f_4, f_5\}$. Then the concerns raised by this fact situation are $Concern_{\mathcal{H}}(X_1) = \{p/p', q/q', r/r'\}$ while the concerns of degree 1 in the hierarchy are $Degree^{1}_{\mathcal{H}} = \{p/p', r/r'\}$, so that the concerns of degree 1 raised by this fact situation are $Concern^{1}_{\mathcal{H}}(X_1) = \{p/p', r/r'\}$.

Next, where $s/\bar{s}$ is a concern, we define a *resolution of this concern based on the fact situation X* as a decision either of the form $d = \langle X, r, s\rangle$ or of the form $d' = \langle X, r', \bar{s}\rangle$. A resolution of a concern based on a fact situation, then, is simply a decision for one side or the other of the concern, together with a rule justifying the decision that is applicable in that fact situation. Given this concept of a resolution of an individual concern, we can now define a *complete resolution of some set of concerns based on a fact situation* as a set containing, for each concern from the set, exactly one resolution of that concern based on that fact situation.

To illustrate: We have seen that the set $Concern^1_{\mathcal{H}}(X_1) = \{p/p', r/r'\}$ contains the concerns of degree 1 raised by the fact situation X_1. These two concerns might be resolved on the basis of this fact situation by, for example, the decisions $d_1 = \langle X_1, r_1, p\rangle$ and $d_2 = \langle X_1, r_2, r'\rangle$, where $r_1 = \{f_1\} \rightarrow p$ and $r_2 = \{f_5\} \rightarrow r'$. The first of these represents a decision for p on the basis of the reason $\{f_1\}$ in spite of the fact that the situation X_1 also contains the factor f_3 favoring p', and the second represents a decision for r' on the basis of the reason $\{f_5\}$ in spite of the fact that X_1 also contains the factor f_4 favoring r. Since d_1 resolves the concern p/p' and d_2 resolves the concern r/r', the set $\{d_1, d_2\}$ containing both of these decisions is a complete resolution of the particular set of concerns from $Concern^1_{\mathcal{H}}(X_1)$ on the basis of X_1.

Finally, we define a procedure through which the resolutions of the various concerns encountered by a reasoner as it works its way through a factor hierarchy can be assembled together into an opinion:

Definition 5 (Opinion based on a fact situation). Given a single-issue factor hierarchy $\mathcal{H}$ and a base-level fact situation $X \subseteq Base_{\mathcal{H}}$, an opinion based on X in the context of $\mathcal{H}$ is defined as an output of the following procedure:

(1) Input the fact situation X
(2) Set
 (a) $Y_0 = X$
 (b) $m = Degree_{\mathcal{H}}(\pi/\delta)$
(3) For $n = 0$ to $m - 1$,
 (a) Let $Resolve_{n+1}$ be some complete resolution of $Concern^{n+1}_{\mathcal{H}}(Y_n)$ based on the fact situation Y_n
 (b) Set $Y_{n+1} = Y_n \cup \{Outcome(d) : d \in Resolve_{n+1}\}$
(4) Output $o = \langle Resolve_1, Resolve_2, \ldots, Resolve_m\rangle$ as an opinion based on X

This procedure works as follows: First, it inputs fact situation X in Step 1. Next, it initializes the variable Y_0 to X in Step (2a), and then in Step (2b) sets the parameter m to the degree of the top-level issue from the hierarchy—given our simplifying assumption that the single issue in a single-issue hierarchy is always π/δ, this will be the degree of that issue. The procedure then enters its central subprocedure in Step 3, which it iterates m times, for n from 0 to $m - 1$. Each iteration has two steps. First, in Step (3a), the procedure calculates the set $Concern^{n+1}_{\mathcal{H}}(Y_n)$ of concerns of degree $n + 1$ raised by the fact situation Y_n, and selects some complete resolution $Resolve_{n+1}$ of this set of concerns based on this fact situation—this is the nondeterministic part of the procedure, since $Concern^{n+1}_{\mathcal{H}}(Y_n)$ may have multiple complete resolutions. Then, in Step (3b), it augments the situation Y_n with the outcomes of the various resolutions belonging to $Resolve_{n+1}$ to form a richer characterization Y_{n+1} of the current fact situation—this richer characterization will then, of course, raise new concerns, which are resolved during later iterations of the subprocedure. Once this series of m iterations of the subprocedure is complete, the main procedure then outputs in Step 4 an m-tuple consisting of the complete resolutions of the concerns raised at each step as an opinion based on the initial fact situation X.

For a concrete illustration, we trace one possible run of the procedure, in the context of the hierarchy $\mathcal{H}$ from Figure 3 and taking the previous situation $X_1 = \{f_1, f_3, f_4, f_5\}$ as input. The procedure begins by setting Y_0 to X_1, the initial fact situation, and setting m to 3, the degree of the issue π/δ in the hierarchy. It then iterates its central subprocedure 3 times, for n from 0 to 2, as follows:

- During the first iteration, the procedure calculates $Concern^1_{\mathcal{H}}(Y_0) = \{p/p', r/r'\}$ as the set of concerns of degree 1 raised by Y_0, and selects a complete resolution of this set based on Y_0. Let us suppose that the complete resolution it selects is $Resolve_1 = \{d_1, d_2\}$, with $d_1 = \langle Y_0, r_1, p\rangle$ where $r_1 = \{f_1\} \rightarrow p$ and with $d_2 = \langle Y_0, r_2, r'\rangle$ where $r_2 = \{f_5\} \rightarrow r'$. It then sets $Y_1 = Y_0 \cup \{p, r'\}$, supplementing the initial characterization of the fact situation with the outcomes of its new decisions.
- During the second iteration, it calculates $Concern^2_{\mathcal{H}}(Y_1) = \{q/q'\}$ as the set of concerns of degree 2 raised by Y_1, and selects a complete resolution of this set based on Y_1. Let us suppose that the complete resolution it selects is $Resolve_2 = \{d_3\}$, with $d_3 = \langle Y_1, r_3, q\rangle$ where $r_3 = \{p\} \rightarrow q$. It then sets $Y_2 = Y_1 \cup \{q\}$, supplementing the current characterization of the fact situation with the outcome of its new decision.
- During the third iteration, it calculates $Concern^3_{\mathcal{H}}(Y_2) = \{\pi/\delta\}$ as the set of concerns of degree 3 raised by Y_2, and selects a complete resolution of this set based on Y_2. Let us suppose that the complete resolution it selects is $Resolve_3 = \{d_4\}$, with $d_4 = \langle Y_2, r_4, \pi\rangle$ where $r_4 = \{q\} \rightarrow \pi$. It then sets $Y_3 = Y_2 \cup \{\pi\}$, supplementing the current characterization of the fact situation with the outcome of its new decision.

After three passes through this iterative subprocedure, the main procedure terminates and outputs the tuple $o_1 = \langle Resolve_1, Resolve_2, Resolve_3\rangle$ containing the complete resolutions just described as an opinion based on the initial fact situation X_1.

Where $o = \langle Resolve_1, Resolve_2, \ldots, Resolve_m\rangle$ is an opinion developed in the setting of a single-issue hierarchy centered around the issue π/δ, the *final resolution*—that is, $Resolve_m$—of this opinion will contain a single decision of the form $d = \langle Y_{m-1}, r, s\rangle$, where s is either π or δ. We say that the opinion o itself *supports* π or δ as its outcome accordingly. This idea is illustrated by the opinion o_1 above, containing, as we have seen, the set $Resolve_3 = \{d_4\}$ with $d_4 = \langle Y_2, r_4, \pi\rangle$ as its final resolution, so that the opinion o_1 itself supports π as the ultimate outcome.

We can now, at last, define a *case* in the context of a hierarchy as a triple of the form $c = \langle X, o, s\rangle$, where X is a base-level fact situation, where o is an opinion based on X in that hierarchy, and where s is the outcome supported by o. Given such a case c, we have, as usual, $Facts(c) = X$ and $Outcome(c) = s$, and we can now define *Opinion* as a function extracting the opinion of a case, so that $Opinion(c) = o$. These definitions can be illustrated with the concrete case $c_1 = \langle X_1, o_1, \pi\rangle$, where $X_1 = \{f_1, f_3, f_4, f_5\}$ is our original fact situation, where $o_1 = \langle Resolve_1, Resolve_2, Resolve_3\rangle$ is the opinion based on X_1 specified above, and where π is the outcome supported by this opinion. The opinion o_1, then, represents a justification of the outcome π in the situation X_1 according to which it is decided, in stages, that this initial fact situation should be further characterized through the additional factors p and r', then q, and at last π, a decision for the plaintiff.

Finally, we define a case base Γ in the context of a hierarchy as a set of cases in the context of that hierarchy.

4 CONSTRAINT

In the standard setting, where the issue π/δ between the plaintiff and the defendant is the only concern, and the entire set of factors is partitioned into those favoring one side or the other, the reason model of constraint is based on the idea that decisions by earlier courts generate an ordering among reasons for one of these sides, with later courts then required to reach decisions that preserve consistency of that ordering. Exactly the same idea can be adapted to the hierarchical setting, but with two additional wrinkles. First, an opinion in this new setting will generally address, not just the single issue π/δ, but a variety of intermediate concerns, and so generate an ordering on the reasons favoring one side or the other of these various intermediate concerns as well. Second, in the hierarchical setting, it is possible that the same factors, and so the same reasons, might bear on different concerns. Previous decisions therefore generate a priority ordering on reasons, not absolutely, but only as they bear on particular concerns—and indeed, the very same reasons might be ordered differently with respect to different concerns.[7]

In order to implement the reason model in the hierarchical setting, then, we begin by defining the priority ordering of reasons bearing on a particular concern to be derived from a particular decision:

Definition 6 (Priority ordering relative to concern derived from decision). In the context of a hierarchy, let $s/\bar{s}$ be a concern, U and V reasons for the sides $\bar{s}$ and s of this concern respectively, and $d = \langle X, r, s\rangle$ a decision. Then the relation $<_d^{s/\bar{s}}$ representing the priority ordering on reasons relative to $s/\bar{s}$ and derived from d is defined by stipulating that $U <_d^{s/\bar{s}} V$ if and only $U \subseteq X$ and $Premise(r) \subseteq V$.

Now, how do we lift this ordering derived from an individual decision to an ordering derived from a case, and then to an ordering derived from a full case base? There are four steps. First, we collect together the individual decisions belonging to an opinion by defining the *merge* of that opinion as a set containing each decision from each complete resolution contained in that opinion. More exactly, where $o = \langle Resolve_1, Resolve_2, \ldots, Resolve_m\rangle$ is an opinion, we define

$$Merge(o) = \bigcup \{Resolve_i : 1 \leq i \leq m\}.$$

Second, since the function *Opinion* extracts the opinion from a case and *Merge* collects together the decisions belonging to an opinion, we can define the set containing the decisions belonging to the opinion from a case $c = \langle X, o, s\rangle$ simply as

$$Merge(Opinion(c)).$$

Where c is a precedent case, the various decisions from $Merge(Opinion(c))$ can be thought of as *precedent constituents*, to use Branting's phrase. As our third step, then, we can define a priority ordering derived from a precedent case in a way that takes each of its constituents into account, by stipulating that one reason has higher priority than another according to the case whenever that priority is supported by some decision belonging to the opinion from that case:

Definition 7 (Priority ordering relative to concern derived from case). In the context of a hierarchy, let $s/\bar{s}$ be a concern, U and V reasons for the sides $\bar{s}$ and s of this concern respectively, and $c = \langle X, o, s\rangle$ a case. Then the relation $<_c^{s/\bar{s}}$ representing the priority ordering on reasons relative to $s/\bar{s}$ and derived from c is defined by stipulating that $U <_c^{s/\bar{s}} V$ if and only if there is some decision d from $Merge(Opinion(c))$ such that $U <_d^{s/\bar{s}} V$.

Fourth and finally, the priority ordering on reasons can then be lifted from cases to case bases exactly as in the standard reason model:

Definition 8 (Priority ordering relative to concern derived from case base). In the context of a hierarchy, let $s/\bar{s}$ be a concern, U and V reasons favoring opposite sides of this concern, and Γ a case base. Then the relation $<_\Gamma^{s/\bar{s}}$ representing the priority ordering on reasons relative to $s/\bar{s}$ and derived from Γ is defined by stipulating that $U <_\Gamma^{s/\bar{s}} V$ if and only if there is some case c from Γ such that $U <_c^{s/\bar{s}} V$.

Once this definition is in place, we can define a case base as consistent just in case there is no concern relative to which the case base supports conflicting information about the priority of reasons—telling us that, for some pair of reasons, each has higher priority than the other relative to that concern:

Definition 9 (Case base consistency). Let Γ be a case base in the context of a particular hierarchy. Then Γ is inconsistent if and only if there is some concern $s/\bar{s}$ with U and V favoring opposite sides of this concern such that $U <_\Gamma^{s/\bar{s}} V$ and $V <_\Gamma^{s/\bar{s}} U$, and consistent otherwise.

This notion of case base consistency allows us to introduce the reason model in the usual way, first specifying the opinions on the basis of which the court is permitted to justify its judgments as those that maintain consistency of the background case base:

Definition 10 (Reason model: permitted opinions). In the context of a hierarchy, let Γ be a consistent case base and X a base-level fact situation confronting the court. Then against the background of Γ, the reason model permits the court to justify its judgment in X with the opinion o, based on X and supporting the side s, if and only if the augmented case base $\Gamma \cup \{\langle X, o, s\rangle\}$ is consistent.

And given this notion of permission, we can then say that the court is required—or *constrained*—simply to reach some permitted judgment.

These definitions can be illustrated with an abstract example. Still working in the context of the hierarchy from Figure 3, and against

[7] A longer version of this paper contains an intuitive example illustrating this possibility.

the background of the case base $\Gamma_1 = \{c_1\}$ with $c_1 = \langle X_1, o_1, \pi\rangle$ as the case described earlier, consider the new fact situation $X_2 = \{f_2, f_6\}$, and suppose the court would like to settle this situation on the basis of the opinion $o_2 = \langle \mathit{Resolve}_1, \mathit{Resolve}_2, \mathit{Resolve}_3\rangle$, where $\mathit{Resolve}_1 = \{d_5, d_6\}$, with $d_5 = \langle Y_0, r_5, p\rangle$ where $Y_0 = X_2$ and $r_5 = \{f_2\} \rightarrow p$, and with $d_6 = \langle Y_0, r_6, r'\rangle$ where $Y_0 = X_2$ and $r_6 = \{f_6\} \rightarrow r'$; where $\mathit{Resolve}_2 = \{d_7\}$, with $d_7 = \langle Y_1, r_7, q\rangle$ where $Y_1 = Y_0 \cup \{p, r'\}$ and $r_7 = \{p\} \rightarrow q$; and where $\mathit{Resolve}_3 = \{d_8\}$, with $d_8 = \langle Y_2, r_8, \delta\rangle$ where $Y_2 = Y_1 \cup \{q\}$ and $r_8 = \{r'\} \rightarrow \delta$. The result would be the case $c_2 = \langle X_2, o_2, \delta\rangle$ as a representation of the court's judgment.

By Definition 11, however, the reason model does not permit the court to proceed in this way, because the augmented case base $\Gamma_2 = \Gamma_1 \cup \{c_2\} = \{c_1, c_2\}$ would be inconsistent. To see this, we first note that $\{r'\} <^{\pi/\delta}_{d_4} \{q\}$ and $\{q\} <^{\pi/\delta}_{d_8} \{r'\}$ by Definition 6. From this it follows that $\{r'\} <^{\pi/\delta}_{c_1} \{q\}$ and $\{q\} <^{\pi/\delta}_{c_2} \{r'\}$ by Definition 7, since d_4 belongs to $\mathit{Merge}(\mathit{Opinion}(c_1)) = \{d_1, d_2, d_3, d_4\}$ while d_8 belongs to $\mathit{Merge}(\mathit{Opinion}(c_2)) = \{d_5, d_6, d_7, d_8\}$. It then follows that $\{r'\} <^{\pi/\delta}_{\Gamma_2} \{q\}$ and $\{q\} <^{\pi/\delta}_{\Gamma_2} \{r'\}$ by Definition 8, since both c_1 and c_2 belong to Γ_2, so that Γ_2 is inconsistent by Definition 9.

For a less abstract illustration, we turn to Prakken and Sartor's hypothetical example concerning the issue whether an individual who has spent time abroad has changed fiscal domicile with respect to income tax. Here, the plaintiff is the individual's native country, which is arguing against change of domicile in order to tax the individual's income, and the defendant is the individual, who is arguing for change of domicile in order to pay, we suppose, the lower tax rates available in the foreign country. For simplicity, we consider just a fragment of Prakken and Sartor's example, depicted in Figure 4, in which the only matters bearing on the ultimate issue π/δ between the plaintiff and defendant are whether or not the individual showed an intention to return and whether or not, while abroad, the individual was employed by a domestic company. The first of these is represented as the intermediate concern f_{19}/f'_{19}, where the factor f_{19} represents intention to return and f'_{19} is its contrary; the second is represented as the intermediate concern f_3/f'_3, where f_3 indicates that the company is domestic and f'_3 that it is not. In addition, there are a number of base-level factors bearing directly on these two intermediate concerns, as indicated in the figure. These are: f_1, indicating that the individual retained their domestic home; f_{20}, indicating that the individual sold their domestic car; f_{21}, indicating that the individual sold their domestic bicycle; f_{11}, indicating that the company the individual worked for while abroad had a domestic president; f_{10}, indicating that the company the individual worked for while abroad had foreign headquarters; and f_{22}, indicating that the company the individual worked for while abroad had a largely foreign workforce.[8]

Now suppose that the first fact situation in this domain coming before a court is $X_3 = \{f_1, f_{20}, f_{11}, f_{10}\}$, representing an individual who has retained their domestic house but not their car, and who,

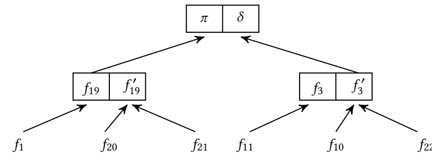

Figure 4: Change of fiscal domicile?

while abroad, worked for a company with a domestic president but foreign headquarters. Suppose the court resolves this case for the plaintiff on the grounds that the retained house showed intention to return and that, while the foreign headquarters showed that the company was foreign, intention to return is stronger as a reason for the plaintiff than foreign company is as a reason for the defendant. More exactly, suppose the court issues an opinion $o_3 = \langle \mathit{Resolve}_1, \mathit{Resolve}_2\rangle$, where $\mathit{Resolve}_1 = \{d_9, d_{10}\}$, with $d_9 = \langle Y_0, r_9, f_{19}\rangle$ where $Y_0 = X_3$ and $r_9 = \{f_1\} \rightarrow f_{19}$, and with $d_{10} = \langle Y_0, r_{10}, f'_3\rangle$ where $Y_0 = X_3$ and $r_{10} = \{f_{10}\} \rightarrow f'_3$; and where $\mathit{Resolve}_2 = \{d_{11}\}$, with $d_{11} = \langle Y_1, r_{11}, \pi\rangle$ where $Y_1 = Y_0 \cup \{f_{19}, f'_3\}$ and $r_{11} = \{f_{19}\} \rightarrow \pi$. This opinion would lead to $c_3 = \langle X_3, o_3, \pi\rangle$ as a settled case.

Next, imagine that, against the background of the case base $\Gamma_3 = \{c_3\}$ containing this case, another court confronts the situation $X_4 = \{f_1, f_{20}, f_{11}, f_{10}, f_{22}\}$, describing an individual exactly like the individual described in X_3, except that the company employing this new individual while abroad has, in addition, a foreign workforce. And suppose this new court would like to decide this situation on the basis of the opinion $o_4 = \langle \mathit{Resolve}_1, \mathit{Resolve}_2\rangle$, where $\mathit{Resolve}_1 = \{d_{12}, d_{13}\}$, with $d_{12} = \langle Y_0, r_{12}, f_{19}\rangle$ where $Y_0 = X_4$ and $r_{12} = \{f_1\} \rightarrow f_{19}$, and with $d_{13} = \langle Y_0, r_{13}, f'_3\rangle$ where $Y_0 = X_4$ and $r_{13} = \{f_{10}, f_{22}\} \rightarrow f'_3$; and where $\mathit{Resolve}_2 = \{d_{14}\}$, with $d_{14} = \langle Y_1, r_{14}, \delta\rangle$ where $Y_1 = Y_0 \cup \{f_{19}, f'_3\}$ and $r_{14} = \{f'_3\} \rightarrow \delta$. The result of this opinion would be the case $c_4 = \langle X_4, o_4, \delta\rangle$, leading to the augmented case base $\Gamma_4 = \Gamma_3 \cup \{c_4\}$.

Can the court rule as it wishes? The answer is No—the new case base Γ_4 would be inconsistent. The reader is invited to verify this fact formally, but even from an informal standpoint the point should be clear: Although the new opinion o_4 puts forth, in a sense, a stronger reason than o_3 for the intermediate conclusion that the individual worked for a foreign company, it had previously been decided in o_3 that foreign company as a reason for the defendant is weaker than intention to return as a reason for the plaintiff, while o_4 relies on exactly the opposite judgment—that intention to return for the plaintiff is weaker than foreign company for the defendant.

Finally, imagine that, again working against the background of the case base $\Gamma_3 = \{c_3\}$, a new court confronts the situation $X_5 = \{f_1, f_{20}, f_{21}, f_{11}, f_{10}\}$, again describing an individual like the individual from X_3 except that this time the individual sold their domestic bicycle as well as their domestic car. And suppose this new court would like to decide this situation on the basis of the opinion $o_5 = \langle \mathit{Resolve}_1, \mathit{Resolve}_2\rangle$, where $\mathit{Resolve}_1 = \{d_{15}, d_{16}\}$, with $d_{15} = \langle Y_0, r_{15}, f'_{19}\rangle$ where $Y_0 = X_5$ and $r_{15} = \{f_{20}, f_{21}\} \rightarrow f'_{19}$, and with $d_{16} = \langle Y_0, r_{16}, f'_3\rangle$ where $Y_0 = X_5$ and $r_{16} = \{f_{10}\} \rightarrow f'_3$; and where $\mathit{Resolve}_2 = \{d_{17}\}$, with $d_{17} = \langle Y_1, r_{17}, \delta\rangle$ where $Y_1 = Y_0 \cup \{f'_{19}, f'_3\}$ and $r_{17} = \{f'_3\} \rightarrow \delta$. The result of this opinion would be the case $c_5 = \langle X_5, o_5, \delta\rangle$, leading to the augmented case base $\Gamma_5 = \Gamma_3 \cup \{c_5\}$.

[8] To the extent possible, the numbering of these factors corresponds to that in Prakken and Sartor's original description of this example, though we have introduced a few new factors not present in their description, with new numbers. In addition, Prakken and Sartor described their factor hierarchy as bearing on the issue of whether or not the individual in question has changed fiscal domicile, where the resolution of this issue then determines a judgment for plaintiff or defendant; in keeping with our convention, however, we present the hierarchy as bearing directly on the ultimate issue π/δ between plaintiff and defendant.

One again, can the court rule as it wishes? This time the answer is Yes—Γ_5 is consistent. Again, a formal verification will be left to the reader, but the point is simple: the fact that the individual sold their bike as well as car is enough to outweigh the fact that they retained their house, allowing the court to reject intention to return as an intermediate conclusion, and since it was intention to return that was previously taken to outweigh foreign company as a reason for the defendant, the path is now open to consistently conclude for the defendant on the basis of foreign company. While both X_4 and X_5 present, in a certain sense, stronger cases for the defendant than X_3, it is the distribution of factors among intermediate concerns that allows a decision for the defendant in X_5, but not in X_4.

5 FLATTENING AND CONSTRAINT

We have developed a version of the reason model applicable in a rich hierarchical setting, with a variety of intermediate concerns lying between base-level factors and ultimate issues. The theory set out here is a conservative extension of the standard reason model, applicable only in a flat setting, but we have also suggested—following Branting, and Bench-Capon and Atkinson—that the presence of these intermediate factors makes an important difference. But what kind of difference? One idea might be that the importance of these intermediate factors is entirely cognitive, rather than logical. Appeal to intermediate factors may help a court to structure its reasoning as it moves from a base-level fact situation to an overall judgment, but does not affect the meaning of the judgment itself; once a court has reached its judgment on the basis of a particular opinion, the intermediate decisions that guided the court's reasoning can be forgotten and the justifying opinion compressed—or *flattened*—into a single step.[9]

Our goal in this section is to show that this reductive suggestion is wrong: intermediate factors and complex opinions have logical, not just psychological, significance. In order to establish this result precisely, we define a notion of flattening that maps rich hierarchies, cases, and case bases into corresponding flattened structures. We then show that the reductive suggestion fails in two directions: First, constraint in a rich hierarchical setting does not entail constraint in the corresponding flattened setting. Second, and more surprisingly, new relations of constraint might appear among the corresponding flattened structures which were not present in the original hierarchical setting.

We define our notion of flattening, first of all, for hierarchies. Here it is useful to recall that we are focusing throughout on single-issue hierarchies, using the convention that such a hierarchy is organized around the single issue π/δ. The idea, then, is that the flattening of a hierarchy removes all intermediate factors and links every base-level factor directly to the top-level factor it favors:

Definition 11 (Flattening of a hierarchy). Where $\mathcal{H}$ is a single-issue hierarchy, the flattening of this hierarchy is the set $\mathit{Flatten}(\mathcal{H}) = \{s \rightarrow t : t \in \{\pi, \delta\}$ and $s \in \mathit{Favor}^t_{\mathcal{H}} \cap \mathit{Base}_{\mathcal{H}}\}$.

It is easy to see that the flattening of any such hierarchy is a standard hierarchy, in the sense of Definition 4. To illustrate, where $\mathcal{H}$ is the hierarchy from Figure 3, its flattening is $\mathit{Flatten}(\mathcal{H}) = \{f_1 \rightarrow \pi, f_2 \rightarrow \pi, f_3 \rightarrow \delta, f_4 \rightarrow \pi, f_5 \rightarrow \delta, f_6 \rightarrow \delta\}$.

Next, we extend the notion of flattening to cases from the hierarchical setting, beginning with the opinions found in these cases. Where o is the opinion from such a case $c = \langle X, o, s\rangle$, then, the flattening $\mathit{Flatten}(o)$ of this opinion is another opinion of the form $\langle \mathit{Resolve}\rangle$ where $\mathit{Resolve} = \{d\}$ with $d = \langle X, r, s\rangle$. Can it be any opinion of this form? Well, no: in order to be the flattening of o—as opposed to some random single step opinion based on X and favoring s—the rule r from the decision $d = \langle X, r, s\rangle$ contained in this opinion should be arrived at by keeping track of the reason that justifies s as an outcome in the single decision belonging to the final resolution of the original opinion o. To do this, we *project* the reason in question into the base-level fact situation X through the favoring relation, as follows:

Definition 12 (Projection of a reason). Let $c = \langle X, o, s\rangle$ be a case in the context of a hierarchy $\mathcal{H}$, with $d = \langle Y, r, s\rangle$ the single decision belonging to the final resolution of the opinion o from this case; further, let $r = V \rightarrow s$, so that V is the reason justifying s in the decision d. Then the projection

$$\mathit{Project}(V, X) = X \cap \bigcup_{v \in V} \mathit{Favor}^v_{\mathcal{H}}.$$

of V into the base-level fact situation X is the set of factors from X that favor some factor from V.

Once we understand what it means to project a high-level reason for an outcome onto a base-level reason for the same outcome that holds in the initial fact situation from a case, we can define the flattening of an opinion as follows:

Definition 13 (Flattening of an opinion). Let $c = \langle X, o, s\rangle$ be a case in the context of a hierarchy, with $d = \langle Y, r, s\rangle$ the single decision belonging to the final resolution of the opinion o from this case, and let $r = V \rightarrow s$. Then $\mathit{Flatten}(o)$ is the opinion $\langle \mathit{Resolve}\rangle$ where $\mathit{Resolve} = \{d'\}$, with $d' = \langle X, r', s\rangle$ where $r' = \mathit{Project}(V, X) \rightarrow s$.

This definition can be illustrated by considering the opinion o_1 from the earlier case $c_1 = \langle X_1, o_1, \pi\rangle$ in the context of the hierarchy $\mathcal{H}$ from Figure 3 and based on the fact situation $X_1 = \{f_1, f_3, f_4, f_5\}$, which contains $\mathit{Resolve}_3 = \{d_4\}$ as its final resolution, with $d_4 = \langle Y_2, r_4, \pi\rangle$ where $Y_2 = X_1 \cup \{p, r', q\}$ and $r_4 = \{q\} \rightarrow \pi$. Here, we can first identify $\mathit{Project}(\{q\}, X_1) = \{f_1\}$ as the projection of the high-level reason $\{q\}$ into the base-level fact situation X_1, so that the flattening of the opinion o_1 is $\mathit{Flatten}(o_1) = \langle \mathit{Resolve}\rangle$, where $\mathit{Resolve} = \{d'_4\}$ with $d'_4 = \langle X_1, r'_4, \pi\rangle$ where $r'_4 = \{f_1\} \rightarrow \pi$.

Having defined the notion of flattening of an opinion, we can now lift that notion to corresponding notions of flattening for a case and a case base in the expected way:

Definition 14 (Flattening of a case and of a case base). Let $c = \langle X, o, s\rangle$ be a case and Γ a case base in the context of a hierarchy $\mathcal{H}$. Then the flattening of c is $\mathit{Flatten}(c) = \langle X, \mathit{Flatten}(o), s\rangle$, and the flattening of the case base Γ is $\mathit{Flatten}(\Gamma) = \{\mathit{Flatten}(c) : c \in \Gamma\}$, both in the context of $\mathit{Flatten}(\mathcal{H})$.

These concepts can be illustrated by noting that $\mathit{Flatten}(c_1) = \langle X_1, \mathit{Flatten}(o_1), \pi\rangle$, and that, where $\Gamma_1 = \{c_1\}$, we have $\mathit{Flatten}(\Gamma_1) = \{\mathit{Flatten}(c_1)\}$.

At this point, we are ready to establish the central result of this section: the intermediate concepts involved in complex opinions

[9] This suggestion mirrors Goodhart's proposal [11] that the meaning, or *ratio decidendi*, of a case is exhausted by the connection between base-level factors and the ultimate judgment. Something along these lines also seems to be endorsed by Horty and Bench-Capon, who write concerning a complex chain of reasoning leading from a fact situation to an ultimate outcome that "a rule moving directly from the base level factors to the issue would be equivalent from a logical point of view" [13, p. 208].

have logical significance. More exactly, we establish, first of all, that constraint in a rich hierarchical setting does not entail constraint in the corresponding flattened setting—that, even though some judgment may not be permitted against the background of a case base, the flattened version of that judgment may be permitted against the background of the flattened version of that case base. And second, that new relations of constraint might appear among the corresponding flattened structures which were not present in the original setting—that, even though some judgment might be permitted against the background of a case base, the flattened version of that judgment may not be permitted against the background of the flattened version of that case base. Since the reason model defines a judgment as permitted whenever that judgment does not introduce inconsistency into a consistent case base, our two results follow from the following two-part fact:

Observation 1. *(1) It is possible for there to be a case base* Γ *and a case* c *in the context of the hierarchy* $\mathcal{H}$*, with* $Flatten(\Gamma)$ *and* $Flatten(c)$ *their flattenings in the context of* $Flatten(\mathcal{H})$*, such that* $\Gamma \cup \{c\}$ *is inconsistent while* $Flatten(\Gamma) \cup \{Flatten(c)\}$ *is consistent. (2) It is also possible for there to be a case base* Γ *and a case* c *in the context of the hierarchy* $\mathcal{H}$*, with* $Flatten(\Gamma)$ *and* $Flatten(c)$ *their flattenings in the context of* $Flatten(\mathcal{H})$*, such that* $\Gamma \cup \{c\}$ *is consistent while* $Flatten(\Gamma) \cup \{Flatten(c)\}$ *is inconsistent.*

We can verify the first part of Observation 1 by going back to the earlier cases $c_1 = \langle X_1, o_1, \pi\rangle$ and $c_2 = \langle X_2, o_2, \delta\rangle$ in the context of the hierarchy $\mathcal{H}$ from Figure 3 and based on the fact situations $X_1 = \{f_1, f_3, f_4, f_5\}$ and $X_2 = \{f_2, f_6\}$ respectively. We have already seen in Section 4 that, where $\Gamma_1 = \{c_1\}$, the augmented case base $\Gamma_1 \cup \{c_2\}$ in the context of $\mathcal{H}$ is inconsistent. What about the case base $Flatten(\Gamma_1) \cup \{Flatten(c_2)\}$ in the context of the flattened hierarchy $Flatten(\mathcal{H})$? Well, first, recall that $Flatten(\Gamma_1)$ is the case base $\{Flatten(c_1)\}$, where $Flatten(c_1) = \langle X_1, Flatten(o_1), \pi\rangle$ is the case where $Flatten(o_1)$ is the opinion $\langle Resolve\rangle$, where $Resolve = \{d'_4\}$ with $d'_4 = \left\langle X_1, r'_4, \pi\right\rangle$ where $r'_4 = \{f_1\} \rightarrow \pi$. Note also that it follows from our definitions that $Flatten(c_2) = \langle X_2, Flatten(o_2), \delta\rangle$ is the case where $Flatten(o_2)$ is the opinion $\langle Resolve\rangle$, where $Resolve = \{d'_8\}$ with $d'_8 = \left\langle X_2, r'_8, \delta\right\rangle$ where $r'_8 = \{f_6\} \rightarrow \delta$. At this point, we can simply observe that, because the fact situations X_1 and X_2 have no factors in common, the priority orderings among reasons derived from the decisions d'_4 and d'_8—and so, the priority orderings among reasons derived from the corresponding cases $Flatten(c_1)$ and $Flatten(c_2)$—cannot possibly support conflicting information. But then the case base $Flatten(\Gamma_1) \cup \{Flatten(c_2)\}$ has to be consistent.

Now, by itself, the first part of our result could be interpreted as confirming the intuition, convincingly defended by Bench-Capon and Atkinson, that "using issues rather than whole cases to constrain decisions will enable us to decide more cases" [4, p. 17]. This intuition is supported by the observation that fact situations that can be distinguished when described in terms of fine-grained, base-level factors may turn out to be legally equivalent when described in terms of higher-level concepts. This is exactly what happens in the case of the fact situations X_1 and X_2 from our previous example: When we ignore the hierarchy $\mathcal{H}$ and consider only base-level factors, X_2 is completely unrelated to—and thus distinguishable from—X_1. But when we consider how X_1 was decided in the context of the full hierarchy $\mathcal{H}$ and which intermediate-level concepts from that hierarchy are supported by X_2, the two fact situations turn out to be equivalent—according to the opinion o_1 justifying the case c_1, X_1 supports the intermediate-level factors p, q, and r', which are exactly the intermediate-level factors supported by X_2.

If we look at things in this way, it is indeed very intuitive to think that, in general, constraint in the full hierarchical setting must be stronger than constraint in the corresponding flattened setting. Surprisingly, however, it turns out that, at times, the opposite is true: according to the second part of Observation 1, there are situations in which constraint is stronger in the flattened setting. A simple example of a situation of this kind can be constructed, once again, in the context of the factor hierarchy $\mathcal{H}$ from Figure 3.

Suppose that c_6 is the case $c_6 = \langle X_6, o_6, \pi\rangle$ in the context of $\mathcal{H}$, where $X_6 = \{f_1, f_4, f_5\}$ and $o_6 = \langle Resolve_1, Resolve_2, Resolve_3\rangle$, where $Resolve_1 = \{d_{18}, d_{19}\}$, with $d_{18} = \langle Y_0, r_{18}, p\rangle$ where $Y_0 = X_6$ and $r_{18} = r_1 = \{f_1\} \rightarrow p$, and with $d_{19} = \langle Y_0, r_{19}, r\rangle$ where $Y_0 = X_6$ and $r_{19} = \{f_4\} \rightarrow r$; where $Resolve_2 = \{d_{20}\}$, with $d_{20} = \langle Y_1, r_{20}, q\rangle$ where $Y_1 = X_6 \cup \{p, r\}$ and $r_{20} = r_3 = \{p\} \rightarrow q$; and where $Resolve_3 = \{d_{21}\}$, with $d_{21} = \langle Y_2, r_{21}, \pi\rangle$ where $Y_2 = X_6 \cup \{p, r, q\}$ and $r_{21} = r_4 = \{q\} \rightarrow \pi$. Let $\Gamma_6 = \{c_6\}$ be a new case base and $X_7 = \{f_1, f_5\}$ a new fact situation. Suppose the court would like to settle this situation on the basis of the opinion $o_7 = \langle Resolve_1, Resolve_2, Resolve_3\rangle$, where: $Resolve_1 = \{d_{22}, d_{23}\}$, with $d_{22} = \langle Y_0, r_{22}, p\rangle$, where $Y_0 = X_7$ and $r_{18} = r_1 = \{f_1\} \rightarrow p$, and with $d_{23} = \langle Y_0, r_{23}, r\rangle$, where $Y_0 = X_7$ and $r_{23} = r_2 = \{f_5\} \rightarrow r'$; $Resolve_2 = \{d_{24}\}$, with $d_{24} = \langle Y_1, r_{24}, q\rangle$, where $Y_1 = X_7 \cup \{p, r'\}$ and $r_{24} = r_3 = \{p\} \rightarrow q$; and $Resolve_3 = \{d_{25}\}$, with $d_{25} = \langle Y_2, r_{25}, \delta\rangle$, where $Y_2 = X_7 \cup \{p, r', q\}$ and $r_{25} = r_8 = \{r'\} \rightarrow \delta$. Is the court permitted, in the context of $\mathcal{H}$, to decide X_7 for δ on the basis of the opinion o_7? The answer is Yes—where c_7 is the case $\langle X_7, o_7, \delta\rangle$, the augmented case base $\Gamma_7 = \Gamma_6 \cup \{c_7\}$ in the context of $\mathcal{H}$ is consistent. A detailed verification is left to the reader, but the key point is that, because the only concerns that are resolved differently in o_6 and o_7 are r/r' and π/δ, the only pairs of decisions that might make Γ_7 inconsistent are either d_{19} and d_{23} or d_{21} and d_{25}. But no inconsistency derives from d_{19} and d_{23}, because the reason $\{f_4\}$, which grounds a decision for r in d_{19}, does not hold in $Facts(d_{22}) = X_7$; and no inconsistency derives from d_{21} and d_{23} either, because the reason $\{r'\}$, which grounds a decision for δ in d_{25}, does not hold in $Facts(d_{21}) = X_6 \cup \{p, r, q\}$.

We can see informally that things are different in the flattened setting: First, $Flatten(c_6) = \langle X_6, Flatten(o_6), \delta\rangle$ is the case in the context of $Flatten(\mathcal{H})$ where $Flatten(o_6) = \langle Resolve\rangle$, where $Resolve = \{d'_{21}\}$ with $d'_{21} = \left\langle X_6, r'_{21}, \pi\right\rangle$ where $r'_{21} = r'_4 = \{f_1\} \rightarrow \pi$. Similarly, $Flatten(c_7) = \langle X_7, Flatten(o_7), \delta\rangle$ is the case in the context of $Flatten(\mathcal{H})$ where $Flatten(o_7) = \langle Resolve\rangle$, where $Resolve = \{d'_{25}\}$ with $d'_{25} = \left\langle X_7, r'_{25}, \delta\right\rangle$ where $r'_{25} = r'_8 = \{f_5\} \rightarrow \delta$. At this point we can apply Definition 6 and observe that, since $\{f_5\}$ is a reason for δ that holds in X_6, the reason $\{f_1\}$ has higher priority than $\{f_5\}$ relative to the concern π/δ according to the case $Flatten(c_6)$; and, in turn, since $\{f_1\}$ is a reason for π that holds in X_7, the reason $\{f_5\}$ has higher priority than $\{f_1\}$ relative to the concern π/δ according to the case $Flatten(c_7)$. But then, where $Flatten(\Gamma_6) = \{Flatten(c_6)\}$, the case base $Flatten(\Gamma_6) \cup \{Flatten(c_7)\}$ in the context of the flattened hierarchy $Flatten(\mathcal{H})$ is inconsistent.

6 CONCLUSION

Our goal in this paper has been twofold: first, to develop the reason model of constraint in the context of a hierarchy of intermediate legal concerns lying between base-level factors and ultimate issues and, second, to explore the relation between constraint in the full hierarchical setting and constraint in the corresponding flattened setting. Our central finding has been that the presence of intermediate legal concerns has logical significance, both in the sense that, as suggested by Branting as well as Bench-Capon and Atkinson, there might be constraint in the hierarchical setting that is lost in the corresponding flattened setting and, more surprisingly, in the sense that new patterns of constraint might appear in the flattened setting that were not present in the original hierarchical setting. The theory described in this paper is currently being implemented in Python; the implementation verifies properties of a factor hierarchy, such as acyclicity and factor uniformity, and then, given a case base and a new fact situation, generates opinions based on that fact situation, and checks whether the priorities derived from these opinions are consistent with those derived from the case base.

Our theory opens a number of problems, both technical and conceptual. We close simply by mentioning two of these.

As an example of a technical problem: The standard reason model of constraint depends on factors; but where do these factors come from? A good deal of recent research in AI and Law has been devoted to discovering the set of factors at work in a domain using NLP and machine learning techniques [7, 8, 15]. If, however, the hierarchical version of the reason model is right and legal reasoning is best understood as involving, not just a flat set of factors, but an entire factor hierarchy, then this raises the difficult issue of how NLP and machine learning techniques could be adapted to discovering this kind of hierarchical information.

As an example of a conceptual problem: The opinions we have described here could be called *extensive opinions*, because they include every decision—or precedent constituent—that might possibly affect the ultimate outcome. But of course, we realize that not every decision from an opinion is binding on future courts—not every precedent constituent is part of the *ratio* of a case. Indeed, the jurisprudential literature contains a variety of suggestions for isolating the components of opinions that should be viewed as part of the *ratio*.

One interesting example is the test introduced by Wambaugh [19], discussed in the legal literature by Cross and Harris [9] and in AI and Law by Branting, according to which, as Branting writes, "if the deciding court could have believed the negation of [some] proposition without changing the outcome of the case, the proposition is *dictum* rather than *ratio*" [6, p. 41]. We can illustrate the effect of this test in the current framework by returning to the opinion o_1 from the case $c_1 = \langle X_1, o_1, \pi \rangle$. This opinion contains a number of decisions, including d_2, which, as we have seen, tells us that the reason $\{f_5\}$ has higher priority than the reason $\{f_4\}$ relative to the concern r/r', so that this priority should constrain future courts. But it is easy to see that, if this concern had been decided the other way, the ultimate decision in this case would have remained the same, since the case c_1 was ultimately decided for π and resolving the intermediate concern r/r' for r rather than r' simply strengthens the overall argument for π. On a natural interpretation of Wambaugh's test, then, the priority ordering derived from d_2 should be viewed as dictum, and should not constrain future courts.

Wambaugh's test can be thought of as one way of trimming, or pruning, the very extensive opinions defined in this paper to yield something more closely approximating the traditional concept of a *ratio*. There are surely other sensible modifications—we have not tried to explore the underlying jurisprudential questions in any depth here. Instead, we hope only to have introduced a framework in which different definitions of the concept, and so different answers to these jurisprudential questions, can be explored in a precise way.

ACKNOWLEDGMENTS

This research is partly supported by the project "Responsible Artificial Agency: A Logical Perspective," funded by the Research Priority Area Human(e) AI at the University of Amsterdam.

REFERENCES

[1] Vincent Aleven. 1997. *Teaching Case-Based Argumentation Through a Model and Examples.* Ph.D. Dissertation. Intelligent Systems Program, University of Pittsburgh.
[2] Vincent Aleven and Kevin Ashley. 1997. Evaluating a learning environment for case-based argumentation skills. In *Proceedings of the Sixth International Conference on Artificial Intelligence and Law (ICAIL-97).* The Association for Computing Machinery Press, 170–179.
[3] Kevin Ashley. 1990. *Modeling Legal Argument: Reasoning with Cases and Hypotheticals.* The MIT Press.
[4] Trevor Bench-Capon and Katie Atkinson. 2021. Precedential constraint: the role of issues. In *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Law (ICAIL 2021).* The Association for Computing Machinery Press, 12–21.
[5] L. Karl Branting. 1991. Reasoning with portions of precedents. In *Proceedings of the Third International Conference on Artificial Intelligence and Law (ICAIL-91).* ACM Press, 145–154.
[6] L. Karl Branting. 1993. A reduction-graph model of ratio decidendi. In *Proceedings of the Fourth International Conference on Artificial Intelligence and Law (ICAIL-93).* ACM Press, 40–49.
[7] L. Karl Branting, Craig Pfeifer, Bradford Brown, Lisa Ferro, John Aberdeen, Brandy Weiss, Mark Pfaff, and Bill Liao. 2021. Scalable and explanable legal prediction. *Artificial Intelligence and Law* 29 (2021), 213–238.
[8] Stefanie Brüninghaus and Kevin Ashley. 1993. Predicting outcomes of case based legal arguments. In *Proceedings of the Fourth International Conference on Artificial Intelligence and Law (ICAIL-93).* The Association for Computing Machinery Press, 233–242.
[9] Rupert Cross and J. W. Harris. 1991. *Precedent in English Law.* Oxford University Press. Fourth edition.
[10] Jonathan Dancy. 2004. *Ethics Without Principles.* Oxford University Press.
[11] Arthur Goodhart. 1930. Determining the *ratio decidendi* of a case. *Yale Law Journal* 40 (1930), 161–183.
[12] John Horty. 2011. Rules and reasons in the theory of precedent. *Legal Theory* 17 (2011), 1–33.
[13] John Horty and Trevor Bench-Capon. 2012. A factor-based definition of precedential constraint. *Artificial Intelligence and Law* 20 (2012), 181–214.
[14] John Horty, Richmond Thomason, and David Touretzky. 1990. A skeptical theory of inheritance in nonmonotonic semantic networks. *Artificial Intelligence* 42 (1990), 311–348. A preliminary version appears in *Proceedings of the Sixth National Conference on Artificial Intelligence (AAAI-87)*, Morgan Kaufmann Publishers, 1987, pp. 358–363.
[15] Jack Mumford, Katie Atkinson, and Trevor Bench-Capon. 2022. Reasoning with legal cases: A hybrid ADF-ML approach. In *Legal Knowledge and Information Systems. JURIX 2022: The Thirty-Fifth Annual Conference (Frontiers in Artificial Intelligence and Applications)*, Vol. 363. IOS Press, 90–102.
[16] Henry Prakken and Giovanni Sartor. 1998. Modelling reasoning with precedents in a formal dialogue game. *Artificial Intelligence and Law* 6 (1998), 231–287.
[17] Edwina Rissland and Kevin Ashley. 1987. A case-based system for trade secrets law. In *Proceedings of the First International Conference on Artificial Intelligence and Law (ICAIL-87).* The Association for Computing Machinery Press, 60–66.
[18] Abraham Cornelis Roth. 2003. *Case-based reasoning in the law: A formal theory of reasoning by case comparison.* Ph.D. Dissertation. Universiteit Maastricht.
[19] Eugene Wambaugh. 1891. *The Study of Cases.* Little, Brown (Boston).