

Open Texture and Precedential Constraint

John Horty

*** DRAFT ***

Version of: February 28, 2023

Contents

1	Introduction	1
2	Basic concepts	7
2.1	Factors and fact situations	7
2.2	Reasons, rules, cases, case bases	9
3	Constraint by reasons	11
3.1	A priority ordering on reasons	12
3.2	Constraint	14
3.3	Requirements and permissions	17
4	Open texture	19
4.1	A semantic account	19
4.2	The Super Scoop	22
4.3	A comparison to Hart	26
5	Discussion	30

1 Introduction

Suppose that two parents have established a household rule according to which their children are allowed to go out and play Saturday mornings only if their rooms are clean. What does it mean, in this setting, for a child’s room to be “clean”—how can it be determined whether this predicate applies to a child’s room? The parents might attempt a definitional account, perhaps stipulating that a child’s room is clean just in case the floor is vacuumed and the bed is made up with fresh sheets. But what if the shelves have not been dusted and are covered with clutter? On the other hand, what if the bed is not made up with fresh sheets, but the reason is that no fresh sheets are available and the washing machine is in use? The parents might refine their initial definition, perhaps leading to: a child’s room is clean just in case the floor is vacuumed, shelves are dusted, and the bed is made up with fresh sheets unless no fresh sheets are available and the laundry room is busy. But what if the trash has not been emptied? What if clothes are not folded and put away? Given the unbounded collection of possible complicating considerations, it is hard to imagine how any definitional account of what it means for a child’s room to be clean could be successful.¹

The phenomenon at work in this example is what H. L. A. Hart describes as *open texture*, a feature of ordinary predicates that he illustrates with his famous example of vehicles in the park:

A legal rule forbids you to take a vehicle into the public park. Plainly this forbids

¹Recent discussions of the problems confronting definitional accounts of predicate meaning can be found in Chapter 1 of Elbourne (2011) and throughout Ludlow (2014); a classic account, focusing on theories of sentence comprehension and concept learning, is provided by Fodor et al. (1980).

an automobile, but what about bicycles, roller skates, toy automobiles? What about airplanes? Are these, as we say, to be called “vehicles” for the purpose of the rule or not?

And just as famous as Hart’s example is his semantic proposal—involving a “core” and a “penumbra”—for understanding the meaning of open-textured predicates:

If we are to communicate with each other at all, and if, as in the most elementary form of law, we are to express our intentions that a certain type of behavior be regulated by rules, then the general words we use—like “vehicle” in the case I consider—must have some standard instance in which no doubts are felt about its application. There must be a core of settled meaning, but there will be, as well, a penumbra of debatable cases in which words are neither obviously applicable nor obviously ruled out. These cases will each have some features in common with the standard case; they will lack others or be accompanied by features not present in the standard case.²

According to Hart’s proposal, then, an open-textured predicate—such as “clean,” applied to a child’s room—is associated with a core of settled meaning, which determines a set of cases to which the predicate clearly applies, as well as a set of cases to which it clearly fails to apply. The predicate would clearly apply, for example, to a glittering room: bed crisply

²The passages quoted in this paragraph are from Section 3 of Hart (1958), where he first discusses the concept of an open-textured predicate, although in this paper he describes these predicates using the phrase “open character” instead. This discussion is then elaborated upon and extended in Chapter 7 of Hart (1961), where the concept of open character is now described as “open texture,” a phrase that Hart adopted from Waismann (1945).

made, fresh sheets, floor perfectly vacuumed, clothes neatly folded and put away, shelves dusted, trash properly disposed of. The predicate would clearly fail to apply to a filthy and chaotic room: bed unmade, dirty sheets, clothes and trash scattered around an unvacuumed floor, cluttered, dust-covered shelves. In addition to these clear cases, however, Hart's view allows for a range of penumbral cases to which the predicate neither clearly applies nor clearly fails to apply. It is not hard to imagine that the room of a typical child would fall within this penumbra: bed sloppily made though perhaps with fresh sheets, floor vacuumed toward the center but debris visible around the edges, trash disposed of, shelves still cluttered but haphazardly dusted.

Although Hart illustrates his concept of open texture with the hypothetical example of vehicles in a park, the problems of determining applicability of particular open-textured predicates in various penumbral situations are common in the law. Sometimes these problems can seem to be comical, even ludicrous, until the stakes are appreciated. For example, the British court system once considered the question whether Pringles could properly be classified as "potato chips." The reason this question found its way into the courts is that, in the United Kingdom, food is generally exempt from the value-added tax, with only a few exceptions—including potato chips. In an effort to avoid this tax, amounting to roughly \$160 million, the manufacturers of Pringles were therefore intent on establishing that Pringles should be classified, not as potato chips, but rather as "savory snacks," on the grounds that they contain corn, rice, and wheat, in addition to potato flour.³ At other times, the importance of the problems involved in determining the applicability of open-textured predicates is

³After multiple levels of appeal, this effort failed, with the result that Pringles were officially classified as "potato chips" and the manufacturers were forced to pay a value-added tax; see Cohen (2009).

almost self-evident. These include the various cases in employment law testing the distinction between “employees” and “contractors,” as well as the range of cases exploring applicability conditions for socially fraught predicates such as “marriage” or “rape” or “person.”⁴

Because of the intrinsic interest and practical importance of the issues surrounding open-textured predicates, a substantial literature on the topic has evolved within legal theory.⁵ For the most part, however, this literature focuses on what might be thought of as broader issues related to open texture—the role of defeasible legal rules, policy arguments concerning the application of these rules, the impact of open-textured predicates on theories of legal interpretation. The legal literature on the topic does not provide anything like a semantic account of open-textured predicates, or at least, not in the sense that a contemporary semanticist would recognize.

The goal of the present paper is to offer such an account, particularly of open-textured predicates in the law, but an account that may be applicable to uses of these predicates in language more broadly.⁶ The central idea is that judgments involving open-textured predicates—whether Pringles are potato chips, whether a child’s room is clean—are evaluated against a background set of previous authoritative decisions involving these predicates, and

⁴A useful discussion of the changing conditions for applicability of the predicates “rape” and “person” can be found in Schiappa (2003).

⁵Some highlights include Baker (1977), Bix (1991), MacCormick (1991), Lyons (1999), Tur (2001), Schauer (2008), and Schauer (2013).

⁶The idea that an account along these lines can be applied in language more broadly, not just legal language, is originally due to Cumming, and currently under development in joint work; see Cumming (20xx) for an initial proposal. A difference between Cumming’s work and the current paper is that Cumming develops his proposal using a full default logic—much more expressive but also more complicated than the special-purpose formalism employed here, which corresponds only to a fragment of default logic.

that these previous decisions then constrain later applications of the same predicates in exactly the way that precedent cases constrain later decisions in the common law.

Because this account draws on the mechanism of precedential constraint to help explain the use of open-textured predicates, it falls within a strong tradition of research connecting work in the philosophy of language with issues in legal theory.⁷ Much of this work concentrates on the illumination, or lack thereof, to be derived from an application within legal theory of ideas originally developed in logic or the philosophy of language—such as formal treatments of vagueness, for example.⁸ The present paper moves in the opposite direction, applying ideas first developed in the study of legal reasoning to illuminate an issue within the philosophy of language itself, the phenomenon of open texture. What makes this shift in explanatory direction possible is a growing body of research in the field of artificial intelligence and law that has brought a new precision to the study of legal reasoning, and led to the development of ideas and tools that can then be applied elsewhere.

In particular, the account of open texture presented here is based on a treatment of precedential constraint—characterized as the *reason model* of constraint—that is derived directly from recent research in artificial intelligence and law, as well as from an earlier proposal due to Grant Lamond.⁹ According to the reason model, what matters about a precedent

⁷See Endicott (2022) for an overview.

⁸A study of vagueness in the law from a perspective that combines legal, linguistic, and logical considerations is presented in Endicott (2000); a later collection on the same topic is found in Keil and Poscher (2016).

⁹See Lamond (2005) for his initial proposal. The first version of the account presented here is found in Horty (2011), later developed in Horty (2015); a book-length exposition is presented in Horty (20xx). This account has been related to other approaches from artificial intelligence and law in Horty and Bench-Capon (2012), compared to analogical approaches in Rigoni (2014), limited in scope in Broughton (2019), and

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

Because the goal of this paper is to show how the reason model of precedential constraint can be adapted to provide a semantic account of open-textured predicates, we begin with a brief but precise formulation of the reason model itself. This is accomplished in the next two sections, with Section 2 presenting the formal framework within which this model is developed, and then Section 3 moving through the series of definitions that constitute the model. Section 4 then shows how the ideas from the reason model can be generalized to a semantic treatment of open-textured predicates, by providing an account of the constraints governing their applicability. Section 5 concludes with a discussion of some open issues and directions for future work.

explored from a formal perspective in Prakken (2021). More recently, a different interpretation of Lamond's original proposal, and one that connects it more closely with traditional ideas from legal theory, has been presented in Mullins (2020).

2 Basic concepts

2.1 Factors and fact situations

We suppose that a situation presented to a court for decision can be represented as a set of *factors*, where a factor is a legally significant fact or pattern of facts bearing on that decision. This style of representation has been used to analyze case-based reasoning in a number of complex legal domains within artificial intelligence and law, where it originated in the work of Edwina Rissland and Kevin Ashley.¹⁰ Cases in different areas of the law will be characterized by different sets of factors, of course. In the domain of trade-secrets law, for example, where the factor-based analysis has been explored most extensively, a case typically concerns the issue of whether the defendant has gained an unfair competitive advantage over the plaintiff through the misappropriation of a trade secret; and here the factors involved might turn on, say, questions concerning whether the plaintiff took measures to protect the trade secret, whether a confidential relationship existed between the plaintiff and the defendant, whether the information acquired was reverse-engineerable or in some other way publicly available, and the extent to which this information did, in fact, lead to a real competitive advantage for the defendant.¹¹

¹⁰See Rissland and Ashley (1987) and Ashley (1989) for their initial proposals, Rissland (1990) for an overview of research in artificial intelligence and law that places this work in a broader perspective, Ashley (1990) for a canonical presentation, and then Rissland and Ashley (2002) for later reflections on the factor-based representation of legal information.

¹¹The most detailed analysis in this domain is presented by Alevan (1997), who analyzed 147 cases from trade-secrets law in terms of a factor hierarchy that includes five high-level issues, eleven intermediate-level concerns, and twenty-six base-level factors. The resulting knowledge base is used in an intelligent tutoring system for teaching elementary skills in legal argumentation, which has achieved results comparable to

Many factors can naturally be taken to have polarities, favoring one side or another. In the domain of trade-secrets law, the presence of security measures likewise favors the plaintiff, since it strengthens the claim that the information secured was a valuable trade secret; reverse-engineerability favors the defendant, since it suggests that the product information might have been acquired through legitimate means. As a simplification, we will assume, not just that many, or even most, factors have polarities, but that all factors are like this, favoring one particular side. In addition, we rely on the further simplifying assumption that the reasoning under consideration involves only a single step, proceeding at once from the factors present in a situation to a decision—directly in favor of the plaintiff or the defendant—rather than moving through a series of intermediate legal concepts.

Formally, then, we start by postulating a set of legal factors bearing on some particular issue. We will let $F^\pi = \{f_1^\pi, \dots, f_n^\pi\}$ represent the set of factors favoring the plaintiff and $F^\delta = \{f_1^\delta, \dots, f_m^\delta\}$ the set of factors favoring the defendant. Given our assumption that each factor favors one side or the other, the entire set $F^{\pi/\delta}$ of legal factors will be exhausted by those favoring the plaintiff together with those favoring the defendant: $F^{\pi/\delta} = F^\pi \cup F^\delta$. As this notation suggests, we take π and δ to represent the two sides in a dispute, plaintiff and defendant, and where s is one of these sides, we let \bar{s} represent the other: $\bar{\pi} = \delta$ and $\bar{\delta} = \pi$.

Based on this set $F^{\pi/\delta}$ of factors, we define a *fact situation* X of the sort presented to the court for judgment simply as some particular subset of the factor set: $X \subseteq F^{\pi/\delta}$. And where X is a fact situation of this kind, we let X^s represent the factors from X that support the side s , so that: $X^\pi = X \cap F^\pi$ and $X^\delta = X \cap F^\delta$. Of course, any interesting situation will contain factors favoring both sides of a given dispute. For example, the situation $X_1 =$

 traditional methods of instruction in controlled studies; see Aleven and Ashley (1997).

$\{f_1^\pi, f_2^\pi, f_1^\delta, f_2^\delta\}$ contains two factors each favoring the plaintiff and the defendant, with those factors favoring the plaintiff contained in $X_1^\pi = \{f_1^\pi, f_2^\pi\}$ and those favoring the defendant contained in $X_1^\delta = \{f_1^\delta, f_2^\delta\}$.

2.2 Reasons, rules, cases, case bases

When presented with a fact situation, a court's primary task is to reach a decision, or determine an outcome. Given our assumption that reasoning proceeds in a single step, we can suppose that the *outcome* of a case is a decision either in favor of the plaintiff or in favor of the defendant, with these two outcomes represented as π or δ respectively.

In addition to reaching a decision for one side or the other, we generally expect the court to supply a rule, or principle, to serve as justification for its decision.¹² Rules of this kind will be characterized in terms of reasons, where a *reason for a side* is some set of factors uniformly favoring that side; a *reason* can then be defined as a set of factors uniformly favoring one side or another. To illustrate: $\{f_1^\pi, f_2^\pi\}$ is a reason favoring the plaintiff, and so a reason.

Since reasons, like fact situations, are sets of factors, we can stipulate that a reason U *holds* in a situation X just in case each factor from U belongs to X , so that U is a subset of X , or $U \subseteq X$. And we can also define a relation of strength among reasons for a side according to which, where U and V are reasons for the same side, then V *is at least as strong a reason as U for that side* just in case U is a subset of V , or $U \subseteq V$. To illustrate: the

¹²Although I will refer to case rules as "rules," I take no stand on the question whether they should actually be classified as rules or as principles; I think of these case rules as relatively specific, a property associated with rules, as opposed to principles, by Raz (1972), but also as defeasible, a property associated with principles, as opposed to rules, by Dworkin (1967).

reason $\{f_1^\pi\}$ holds in the fact situation $X_1 = \{f_1^\pi, f_2^\pi, f_1^\delta, f_2^\delta\}$, since $\{f_1^\pi\} \subseteq X_1$, and that, of the two reasons $\{f_1^\pi\}$ and $\{f_1^\pi, f_2^\pi\}$, the second favors the plaintiff at least as strongly as the first, since $\{f_1^\pi\} \subseteq \{f_1^\pi, f_2^\pi\}$.

Given this notion of a reason, a *rule* can now be defined as a statement of the form $U \rightarrow s$, where U is a reason supporting the side s . For convenience, we introduce two auxiliary functions—*Premise* and *Conclusion*—picking out the premise and conclusion of a rule, so that, if r stands for the rule just mentioned, we would have $Premise(r) = U$ and $Conclusion(r) = s$. And we will say that a rule is *applicable* in a situation whenever the reason that forms its premise holds in that situation. To illustrate: The statement $\{f_1^\pi\} \rightarrow \pi$ is a rule, since $\{f_1^\pi\}$ is a reason supporting the plaintiff. If we take r_1 to stand for this rule, we would have $Premise(r_1) = \{f_1^\pi\}$ and $Conclusion(r_1) = \pi$. And r_1 is applicable in the situation X_1 above, since $Premise(r_1)$ holds in this situation.

The rules defined here are to be interpreted as defeasible, telling us that their premises entail their conclusions, not as a matter of necessity, but only by default. Continuing with our illustration, what the rule $r_1 = \{f_1^\pi\} \rightarrow \pi$ means, very roughly, is that, whenever the premise $\{f_1^\pi\}$ of the rule holds in some situation, then, as a default, the court ought to decide that situation for the conclusion π of the rule—or perhaps more simply, that the premise of the rule provides the court with a pro tanto reason for deciding in favor of its conclusion.¹³

On the basis of the concepts introduced so far—fact situations, rules, outcomes—a *case* can be defined as a situation together with an outcome and a rule through which that outcome is justified: such a case can be specified as a triple of the form $c = \langle X, r, s \rangle$, where X is a

¹³The connections among default rules, reasons, and oughts sketched in this paragraph are developed in detail in Horty (2012).

situation containing the factors presented to the court, r is a rule, and s is an outcome.¹⁴ For illustration, consider the case $c_1 = \langle X_1, r_1, s_1 \rangle$, where the fact situation of this case is the familiar $X_1 = \{f_1^\pi, f_2^\pi, f_1^\delta, f_2^\delta\}$, where the case rule is the familiar $r_1 = \{f_1^\pi\} \rightarrow \pi$, and where the outcome of the case is $s_1 = \pi$, a decision for the plaintiff. This particular case, then, represents a situation in which the court, when confronted with the fact situation X_1 , decided for the plaintiff by applying or introducing the rule r_1 , according to which the presence of the factor f_1^π —that is, the reason $\{f_1^\pi\}$ —leads, by default, to a decision for the plaintiff.

Finally, with this notion of a case in hand, we can now define a *case base* as a set Γ of precedent cases. It is a case base of this sort—a set of precedent cases—that will be taken to represent the common law in some area, and to constrain the decisions of future courts.

3 Constraint by reasons

According to the reason model, we recall, what matters about a precedent case is the precedent court’s assessment of the relative importance of the reasons presented by that case for each of the opposing sides. This assessment can be represented as a priority ordering on reasons, with later courts then required to reach decisions that are consistent with the

¹⁴Our representation of cases embodies the simplifying assumption that the particular rule underlying a court’s decision is plain, ignoring the extensive literature on methods for determining the *ratio decidendi* of a case; and we suppose, as a further simplification, that a case always contains a single rule, ignoring situations in which a court might offer several rules for a decision, or in which a court reaches a decision by majority, with different members of the court offering different rules, or in which a court might simply render a decision in a case without setting out any general rule at all.

priority ordering derived from the decisions of earlier courts.

In order to develop this idea, we need to explain how a priority ordering on reasons can be derived from the decisions of earlier courts, and then what it means for the decision of a later court to be consistent with that ordering.

3.1 A priority ordering on reasons

To begin with, then, let us return to the case $c_1 = \langle X_1, r_1, s_1 \rangle$ —where $X_1 = \{f_1^\pi, f_2^\pi, f_1^\delta, f_2^\delta\}$, where $r_1 = \{f_1^\pi\} \rightarrow \pi$, and where $s_1 = \pi$ —and ask what information is carried by this case; what is the court telling us with its decision? Well, two things. First of all, with its decision for the plaintiff on the basis of the rule r_1 , the court is registering its judgment that *Premise*(r_1), the reason for its decision, is more important—or has higher *priority*—than any reason for the defendant that holds in X_1 , the fact situation of the case.¹⁵ How do we know this? Because if the court had viewed some reason for the defendant that held in the situation X_1 as more important, or higher in priority, than *Premise*(r_1), the court would have found for the defendant on the basis of that reason, rather than for the plaintiff on the basis of *Premise*(r_1). And second, if the court is telling us explicitly that the reason *Premise*(r_1) itself has higher priority than any reason for the defendant that holds in X_1 ,

¹⁵When comparing the relative importance of reasons, it is more common to say that one carries greater weight than the other, or that one is weightier than the other. I prefer to speak in terms of priority, rather than weight, for two reasons: first, the priority ordering on reasons to be defined here is nonlinear, while the concept of weight tends to suggest linearity; second, the ordering to be defined here allows only ordinal comparisons among reasons, while the concept of weight suggests that cardinal comparisons must be available as well.

plaintiff that is at least as strong as $Premise(r_1)$ must likewise have a higher priority than any reason for the defendant that holds in this situation.

We can recall that a reason U for the defendant holds in the situation X_1 just in case $U \subseteq X_1$, and that a reason V for the plaintiff is at least as strong for the plaintiff as the reason $Premise(r_1)$ just in case $Premise(r_1) \subseteq V$. If we let the relation $<_{c_1}$ represent the priority ordering on reasons derived from the particular case c_1 , then, the force of the court's decision in this case is simply that: where U is a reason favoring the defendant and V is a reason favoring the plaintiff, we have $U <_{c_1} V$ just in case $U \subseteq X_1$ and $Premise(r_1) \subseteq V$. To illustrate: Consider the reason $\{f_1^\delta\}$ for the defendant and the reason $\{f_1^\pi, f_2^\pi, f_3^\pi\}$ for the plaintiff. Here, we have $\{f_1^\delta\} \subseteq X_1$ as well as $Premise(r_1) \subseteq \{f_1^\pi, f_2^\pi, f_3^\pi\}$. It therefore follows that $\{f_1^\delta\} <_{c_1} \{f_1^\pi, f_2^\pi, f_3^\pi\}$ —the court's decision in the case c_1 entails that the reason $\{f_1^\pi, f_2^\pi, f_3^\pi\}$ favoring the plaintiff is to be assigned a higher priority than the reason $\{f_1^\delta\}$ favoring the defendant.

Generalizing from this example, we reach the following definition of the priority ordering among reasons derived from a single case:

Definition 1 (Priority ordering derived from a case) Where $c = \langle X, r, s \rangle$ is a case and U and V are reasons favoring the sides \bar{s} and s respectively, the relation $<_c$ representing the priority ordering on reasons derived from the case c is defined by stipulating that $U <_c V$ if and only if $U \subseteq X$ and $Premise(r) \subseteq V$.

This priority ordering on reasons derived from a single case can be lifted to an ordering derived from an entire case base in the natural way, through the stipulation that one reason has a higher priority than another according to the case base whenever that priority is

supported by some case from the case base:

Definition 2 (Priority ordering derived from a case base) Where Γ a case base and U and V are reasons, the relation $<_{\Gamma}$ representing the priority ordering on reasons derived from the case base Γ is defined by stipulating that $U <_{\Gamma} V$ if and only if $U <_c V$ for some case c from Γ .

And using this concept of the priority ordering derived from a case base, we can now define a case base itself as inconsistent if the ordering it supports yields conflicting information about the priority among reasons—telling us, for some pair of reasons, that each has a higher priority than the other—and consistent otherwise:

Definition 3 (Inconsistent and consistent case bases) Where Γ is a case base with $<_{\Gamma}$ its derived priority ordering Γ is inconsistent if and only if there are reasons U and V such that $U <_{\Gamma} V$ and $V <_{\Gamma} U$, and consistent otherwise.

3.2 Constraint

We now present the reason model of constraint itself, building on the concept of case base consistency. The guiding idea, once again, is that, in deciding a case, a constrained court is required to preserve the consistency of the background case base. Suppose, more exactly, that a court constrained by a consistent background case base is confronted with a new fact situation. Then what the reason model tells us, in the first instance, is that the court is permitted to base its decision on a particular rule only if augmenting the background case base with a decision based on that rule maintains consistency:

Definition 4 (Reason model constraint on rule selection) Against the background of a consistent case base Γ , the reason model permits a court to base its decision in some situation X on the rule r , applicable in X and supporting the side s , if and only if the augmented case base $\Gamma \cup \{\langle X, r, s \rangle\}$ is consistent.

This definition can be illustrated by imagining that the background case base is $\Gamma_1 = \{c_1\}$, containing as its single member the familiar case $c_1 = \langle X_1, r_1, s_1 \rangle$ —where, again, $X_1 = \{f_1^\pi, f_2^\pi, f_1^\delta, f_2^\delta\}$, where $r_1 = \{f_1^\pi\} \rightarrow \pi$, and where $s_1 = \pi$. Suppose that, against this background, the court confronts the fresh situation $X_2 = \{f_1^\pi, f_2^\pi, f_1^\delta, f_2^\delta, f_3^\delta\}$ and considers finding for the defendant in this situation on the basis of the reason $\{f_1^\delta, f_2^\delta\}$, leading to the decision $c_2 = \langle X_2, r_2, s_2 \rangle$, where X_2 is as above, where $r_2 = \{f_1^\delta, f_2^\delta\} \rightarrow \delta$, and where $s_2 = \delta$. Is the court permitted to carry through with this plan, according to the reason model?

Well, as we can see, $Premise(r_1) = \{f_1^\pi\}$, the reason for the decision in the initial case, holds in the new situation X_2 as well, since $\{f_1^\pi\} \subseteq X_2$. And of course, the new reason $Premise(r_2) = \{f_1^\delta, f_2^\delta\}$ favors the defendant at least as strongly as itself—that is, $Premise(r_2) \subseteq Premise(r_2)$, or $Premise(r_2) \subseteq \{f_1^\delta, f_2^\delta\}$. It therefore follows from Definition 1 that c_2 , the court’s envisaged decision, would assign the reason $\{f_1^\delta, f_2^\delta\}$ for the defendant a higher priority than the reason $\{f_1^\pi\}$ for the plaintiff—that is, $\{f_1^\pi\} <_{c_2} \{f_1^\delta, f_2^\delta\}$. But Γ_1 already contains the case c_1 , from which, in a similar fashion, we can derive the priority relation $\{f_1^\delta, f_2^\delta\} <_{c_1} \{f_1^\pi\}$, telling us exactly the opposite. Since the augmented case base

$$\begin{aligned} \Gamma_2 &= \Gamma_1 \cup \{c_2\} \\ &= \{c_1, c_2\} \end{aligned}$$

resulting from the court's envisaged decision contains both these cases, we would then have both $\{f_1^\delta, f_2^\delta\} <_{r_2} \{f_1^\pi\}$ and $\{f_1^\pi\} <_{r_2} \{f_1^\delta, f_2^\delta\}$ by Definition 2, so that, by Definition 3, this augmented case base would be inconsistent. By Definition 4, then, we can conclude that the court is not permitted to decide for the defendant in the situation X_2 on the basis of the rule r_2 , since c_2 , the resulting decision, would introduce an inconsistency into the background case base.

Of course, it does not follow from the fact that the court is not permitted to decide the situation X_2 for the defendant on the basis of the particular rule r_2 that it is not permitted to decide this situation for the defendant at all—in this situation, there are other rules on the basis of which the court is permitted to reach a decision for the defendant. Suppose, for example, that the court considers finding for the defendant on the basis of the reason $\{f_1^\delta, f_3^\delta\}$, leading to the decision $c_3 = \langle X_3, r_3, s_3 \rangle$, where $X_3 = X_2$, where $r_3 = \{f_1^\delta, f_3^\delta\} \rightarrow \delta$, and where $s_3 = \delta$. The augmented case base

$$\begin{aligned}\Gamma_3 &= \Gamma_1 \cup \{c_3\} \\ &= \{c_1, c_3\}\end{aligned}$$

resulting from this decision would then be consistent. As before, the previous case c_1 supports the priority $\{f_1^\delta, f_2^\delta\} <_{c_1} \{f_1^\pi\}$, and the new decision c_3 would now support the priority $\{f_1^\pi\} <_{c_3} \{f_1^\delta, f_3^\delta\}$, so that we would then have both the case base priorities $\{f_1^\delta, f_2^\delta\} <_{\Gamma_3} \{f_1^\pi\}$ and $\{f_1^\pi\} <_{\Gamma_3} \{f_1^\delta, f_3^\delta\}$. But there is nothing inconsistent about this pair of priorities.

Now imagine that the court does, in fact, decide the situation X_2 in this way, augmenting the background case Γ_1 with the new decision c_3 , leading to the augmented case base $\Gamma_3 = \Gamma_1 \cup \{c_3\}$. According to the reason model, this decision would then represent a step in the normal development of a legal system, which proceeds more generally as follows: A court

confronts a new situation X against the background of a consistent case base Γ , with an associated ordering $<_{\Gamma}$ on reasons. The court is permitted to base its decision only on a rule r supporting an outcome s such that the case base $\Gamma' = \Gamma \cup \{\langle X, r, s \rangle\}$ is consistent, with the result that the background case base is augmented with this new decision. The next court confronting the next new situation Y must then work against the background of the augmented case base Γ' , which gives rise to the strengthened ordering $<_{\Gamma'}$ on reasons. This new court is likewise permitted to base its decision only on a rule r' supporting an outcome s' such that the case base $\Gamma'' = \Gamma' \cup \{\langle Y, r', s' \rangle\}$ is consistent, thus further augmenting the case base, further strengthening the underlying priority ordering on reasons, and the process continues.

The hypothesis of the reason model is that this is how the common law develops in the normal, incremental case—by building up a stronger and stronger priority ordering on reasons through a series of decisions that are, at each stage, consistent with the existing case base.

3.3 Requirements and permissions

Definition 4 characterizes only the rules on the basis of which a court is permitted to justify its decisions. But of course, once this idea is in place, it can be used to define the conditions under which a court is permitted, or required, to reach a decision for one side or another—through the natural stipulation that a court is permitted to reach a decision for a side if some rule on the basis of which it is permitted to justify its decision supports that side, and required to reach a decision for a side if every rule on the basis of which it is permitted to justify its decision supports that side:

Definition 5 (Reason model constraint on decision) Against the background of a consistent case base Γ , the reason model permits a court to decide the situation X for the side s if and only if some rule on the basis of which the court is permitted to decide that situation supports s . Likewise, the reason model requires the court to decide the X for the side s if and only if every rule on the basis of which the court is permitted to decide that situation supports s .

For illustration: We have seen that, against the background of Γ_1 , the court is permitted to decide the fact situation X_2 on the basis of the rule r_3 supporting δ , the defendant. And it is easy to see that the court is likewise permitted to decide this situation on the basis of r_1 , supporting π , the plaintiff. It follows from Definition 5, therefore, that the court is permitted to decide this situation for each side, but not required to decide for either. By contrast, suppose that against the background of the same case base, the court is now faced with the situation $X_4 = \{f_1^\pi, f_1^\delta\}$. It then follows that the only rule on the basis of which the court is permitted to justify its decision is r_1 , so that the court is required to decide this new situation for π , the plaintiff.

It is just worth noting that the notions of requirement and permission introduced in Definition 5 conform to the rules of standard deontic logic. We can see, for example, that a court is required to decide a situation for the side s just in case it is not permitted to decide that situation for \bar{s} , the opposite side. And as long as it is working against the background of a consistent case base, the court will never be required to decide the same situation for one side and also for the other; it will always be required to reach a decision only for one side, or required to reach a decision only for the other, or permitted to reach a decision for

either side.

4 Open texture

4.1 A semantic account

We now turn, at last, to the central task of this paper: showing how the treatment of precedential constraint sketched so far can be adapted to supply a semantic account of open-textured predicates. The first step is to interpret π and δ —previously regarded simply as grammatically indeterminate symbols indicating a decision for the plaintiff or the defendant—explicitly predicates, so that, where X is a fact situation, the application of π to X means that the situation is decided for the plaintiff, while the application of δ to X means that the situation is decided for the defendant. If π and δ are predicates, it seems clear that they must be open-textured predicates, since a judgment about their applicability in some situation is determined, not by appeal to definition, but by assessing the various competing considerations that might favor a decision for the plaintiff or the defendant. And it is clear also that the predicates π and δ are contraries, in the traditional sense that they cannot both apply in a particular situation, but that, at any given point, it may not yet be determined which applies.

Once we have agreed to regard π and δ as open-textured predicates, the next step is simply to generalize the analysis already set out for the particular predicates π and δ to open-textured predicates more broadly. We begin by stipulating that, just as π and δ can be thought of as contraries, each open-textured predicate p is associated with some contrary p' . To illustrate: If p represents the predicate “clean,” applied to a child’s room, then p'

represents the predicate “not clean.” If p represents the predicate “potato chips,” applied to a manufactured comestible, such as Pringles, then p' represents the predicate “not potato chips.” If p represents the predicate “employee,” applied to an individual performing a service for pay, such as an Uber driver, then p' represents the predicate “contractor.” A pair consisting of an open-textured predicate p and its contrary p' represents the two sides of a *dispute*. As before, we will let s range over these two sides, and where s is one of the sides, \bar{s} is the other: $\bar{p} = p'$ and $\overline{p'} = p$.

For each dispute between a pair of open-textured predicates p and p' , we postulate a set $F^p = \{f_1^p, f_2^p, \dots, f_n^p\}$ of factors favoring the decision that the predicate p should be applied to some object or situation under consideration, and a set $F^{p'} = \{f_1^{p'}, f_2^{p'}, \dots, f_n^{p'}\}$ of factors favoring the judgment that, instead, the predicate p' should be applied. If we take p and p' to represent “clean” and “not-clean,” for example, then F^p might include the factors that, in a particular child’s room, the bed is crisply made, or the floor carefully vacuumed, while $F^{p'}$ might include the factors that unfolded clothes are strewn about, or that trash has not been emptied. If we take p and p' to represent the predicates “potato chips” and “not potato chips,” then F^p might include the factor that a particular manufactured comestible contains at least 40% potato flour, while $F^{p'}$ might include the factor that it contains other ingredients as well, such as corn, rice, or wheat flour. If we take p and p' to represent the predicates “employee” and “contractor,” then F^p might include the factors that, for a particular individual, the company directs “when, where, and how” that individual’s work is done or that the individual is required to “undergo company-provided training,” while $F^{p'}$ might include the factors that there is no need for the individual in question to perform “on-site services” or that the individual performs the required services using “independently-

obtained supplies or tools.”¹⁶

Following our earlier pattern, we let $F^{p/p'} = F^p \cup F^{p'}$ represent the entire set of factors bearing on the dispute between p and p' . And we define a fact situation X that *gives rise* to this dispute as some subset of $F^{p/p'}$ —that is, $X \subseteq F^{p/p'}$ —divided into those factors $X^p = X \cap F^p$ favoring application of the predicate p and those factors $X^{p'} = X \cap F^{p'}$ favoring application of the predicate p' . Again, the most interesting situations are those containing factors favoring opposite sides of some dispute, such as the situation presented by a typical child’s room, as described earlier, the situation presented by Pringles, which contain 42% potato flour but substantial amounts of corn, wheat, and rice flour, and the situation presented by Uber drivers, who undergo company-provided training but perform services off site using their independently-provided vehicles.

From this point forward, the account set out already, centered around the dispute between application of the particular open-textured predicates π and δ , generalizes in a straightforward way to any dispute between application of the open-textured predicates p or p' more broadly. A reason U for a side s of the dispute between p and p' is defined as a set of factors uniformly favoring that side—that is, $U \subseteq F^s$ —and a reason bearing on this dispute is defined as a reason for one side of the dispute or the other. As before, the reason U is said to hold in a fact situation X just in case each factor from U belongs to X , or $U \subseteq X$. A rule for the side s of the dispute between p and p' has the form $U \rightarrow s$, where U is a reason for s , and such a rule is applicable in some situation X just in case the reason U that form its premise holds in that situation. A case *bearing on* the dispute between p and p' is

¹⁶These particular factors are extracted from the United States Internal Revenue Service 20-factor test for differentiating employees from contractors.

a structure of the form $c = \langle X, r, s \rangle$, where X is a fact situation giving rise to this dispute and r is a rule applicable in that fact situation and supporting the side s . And a case base Γ bearing on this dispute is a set of cases bearing on the dispute.

Exactly as before, a priority ordering among reasons supporting opposite sides of the dispute between p and p' , and derived from a particular case bearing on this dispute, can be set out as in Definition 1, and then extended to a priority ordering derived from a case base as in Definition 2. The notion of a consistent case base can be set out as in Definition 3. Finally, against the background of a case base Γ bearing on the dispute between p and p' , the rules on the basis of which the court is permitted to arrive at a decision in a particular situation X giving rise to this dispute can be specified as in Definitions 4, and the decisions that the court is required or permitted to reach specified as in Definition 5.

4.2 The Super Scoop

At this point, we shift from a child’s room, Pringles, and Uber drivers to another example, which we consider in detail because it illustrates the account developed here in a particularly clear way. In *Stewart v. Dutra Construction Company, Inc.*, a series of United States federal courts considered the question whether the Super Scoop—a dredge, at the time the largest in the world—could properly be classified as a “vessel.”¹⁷ This question was brought before the courts by Willard Stewart, a marine engineer working on the Super Scoop, who was injured on the job through, as he claimed, the company’s negligence, and sought compensation for damages. Stewart had two routes to recovery. He could file a claim through the Longshoreman and Harbor Workers’ Compensation Act, a federal statute that would

¹⁷543 U.S. 481 (2005).

provide the equivalent of workers' compensation, but would exclude negligence. Or he could file under the Jones Act, another federal statute specifically enacted to protect seamen, due to the extraordinary perils of work at sea, containing the language

Any seaman who shall suffer personal injury in the course of his employment may, at his election, maintain an action for damages at law, with the right of trial by jury. . .

and so allowing recovery for negligence.¹⁸

Because Stewart hoped to claim negligence under the Jones Act, it was necessary for him to establish that he had been employed by Dutra as a “seaman” at the time of his injury. Although this term is not defined in the Jones Act itself, a gloss on the statute specifies that whether or not an individual is a seaman depends on that individual’s connection with a vessel. The nature of Stewart’s connection with the Super Scoop was never an issue, since all parties acknowledged that he had been employed as a member of its crew. The question remained, however, whether the Super Scoop could legitimately be classified as a “vessel”—or more exactly, a “vessel in navigation”—as this predicate was understood in the Jones Act, and on that issue, there were considerations naturally favoring different sides. On one hand, the Super Scoop shared a number of characteristics with more typical vessels. It had a captain and crew, as well as various marine appurtenances, such as ballast tanks and navigation lights; and, importantly, it was registered with and subject to regulations of the United States Coast Guard. On the other hand, the Super Scoop was incapable of self-propulsion, but had to be towed from one location to another, and its primary purpose

¹⁸46 U.S.C. App. §688(a)

was construction, rather than navigation.

Stewart's suit against Dutra began in the District Court of Massachusetts, which found that the Super Scoop was not a vessel, so that Stuart could not proceed under the Jones Act, a decision that was upheld by the First Circuit Court of Appeals.¹⁹ The decision was then appealed again to the United States Supreme court, which reversed the Appeals Court judgment, ruling instead that the Super Scoop was a vessel, and allowing Stewart to proceed with his Jones Act suit.

We will not consider here the reasoning either of the District Court or of the Supreme Court, but focus only on the decision of the First Circuit Court of Appeals, which was explicitly based on the precedent established in *Di Giovanni v. Traylor Bros, Inc*, an earlier case before the same court, and dealing with the same issue.²⁰ This case concerned, not a dredge, but a barge, the Betty F, bearing a crane used for bridge construction. The Betty F was similar, in many ways, to the Super Scoop, with a captain and crew, requiring Coast Guard registration, but without the capacity for self-propulsion, and with construction rather than navigation as its primary business; in addition, at the time of the incident in question, the Betty F had been largely stationary for over a month. This incident occurred when Rocco Di Giovanni, a workman on the Betty F, slipped and fell, due to the negligence of Traylor in failing to address a hydraulic fluid leak. Like Stewart, Di Giovanni hoped to bring suit as a seaman under the Jones Act. Again, the sole point of contention was whether or not the Betty F could be classified as a vessel, a question that had found its way to the First Circuit Court of Appeals, which decided that the Betty F could not be so classified, on

¹⁹230 F.3d 461 (1st Cir. 2000).

²⁰959 F.2d 1119 (1st Cir. 1992).

the grounds that “if a barge, or other float’s ‘purpose or primary business is *not* navigation or commerce,’ then workers assigned thereto . . . are to be considered seamen only when it is in actual navigation or transit.”²¹ Confronted with an analogous issue in *Stewart*, the court felt that it was bound by its own precedent, and so concluded that the Super Scoop could not be classified as a vessel either.

To model, or at least approximate, the situation confronting the First Circuit Court in *Stewart* within the current framework, we let the open-textured predicates v and v' represent the judgments that some marine platform is or is not a vessel. Among the factors favoring v , that the object is a vessel, we let f_1^v indicate that it has a captain and crew and f_2^v that it is subject to Coast Guard regulations. Among the factors favoring v' , that the object is not a vessel, we let $f_1^{v'}$ indicate that it is not capable of self-propulsion, $f_2^{v'}$ that its primary business is not navigation, and $f_3^{v'}$ that it has been largely stationary for at least a month.

Using this notation, the situation presented by the Betty F to the *Di Giovanni* court can be represented as $X_5 = \{f_1^v, f_2^v, f_1^{v'}, f_2^{v'}, f_3^{v'}\}$ —and we simplify by imagining that the court was considering this situation against the background of an empty case base $\Gamma_4 = \emptyset$ containing no decisions at all concerning applicability of the predicate “vessel.”²² Given this information, and reasoning, as we imagine, against this background case base, the court then concluded that the Betty F should not be classified as a vessel on the grounds that its

²¹959 F.2d 1119,1123 (1st Cir. 1992).

²²This is a significant simplification, since, by the time of *Di Giovanni*, there was already a substantial case base concerning applicability of the predicate “vessel.” It is also worth noting that any decision is permitted against the background of an empty case base, so that the court is free to rely on its own judgment to decide whether an open-textured predicate is applicable; once this decision is incorporated into the case base, the court’s judgment then gains legal authority, constraining later decisions.

primary business was not navigation—that is, on the basis of the rule $r_5 = \{f_2^{v'}\} \rightarrow v'$ —leading to the decision $c_5 = \{X_5, r_5, s_5\}$, where X_5 and r_5 are as above and where $s_5 = v'$. The augmented case base resulting from this earlier decision, and constraining the reasoning of the later *Stewart* court, is therefore

$$\begin{aligned}\Gamma_5 &= \Gamma_4 \cup \{c_5\} \\ &= \{c_5\},\end{aligned}$$

with the situation presented by the Super Scoop to the *Stewart* court itself represented as $X_6 = \{f_1^v, f_2^v, f_1^{v'}, f_2^{v'}\}$, differing from that presented by the Betty F situation only in omitting $f_3^{v'}$, and so forming a slightly stronger case for the conclusion v . Nevertheless, as the reader can verify, the reason model of constraint requires a finding for v' in the situation X_6 considered against the background of Γ_5 —that is, a decision that the Super Scoop is not a vessel, just as the *Stewart* court itself concluded. In fact, the court justified its decision through a further application of the *Di Giovanni* rule, leading to $c_6 = \{X_6, r_6, s_6\}$ as the decision in *Stewart*, where X_6 is as above, where $r_6 = r_5$, and where $s_6 = v'$.

4.3 A comparison to Hart

With this semantic account of open-textured predicates before us, we can now draw a comparison with Hart’s own proposal, presented in the canonical passage quoted in the Introduction of this paper. In this passage, we recall, Hart argues that an open-textured predicate has a “core of settled meaning,” which determines a range of cases in which “no doubts are felt about its application,” but that such a predicate may also allow for a range of penumbral cases in which the predicate is neither “obviously applicable nor obviously ruled out.” The current account, however, does not postulate a separate core of settled meaning to determine

situations in which no doubts are felt about the applicability of an open-textured predicate. Instead, it relies only on a background set of precedent cases that requires the application of that predicate in certain situations, and requires the application of its contrary in others. In the same way, the present account does not postulate a set of penumbral situations in which an open-textured predicate is neither obviously applicable nor obviously ruled out, but supposes only that there may be a range of situations in which the background set of precedent cases requires application neither of the predicate nor of its contrary, but permits the decision to go either way.

I think of the current account of open texture as providing a sympathetic reconstruction of Hart, capturing in a formal semantic theory much of what is most important in his proposal. At the same time, I also want to argue that an explicit reliance on a background set of precedent decisions, as in the current account, has advantages over any appeal to a separate core of settled meaning. I will try to establish this point by, first, deflecting Hart's own argument, if it is interpreted as favoring a separate core of settled meaning, and then highlighting one benefit of relying, instead, on a background set of precedents.

We begin, then, with Hart's argument for a core of settled meaning. The argument is brief, and contained in the canonical passage already cited. Here, Hart takes as his premise the claim: "If we are to communicate with each other at all . . . then the general words we use—like 'vehicle' in the case I consider—must have some standard instances in which no doubts are felt about its application." And from this he moves directly to his conclusion: "There must be a core of settled meaning" But this argument fails if it is interpreted as favoring a separate core of settled meaning, apart from the background set of precedent cases. After all, in the example we have just considered, once the court had decided that

the Betty F cannot be classified as a vessel, it follows at once that it is no longer permissible to apply the predicate “vessel” to the Super Scoop either—there is, in Hart’s language, no longer any doubt about application of this predicate to the Super Scoop. Yet this judgment does not depend on any separate core meaning of the predicate “vessel,” but only on the relation between the situation at hand and the background set of precedents.

Turning now to our positive argument: the current account, with its explicit reliance on a background set of precedent decision, seems to allow a better explanation than an account based on a core of settled meaning for the linkage, or coordination, between judgments concerning application of open-textured predicates to different items that were originally in the penumbra. Imagine, for example, the state of affairs as it existed before applicability of the predicate “vessel” had been investigated for either of the two marine platforms under consideration, the Betty F and the Super Scoop—imagine, once again, that the background set of decisions on the issue was simply $\Gamma_4 = \emptyset$. At that point, it is natural to suppose that either decision concerning applicability of the open-textured predicate to each of these items would have been permissible, or in Hart’s terminology, that both would have fallen within the predicate’s penumbra. Once it was decided in *Di Giovanni* that the Betty F should not be classified as a vessel, however—that is, once the background case base had shifted from Γ_4 to $\Gamma_5 = \{c_5\}$ —the later *Stewart* court was required to reach the same decision concerning the Super Scoop, since the *Di Giovanni* rule applied to the Super Scoop as well, and the Super Scoop displayed no features on the basis of which it could be distinguished.

What can explain the Super Scoop’s change of status—from an item lying within the penumbra of the open-textured predicate “vessel,” for which either decision concerning applicability would have been permitted, to an item whose exclusion from the category of

vessels is now required? The current account offers an explanation, since the required classifications depend on the background set of precedent cases, and this set has changed, from Γ_4 to Γ_5 —it now contains *Di Giovanni*, which, in accord with the reason model, requires the judgment that the Super Scoop is not a vessel. It is more difficult to find an explanation for this change of classification on any view according to which the classification of an item—as a vessel, not a vessel, or lying in the penumbra—is supposed to depend on a separate core of settled meaning for the open-textured predicate.

One way to understand the difficulty is to ask: if there is a separate core of settled meaning for the predicate “vessel,” did this core of settled meaning change with the *Di Giovanni* decision? And here we face a dilemma. If the core of settled meaning did not change, and the core of settled meaning is what determines the classification of an item, then, since the Super Scoop lay within the penumbra prior to the *Di Giovanni* decision, it should remain in the penumbra afterward. On the other hand, if the core of settled meaning for the predicate “vessel” did change with the *Di Giovanni* decision, then that could explain the change of classification, of course. But in that case, if the core of settled meaning of an open-textured predicate can vary with the set of precedent decisions concerning applicability of that predicate, and variation in this set of precedent decisions can account for changes of classification all on its own, as in the reason model, then it is reasonable to wonder what additional work the separate core of settled meaning is supposed to be doing.

5 Discussion

The goal of this paper has been to suggest that the reason model of precedential constraint can be generalized to provide a semantic account of open-textured predicates, primarily in a legal setting, but applicable to other uses of open-textured predicates as well. This suggestion could be developed in a number of ways. As a first example, our treatment of open texture could be adapted to provide an account of the closely related phenomenon of vagueness. Here, the idea would be that vague predicates form a special class of open-textured predicates whose applicability is determined by factors keyed to values along dimensions with a particular, often numerical structure—for the vague predicate “tall,” say, the relevant dimension would be that of height, and the set of relevant factors might include being at least 6’2” in height, for instance, or no more than 5’11” in height.²³

As a second example, the current treatment of open texture relies on a set of factors whose own applicability, or not, is assumed to be clear. And often, this assumption is appropriate—it is at least relatively clear, for instance, whether a particular marine platform has a captain, or ballast tanks. In other cases, however, the factors in terms of which open-textured predicates are analyzed may themselves be open-textured. The question whether a marine platform should be classified as a vessel also depends, for instance, on whether it is largely stationary, or whether it has navigation as its primary purpose. But “largely stationary” and “has as its primary purpose” are straightforward examples of monadic and dyadic open-textured predicates. In order to apply to examples like this—with open-textured

²³See Horty (2019) for an investigation of the reason model of constraint based on dimensional factors of this kind, and then Cumming (20xx) for a treatment of vagueness based on these dimensional factors.

predicates analyzed in terms of other open-textured predicates—the current account would have to be extended to apply, not simply to a single open-textured predicate analyzed in terms of a set of concrete underlying factors, but to hierarchies of interlocking open-textured predicates.²⁴

Rather than exploring directions for further technical development, however, I want to close by addressing two more philosophical concerns.

First, I claim to have provided a semantic account of open-textured predicates—but can the account provided here really be characterized as semantic? The dominant approach to semantics in contemporary philosophy is truth-conditional, with the goal of specifying, in a systematic way, the conditions under which sentences are true—so, for example, the conditions under which a sentence like “The Super Scoop is a vessel” is true. The goal of the present account, by contrast, is not to specify the conditions under which sentences are true, but instead, the conditions under which the court is required or permitted to affirm certain statements, as well as the reasons on the basis of which it is permitted to justify its decisions—it tells us whether, for example, against the background of a set of previous decision, the court is required to affirm that the Super Scoop is a vessel, required to affirm that it is not a vessel, or both permitted to affirm that it is a vessel and also permitted to affirm that it is not a vessel. The present account is thus developed at an entirely different level from the standard truth-conditional approach to semantics—it is normative, working at the level of requirements and permissions, rather than factual.

My response to this first concern is that, while truth-conditional semantics may be the dominant semantic approach, there are any number of other approaches—such as proof-

²⁴See Canavotto and Horty (20xx) for a proposal.

theoretic, or verificationist, semantics, conceptual-role semantics, dynamic semantics, or inquisitive semantics, just to mention a few. Stepping back from this tangle of terminology, and the associated tangle of competing semantic theories, I adopt here a perspective most closely associated with the work of Michael Dummett, according to which a theory of meaning is, at bottom, a theory of understanding—the function of a semantic theory is to provide an account of what it is that a speaker knows by virtue of understanding a language.²⁵ From this perspective, the current account can be seen as offering the—arguably reasonable—suggestion that what a speaker knows who understands the meaning of an open-textured predicate is, not some set of necessary and sufficient conditions under which that predicate is applicable, but instead: taking into account the existing authoritative uses of that predicate, what reasons are permitted to justify further applications of the predicate or its contrary, and when the application of that predicate or its contrary is itself either permitted or required.

The second concern involves the generality of the proposed account. The example of the Super Scoop illustrates how, in the legal setting, constraints derived from previous decisions force coordination in the application of open-textured predicates. But I have also suggested that applications of open-textured predicates in natural language are constrained in a similar way: Over the course of a conversation, which can last for seconds or for centuries, a stock of prior applications of an open-textured predicates is established. Individuals who wish to participate in this conversation, rather than starting a new one, are then required to use these predicates in a way that respects the constraints established in their previous applications.

Can the account of open texture presented here really be generalized in this way, from

²⁵This perspective can be found throughout Dummett's work, receiving its earliest extensive discussion in Dummett (1975).

legal predicates to open-textured predicates more generally? I think this is a hypothesis worth considering, but of course, even as a hypothesis, it would have to be explored in much more detail—there are many ways in which the use of open-textured predicates in natural language differs from their use in the law. In the legal setting, for example, the set of precedent cases bearing on the further application of an open-textured predicate is carefully documented and curated; if questions arise, there are recognized methods of argument for determining whether or not some previous decision functions as an authoritative precedent in a new situation. In the more fluid setting of a natural language, by contrast, we could expect the set of precedent cases constraining the use of open-textured predicates to be indefinite, local, and changing; speakers might well exercise creativity by flouting norms, ignoring previous cases that should count as precedents, or granting authority to previous cases that should not. Another difference between open-textured predicates in legal language compared to natural language more generally is that, while the authority of past decisions over present cases in the law is carefully documented, the nature of the authority on the basis of which previous uses of open-textured predicates might constrain current uses in natural language more generally is much less clear. My suspicion is that these constraints result from an unnoticed, or at least underexplored, principle of conversational coordination in natural languages, which leads to coordination in the use of open-textured predicates.²⁶

²⁶For a central account of meaning coordination in natural language, see, of course, Lewis (1969), who himself appeals to precedent but only as a mechanism for aligning mutual expectations. Empirical research along these lines can be found in Clark and Marshall (1981) and Clark and Wilkes-Gibbs (1986), and then in Garrod and Anderson (1987), who introduce the term “entrainment” for the kind of meaning coordination under consideration; this term was later adopted in Clark (1991) and Ludlow (2014). The current suggestion is that a precedent is more than just a kind of signpost for use by speakers to coordinate expectations, but

If this suspicion, or something like it, is correct, then the legal doctrine of precedent, like so much else in the law, can be seen as a more stylized, self-conscious, and rigorous development of a mechanism that is already at work in our everyday interactions.

Acknowledgments

I am very grateful for help from Ilaria Canavotto, Sam Cumming, and an anonymous reviewer for this volume.

References

- Aleven, Vincent (1997). *Teaching Case-Based Argumentation Through a Model and Examples*. PhD thesis, Intelligent Systems Program, University of Pittsburgh.
- Aleven, Vincent and Ashley, Kevin (1997). Evaluating a learning environment for case-based argumentation skills. In *Proceedings of the Sixth International Conference on Artificial Intelligence and Law (ICAAIL-97)*, pages 170–179. The Association for Computing Machinery Press.
- Ashley, Kevin (1989). Toward a computational theory of arguing with precedents: accommodating multiple interpretations of cases. In *Proceedings of the Second International Conference on Artificial Intelligence and Law (ICAAIL-89)*, pages 93–110. The Association for Computing Machinery Press.

that, instead, precedential constraint has a normative force, explicated here by the reason model, in bringing about what Garrod, Anderson, Clark, and Ludlow refer to as meaning entrainment.

- Ashley, Kevin (1990). *Modeling Legal Argument: Reasoning with Cases and Hypotheticals*. The MIT Press.
- Baker, Gordon (1977). Defeasibility and meaning. In Hacker, Peter and Raz, Joseph, editors, *Law, Morality, and Society: Essays in Honour of H. L. A Hart*, pages 26–57. Oxford University Press.
- Bix, Brian (1991). H. L. A Hart and the open texture of language. *Law and Philosophy*, 10:51–72.
- Broughton, Gabriel (2019). Vertical precedents in formal models of precedential constraint. *Artificial Intelligence and Law*, 27:253–307.
- Canavotto, Ilaria and Horty, John (20xx). Reasoning with a hierarchy of open-textured predicates. Manuscript, currently under review.
- Clark, Herbert (1991). Words, the world, and their possibilities. In Pomerantz, James and Lockhead, Gregory, editors, *The Perception of Structure*, pages 263–277. American Psychological Association.
- Clark, Herbert and Marshall, Catherine (1981). Definite reference and mutual knowledge. In Joshi, Aravind, Sag, Ivan, and Webber, Bonnie, editors, *Elements of Discourse Understanding*. Cambridge University Press.
- Clark, Herbert and Wilkes-Gibbs, Deanna (1986). Referring as a collaborative process. *Cognition*, 22:1–39.

- Cohen, Adam (2009). The Lord Justice hath ruled: Pringles are potato chips. *The New York Times*, May 31, 2009.
- Cumming, Samuel (20xx). Semantic reasons. Forthcoming in *Nous*.
- Dummett, Michael (1975). What is a theory of meaning? In Guttenplan, Samuel, editor, *Mind and Language: Wolfson College Lectures 1974*, pages 97–138. Oxford University Press.
- Dworkin, Ronald (1967). The model of rules. *The University of Chicago Law Review*, 35:14–46. Reprinted in Dworkin (1977), pp. 14–45.
- Dworkin, Ronald (1977). *Taking Rights Seriously*. Harvard University Press.
- Elbourne, Paul (2011). *Meaning: A Slim Guide to Semantics*. Oxford University Press.
- Endicott, Timothy (2000). *Vagueness in Law*. Oxford University Press.
- Endicott, Timothy (2022). Law and language. In Zalta, Edward, editor, *The Stanford Encyclopedia of Philosophy (Spring 2022 Edition)*. Stanford University. Available at <http://plato.stanford.edu/archives/spr2022/entries/law-language/>.
- Fodor, Jerry, Garrett, Merrill, Walker, Edward, and Parkes, Cornelia (1980). Against definitions. *Cognition*, 8:263–367.
- Garrod, Simon and Anderson, Anthony (1987). Saying what you mean in dialogue: a study in conceptual and semantic co-ordination. *Cognition*, 27:181–218.
- Hart, H. L. A. (1958). Positivism and the separation of law and morals. *Harvard Law Review*, 71:593–629.

- Hart, H. L. A. (1961). *The Concept of Law*. Oxford University Press. Second edition published in 1994; pagination refers to the second edition.
- Horty, John (2011). Rules and reasons in the theory of precedent. *Legal Theory*, 17:1–33.
- Horty, John (2012). *Reasons as Defaults*. Oxford University Press.
- Horty, John (2015). Constraint and freedom in the common law. *Philosopher's Imprint*, 15(25).
- Horty, John (2019). Reasoning with dimensions and magnitudes. *Artificial Intelligence and Law*, 27:307–345.
- Horty, John (20xx). *The Logic of Precedent: Constraint and Freedom in Common Law Reasoning*. Forthcoming with Cambridge University Press.
- Horty, John and Bench-Capon, Trevor (2012). A factor-based definition of precedential constraint. *Artificial Intelligence and Law*, 20:181–214.
- Keil, Geert and Poscher, Ralf (2016). *Vagueness and Law: Philosophical and Legal Perspectives*. Oxford University Press.
- Lamond, Grant (2005). Do precedents create rules? *Legal Theory*, 11:1–26.
- Lewis, David (1969). *Convention: A Philosophical Study*. Harvard University Press.
- Ludlow, Peter (2014). *Living Words: Meaning Underdetermination and the Dynamic Lexicon*. Oxford University Press.

- Lyons, David (1999). Open texture and the possibility of legal interpretation. *Law and Philosophy*, 18:297–309.
- MacCormick, Neil (1991). On open texture in law. In Amslek, Paul and MacCormick, Neil, editors, *Controversies about Law's Ontology*, pages 72–83. Edinburgh University Press.
- Mullins, Robert (2020). Protected reasons and precedential constraint. *Legal Theory*, 26:40–61.
- Prakken, Henry (2021). A formal analysis of some factor- and precedent-based accounts of precedential constraint. *Artificial Intelligence and Law*, 29:559–581.
- Raz, Joseph (1972). Legal principles and the limits of law. *The Yale Law Journal*, 81:823–854.
- Rigoni, Adam (2014). Common-law judicial reasoning and analogy. *Legal theory*, 20:133–156.
- Rissland, Edwina (1990). Artificial intelligence and law: stepping stones to a model of legal reasoning. *The Yale Law Journal*, 99:1957–1981.
- Rissland, Edwina and Ashley, Kevin (1987). A case-based system for trade secrets law. In *Proceedings of the First International Conference on Artificial Intelligence and Law (ICAAIL-87)*, pages 60–66. The Association for Computing Machinery Press.
- Rissland, Edwina and Ashley, Kevin (2002). A note on dimensions and factors. *Artificial Intelligence and Law*, 10:65–77.
- Schauer, Frederick (2008). A critical guide to vehicles in the park. *New York University Law Review*, 83:1109–1134.

Schauer, Frederick (2013). On the open texture of law. *Grazer Philosophische Studien*, 87:197–215.

Schiappa, Edward (2003). *Defining Reality: Definitions and the Politics of Meaning*. Southern Illinois University Press.

Tur, Richard (2001). Defeasibilism. *Oxford Journal of Legal Studies*, 21:355–368.

Waismann, Frederick (1945). Verifiability. In *Proceedings of the Aristotelian Society, Supplementary Volume 19*, pages 119–150. Harrison and Son. Reprinted in Anthony Flew, editor, *Logic and Language*, pages 117-144, Blackwell Publishing Company, 1951; pagination refers to this version.