What the Epistemic Account of Vagueness Means for Legal Interpretation

[Cite published version: *Law and Philosophy*, Volume 35, Number 1 (February 2016): 29-54]

Luke William Hunt

To suggest that the law is vague is to suggest a platitude. Almost all of our language may be thought of as vague in some sense, and the language of law is no different. What distinguishes legal language from the rest of our language is that questions of life and death may turn upon the interpretation of a word, phrase, or expression. It is in part for this reason that most of the literature on vagueness in the law has focused upon the extent to which legislation could have been written more precisely or conversely the extent to which vagueness might be useful in law.[1] While there is good reason to make sense of these issues, less attention has been given to the question of whether a particular theory of legal interpretation is consistent with a particular account of vagueness. There are likewise good reasons to make sense of this question, not least of which is that a failure to do so lends itself to an *ad hoc* account of vagueness in the law. If we do not adhere to a particular account of vagueness, then our rationale for asserting the content of the law is on shaky ground.

In this paper, then, I will explore the latter question raised—what philosophical accounts of vagueness mean for theories of legal interpretation—in the context of the *epistemic account of vagueness ("epistemicism")*: the thesis that vague statements are true or false even though it is impossible to know which. Specifically, I will argue that if epistemicism is accepted within the domain of the law, then the following three conditions must be satisfied: (1) Interpretative reasoning within the law must adhere to the principle of bivalence and the law of excluded

[1] See Timothy Endicott, "The Value of Vagueness"; Scott Soames, "What Vagueness and Inconsistency Tell Us about Interpretation"; and Jeremy Waldron, "Vagueness and the Guidance of Action", in Andrei Marmor and Scott Soames (ed,), *Philosophical Foundations of Language in the Law*(Oxford: Oxford University Press, 2011).

middle, (2) interpretative reasoning within the law must construe vague statements as an epistemic phenomenon, and (3) epistemicism must be expanded to include normative considerations in order to account for legal theories that are consistent with the first two conditions. The first two conditions may be dealt with more easily because they are *internal* to a particular theory of legal interpretation: The features of a particular legal theory either do or do not map to epistemicism coherently. The third condition poses more difficulty because it is *external* to a particular theory of legal interpretation: Given a particular legal theory that is consistent with the first two conditions, epistemicism itself must be expanded to account for the normativity presupposed by such a legal theory. I will draw out these points by examining two prominent theories of legal interpretation: Ronald Dworkin's theory of *law as integrity,* which is based upon our knowledge of interpretative facts; and Michael Moore's theory that the law is a *functional kind,* which is based upon metaphysical truths about the way the world is.[2] In the recent literature on vagueness and the law, it has been suggested that both theories are consistent with epistemicism.[3] However, I will try to show that the former is inconsistent with epistemicism unless epistemicism is expanded as described in condition (3), and the latter is inconsistent with epistemicism because it does not satisfy condition (2). This conclusion is significant in several ways. It shows that there are legal theories that are internally consistent with the fundamental features of epistemicism. However, within the domain of law—and specifically in the case of legal theories that are internally consistent with epistemicism—it shows that vagueness cannot be explained simply by our ignorance of the meaning and use of

[2] Ronald Dworkin, *Law's Empire* (Cambridge: Harvard University Press, 1986); Michael Moore, "Law as a Functional Kind", in Robert P. George (ed.), Natural Law Theory (Oxford: Oxford University Press, 1992).
[3] See Ralf Poscher, "Ambiguity and Vagueness in Legal Interpretation", in Peter M. Tiersma and Lawrence M. Solan (ed.), *The Oxford Handbook of Language and Law* (Oxford: Oxford University Press, 2012), p.138 (in which Poscher suggests without argument that both Dworkin's theory and Moore's theory correspond to epistemicism).

vague expressions. Rather, epistemicism must also account for ignorance of the requisite normative considerations in legal theories with which it is otherwise consistent.

1. The Threat of Vagueness and the Epistemic Response

In this section, I will provide a sketch of how vagueness threatens two principles of classical logic: the *law of excluded middle* ("*P* or not *P*"), and the principle of *bivalence* (any sentence *P* that expresses a proposition is either true or false). The threat is not limited to classical logic. If vagueness threatens logical principles regarding the truth and falsity of a sentence's propositions—and the law consists of sentences—then vagueness perhaps threatens the principles regarding the truth and falsity of the law.[4] I will illustrate this threat by presenting vagueness in one of its most well-known forms: the *sorites paradox*. I will then examine how the threat is addressed in the *locus classicus* of epistemicism, Timothy Williamson's *Vagueness*,

[4] There is a divergence of opinion about this claim. Alethic logic is the branch of modal logic that is concerned with necessary truth and related modalities, while deontic logic is the branch of modal logic that is concerned with normative modalities like what is obligatory, required, and so on. These divisions are important because the truth conditions for a statement like, "the cat is on the mat," are perhaps of a different nature altogether than those for the directive, "provide reasonable aid to another in need." More specifically, there is a divergence of opinion regarding the extent to which standard notions of logical consequence and truth apply to normative sentences like the latter example. This is because applying alethic modalities to normative sentences does not always result in the consequences we have come to expect in classical logic. Indeed, such applications may very well result in a paradox, which has been discussed thoroughly in the deontic logic canon. See Jorgen Jorgensen,"Imperatives and Logic", *Erkenntnis* 7 (1937), pp. 288–296 (arguing that while imperatives may appear incapable of being true or false, they are capable of working as premises and consequences in logical inferences—which might be explained by the fact that imperative sentences contain both imperative and indicative factors); Alf Ross, "Imperatives and Logic", *Theoria* 7 (1941), pp. 53–71 (describing how if B follows from A, then OB (with "O" being read as an imperative operator) follows from OA cannot be valid. Otherwise, from, "You ought to mail the letter," we could conclude that, "You ought to mail the letter or burn it," meaning that an obligation to mail the letter implies an obligation that may be met by burning the letter.). While I cannot address the voluminous work done on the logic of imperatives since these early paradoxes were introduced, I do want to lay my cards on the table the best I can. The law deals with normative propositions and those propositions are quite often imperatives. Recent work has gone a long way toward establishing a limited logic for imperatives that might help provide a foundation for normative discourse. While the matter is by no means closed, it is likely that this line of thought is most promising with respect to my position that normative propositions can in some sense be said to be true in the context of legal reasoning as far as vagueness is concerned. See Peter Vranas, "New foundations for Imperative Logic I: Logical Connectives, Consistency, and Quantifiers", *Noûs* 42 (2008), pp. 529–72 (arguing for a logic that mixes imperatives and declaratives); Peter Vranas, "In Defense of Imperative Inference", *Journal of Philosophical Logic* 39 (2010): 59–71 (arguing that some inferences have imperatives as premises and conclusions and that issuing imperatives with conflicting permissive presuppositions does not entail changing one's mind).

as well as the way in which Williamson treats vagueness as an epistemic phenomenon.[5] It will be necessary to get clear on these points before moving to epistemicism and theories of legal interpretation in part 2.

1.1. The Sorites Paradox

Suppose a woman, Jane, is out on a jog at night and observes a small child struggling to keep his head above water in a large but shallow lake. Jane stops briefly when she hears the child, walks some yards into the shallow lake, but decides to turn around before reaching the child. The incident is caught on a surveillance camera and Jane is charged with failing to render aid to the child, who drowned. The law that Jane is alleged to have violated states: "A person who knows that another is exposed to grave physical harm shall, to the extent that the same can be rendered without danger or peril to himself or without interference with important duties owed to others, give reasonable assistance to the exposed person unless that assistance or care is being provided by others."[6] Because of the angle of the camera and the poor light, it is impossible to determine the child's exact location in the lake when Jane witnessed him struggling (though audio from the surveillance video confirms that the child can be heard by Jane). The case turns upon the interpretation of "reasonable assistance." Jane, a philosophy major, argues that offering assistance to the child would be unreasonable because the child was too far out in the lake, even though the lake is shallow. She states the following upon taking the witness stand: "Of course, taking some steps into a shallow lake to save a child is clearly reasonable, but walking several yards into a lake when the child is very far away is clearly unreasonable. Any contrary argument would be absurd because its conclusion wouldn't follow from its premise." Somewhat perplexed, the prosecutor asks Jane what she means. Jane replies:

[5] See Timothy Williamson, Vagueness (New York: Routledge, 1994).
[6] 12 Vt. Stat Ann.s. 519.

> We all agree that taking one step into the lake would not be unreasonable assistance, and, moreover, that if it is not unreasonable to impose a duty to walk one step to save a drowning child, then surely it is not unreasonable to require two steps. Simply put, the difference between two steps and one step is morally trivial. But if we continued along this line of argument, we would reach the conclusion that the difference between two steps and three is equally trivial, as is the difference between threes steps and four steps—or between sixteen and seventeen, or between 499 and 500. And, therefore, if it is not unreasonable to take one step into the lake to rescue a child, then it is not unreasonable to take 500 steps into a lake to save a child. But this conclusion is clearly absurd.

Triumphantly, Jane then declares that she is not liable. "Not so fast," the prosecutor replies. "You haven't made the case that you are not liable under the law for failing to save the child," he continues. "You have shown only that the conclusion of a particular kind of argument is paradoxical, which has no bearing on you liability under the law." At a loss, the judge provides the following written instruction to the jurors to guide their deliberation: "The arguments you have heard during this trial perhaps show that there simply cannot be one step into a pond at which point Jane's duty shifted from being reasonable to unreasonable. Rather, the arguments seem to suggest that it is indeterminate whether Jane is legally liable on the basis of what the law says, and, as finders of fact, you must resolve that indeterminacy the best you can."

What is the best way to characterize the sorites paradox that is raised in Jane's case?[7] Are such arguments valid, absurd, indeterminate, or something else altogether? The sorites paradox is a paradox about vague predicates, in this case the predicate "reasonable assistance." Before addressing responses to the problem, I will try to clarify what is at stake. Here is the form of the argument:

[7] I use the example of a so called "Bad Samaritan Law" because it illustrates the sorites paradox vividly. However, one might argue that this example is a bad one, not least because such laws are relatively novel and rarely enforced. And perhaps it's unlikely that sorites reasoning would be used in this sort of case. But these objections do not dissolve the theoretical problem because vagueness in the law is pervasive and much of the law is sorties susceptible. To be sure, the sorites paradox is less like a novel philosophical problem in the law, and more like an undiscovered continent of jurisprudential issues.

- Taking one step into a shallow lake to save a child is not unreasonable.
- If taking one step into a lake to save a child is not unreasonable, then taking two steps is not.
- If taking two steps into a shallow lake to save a child is not unreasonable, then taking three steps is not.
- …
- If taking 499 steps into a shallow lake to save a child is not unreasonable, then taking 500 steps is not.
- _____
- Taking 500 steps into a shallow lake to save a child is not unreasonable.

The argument seems to be valid with respect to the rules of classical logic. In other words, the conclusion (taking 500 steps into a shallow lake to save a child is not unreasonable) must be true if all the premises are true. Moreover, each premise relies upon inference rules of classical logic: *modus ponens* and *Cut*. *Modus ponens* allows us to infer "two steps into a shallow lake is not unreasonable" when it is true that "one step into a shallow lake is not unreasonable" and it is true that "if one step into a shallow lake is not unreasonable, then two steps is not unreasonable." *Cut* is the principle that valid arguments (e.g., multiple applications of *modus ponens*) may be chained together for an overarching application of *modus ponens*, meaning that 499 applications of *modus ponens* results in the conclusion that 500 steps is unreasonable. Hence, if one step is not unreasonable, then 500 steps is not unreasonable. This conclusion seems clearly false because there would seem to be a point at which the number of steps one must take into a lake at night to save another ceases to be a reasonable legal duty. In other words, we want to say that the first premise (taking one step into a shallow lake to save a child is not unreasonable) is true and its conclusion (taking 500 steps into a shallow lake to save a child is not unreasonable) is false, meaning that one of the other conditional premises must be false. But how are we to argue that one of the premises is false—and identify such a premise— if each is based on the above principles of classical logic?

One response to this dilemma is that proposed by Joel Feinberg, who describes a legal duty to rescue scenario similar to Jane's. If there is a borderline zone—say, twenty-five to fifty yards—at which it is not clearly reasonable and not clearly unreasonable to impose a legal duty to save another in a lake, then perhaps the law should impose liability with such a borderline zone in mind. For example, Feinberg suggests that the we divide these sorts of cases into three zones of rescue: "(1) clear cases of opportunity to rescue with no unreasonable risk, cost, or inconvenience whatever; (2) cases of opportunity to rescue but only at clearly unreasonable risk, cost, or inconvenience to the rescuer or others; and (3) everything in the vast no-man's-land of uncertain cases in between the extremes."[8] In order to play it safe, Feinberg suggests that we hold no one liable in the middle, borderline zone, and instead draw the line of demarcation where the first zone (clearly reasonable) ends and the second zone (borderline zone in which it is not clearly reasonable and not clearly unreasonable) begins. The idea is to hold one liable only if one *clearly* deserves to be liable. But Feinberg's solution immediately introduces the problem of *higher-order vagueness*: At what point does it cease to be clearly reasonable to rescue and become unclear as to whether it is clearly reasonable to rescue? In other words, there might be points at which it is unclear whether it is unclear whether it would be reasonable to rescue, leading to an infinite regress.[9] I will not recount in detail the many non-epistemic answers to this problem, but I think it is fair to say that most all non-epistemic answers share the following upshot in one way or another: (1) Either classical logic is inadequate with respect to vagueness and we must deny the law of excluded middle[10] and bivalence,[11] or (2) our language is

[8] Joel Feinberg, *Harm to Others* (New York: Oxford University Press, 1984), pp. 155-56.
[9] Feinberg acknowledges that drawing a line between the two zones may seem difficult and arbitrary, but that "careful draftsmanship of statues could leave it up to juries to decide where reasonable doubts begin." Ibid. p. 156. While this is a perfectly good way to address the problem pragmatically, it of course does not address the underlying phenomenon of higher order vagueness and how one might account for that phenomenon in the law philosophically.
[10] Suppose Jane walked into Feinberg's borderline zone, which constitutes assistance that is not clearly reasonable and not clearly unreasonable. Perhaps the vagueness involved in the predicate "reasonable assistance" invalidates

inadequate with respect to vagueness in that vague sentences do not say anything.[12] If neither of these options is appealing, then one might be led to epistemicism.

---

classical logic such that "Jane offered reasonable assistance or Jane did not offer reasonable assistance" (an instance of the law of excluded middle) is not true. This view might be based on the position that if the disjunction, "Jane offered reasonable assistance or Jane did not offer reasonable assistance" is true, then one of the disjuncts is true. However, in the borderline zone, neither "Jane offered reasonable assistance" nor "Jane did not offer reasonable assistance" is true—thus invalidating the law of excluded middle. This is a rough summary of Russell's account of vagueness; see Williamson, *Vagueness*, Chapter 2, (New York: Routledge, 1994) for a more complete description.

[11] In Jane's case, the principle of *bivalence* holds that "Jane offered reasonable assistance" is either true or false. But what is the truth value of such a sentence if Jane's actions fall within the borderline zone? Non-epistemic responses to the problem of vagueness and bivalence take one of two approaches generally. For instance, one might argue that a third category is needed in addition to the classifications "true" and "false." Accordingly, "Jane offered reasonable assistance" is eligible to fall into the third category of "neither true nor false" ("neutral"). This raises two concerns. First, two-valued (true and false) truth tables of classical logic are no longer adequate because they neither account for the third category of "neutral," nor do they account for "the meaning of basic logical operators and the mechanical test of validity for inference forms involving them," as Williamson puts it. Williamson, *Vagueness*, 99. Second, by adding the third category, neutral, we are still faced with the problem of higher order vagueness described in Feinberg's three categories. In other words, we still cannot establish a precise point at which a proposition changes from true to neutral and a precise point at which a proposition changes from neutral to false. More sophisticated solutions to the problem include infinite-valued logic and supervaluations. See ibid., Chapters 4 and 5. Each of these solutions denies the discussed principles of classical logic to some extent, though I will not address arguments regarding the extent to which such denials are justified.

An alternative, semantic approach to the problem of vagueness and bivalence is as follows. If Jane's actions fall within a borderline zone, then we might be tempted to say that there is no proposition to *know* in the sentence "Jane did not offer reasonable assistance" because of the sentence's inherent vagueness. Under this rationale, one must adopt the position that nothing has been said to be the case in the sentence: "Jane did not offer reasonable assistance." But it seems odd to suggest that the prosecutor's utterance, "Jane did not offer reasonable assistance," has no propositional content and is neither true nor false. It likewise seems odd to suggest that there is no propositional content or truth value in the sentence that Jane utters: "I offered reasonable assistance." On the other hand, objections of this sort may have limited appeal with respect to the law. For example, one might argue that the law is indeterminate in this example, inasmuch as the predicate "reasonable assistance" proposes—but does not establish—a standard for legal adjudication. In other words, the predicate might simply be in need of interpretation that will yield its semantic content and allow one to say whether it is true or false, which will be based upon some particular theory of construction (i.e., perhaps the legislature did not provide sufficient information in the statutory law, requiring the judiciary to turn to additional (legal or extra-legal) considerations in order to apply the law). While this suggestion will not eliminate the problem of vagueness, it might go a long way toward solving the practical problem of deciding Jane's case in a legally satisfactory manner. And it might allow one to embrace indeterminacy without denying bivalence, inasmuch as one takes the position that it is simply unclear whether the statement, "Jane did not offer reasonable assistance," is true. In any event, for a thorough examination of what the *semantic account* of vagueness means for legal interpretation, see Timothy Endicott's seminal work, *Vagueness in Law* (Oxford: Oxford University Press, 1990).

[12] For example, it might be argued that a vague predicate like "reasonable assistance" is simply an incompletely defined predicate in the sense that it does not refer to anything. As such, any sentence containing the vague predicate "reasonable assistance" is neither true nor false. This means that the law of excluded middle would not be valid because "Either Jane provided reasonable assistance or Jane did not provide reasonable assistance" is not true if it does not refer to anything. This is a rough summary of Frege's account of vagueness; see Williamson, *Vagueness*, Chapter 2, for a more complete description.

[12] Williamson, *Vagueness*, p. 99.

1.2. An Answer to the Sorites Paradox

One of the fundamental aspects of Williamson's epistemicism is that it is disquotational, meaning that the sentence "Jane did not offer reasonable assistance" is true if and only if Jane did not offer reasonable assistance.[13] Another fundamental aspect of Williamson's epistemicism is that the denial of bivalence with respect to vague utterances leads to a contradiction. In the case involving Jane, this would involve rejecting the supposition that the utterance "Jane did not offer reasonable assistance" is either true or false. Using Jane's case, here is how Williamson's argument looks:

(1) *Denial of bivalence*: not ("'Jane did not offer reasonable assistance' is true" or "'Jane did not offer reasonable assistance' is false");

(2) *Disquotational scheme for truth*: "Jane did not offer reasonable assistance" is true if and only if Jane did not offer reasonable assistance;

(3) *Substitution in (1)*: not (Jane did not offer reasonable assistance or Jane offered reasonable assistance);

(4) *De Morgan's Laws*: not (Jane did not offer reasonable assistance) and not (Jane offered reasonable assistance);

(5) *Elimination of negation*: Jane did not offer reasonable assistance and Jane offered reasonable assistance.

And so we are left with a contradiction in (5) above.[14] If this is an unacceptable conclusion that lends support to epistemicism, then what is the upshot? First, borderline cases involving utterances such as "Jane did not offer reasonable assistance" state meaningful propositions. In

[13] See ibid., pp. 190-92.
[14] See ibid., pp. 187-89, for Williamson's discussion of the argument, of which he concludes: "The supposition of a counterexample to bivalence has led straight to a contradiction, through the elucidations of truth and falsity and some trivial logic."

other words, there is a fact of the matter to know, namely, whether it is true or false that "Jane did not offer reasonable assistance." Second, if there is a fact of the matter in these sorts of borderline cases, then there must be a sharp cut-off point at which "Jane did not offer reasonable assistance" moves from being true to false (or else we are left with the contradiction in (5) above). Third, although there are sharp cut-off points in borderline cases, we are ignorant of those points. The third point can be illustrated with Williamson's "omniscient speakers" ("OS") thought experiment.[15]

Williamson assumes that ignorance is involved in all borderline cases, regardless of which account of vagueness one endorses. The OS thought experiment draws out this point by illustrating how borderline cases still exist even if there is nothing hidden because one is omniscient. Here is how the thought experiment works. Suppose there are OS(s) (ignorant of nothing regarding borderline cases) such that if Jane did not offer reasonable assistance, then the OS knows that Jane did not offer reasonable assistance; and if Jane did offer reasonable assistance, then the OS knows that Jane did offer reasonable assistance. You ask the OS a series of questions beginning with "Has Jane offered reasonable assistance by taking one step?" and then adding an additional step for each subsequent question to the OS (two steps, threes steps, and so on). The OS will answer "no" after the first few questions, but at a certain point (unknown to you) the OS will stop saying "no." For an OS, then, there would be a certain point at which the number of steps Jane took ceases to be unreasonable, but that point "marks some sort of previously hidden boundary" that you are in no position to know.[16] In other words, the boundary that is known by the OS is hidden from ordinary (non-omniscient) speakers because ordinary speakers are ignorant of the "semantic significance" of borderline cases involving a

[15] See ibid., Section 7.3, for this thought experiment, which I will discuss in the context of Jane's case.
[16] Ibid. p. 200.

predicate like "reasonable assistance."[17] For Williamson, the OS thought experiment suggests that vagueness is intuitive: If an ordinary speaker knew the point at which an OS would cease to answer "no" to the question about Jane, then there would no longer be a phenomenon of vagueness but rather a denial of the phenomenon of vagueness. Perhaps the best way to clarify Williamson's point here is to return to classical logic: By the law of excluded middle, Jane either offered reasonable assistance or Jane did not offer reasonable assistance, and based on bivalence the utterance "Jane offered reasonable assistance" is either true or false. As exemplified by the OS thought experiment: "[These laws] hold[] in borderline cases. It is just that we are in no position to find out which truth-value the vague utterance has."[18] This raises the final piece of Williamson's account of vagueness, namely, the nature of our ignorance.

The ignorance we have in borderline cases is part of a wider phenomenon to which Williamson refers as *inexact knowledge*.[19] I will briefly touch upon but one aspect of inexact knowledge that is relevant to my argument: meaning and natural divisions. Here is an example. If the prosecutor in Jane's case states, "Anyone who takes only 150 steps into a shallow lake has not offered reasonable assistance to rescue another," then he does not know the truth of his utterance if we assume that 150 steps is a borderline case of reasonable assistance. To put the point differently, the extension of a predicate like "reasonable assistance" could have been slightly different such that it would have included everyone who takes 150 steps into a shallow lake. Williamson explains it this way: "What distinguishes vagueness as a source of inexactness is that the margin for error principles to which it gives rise advert to small differences in meaning, not small differences in the objects under discussion."[20] The point is more obvious

[17] Ibid. p. 201.
[18] Ibid.
[19] Ibid. Chapter 8.
[20] Ibid. pp. 230-31.

when we contrast extensions of predicates like "reasonable assistance" with words that are "stabilized by natural divisions, so that a small difference in use would make no difference in meaning."[21] Williamson is describing natural-kind terms like gold. For instance, if we begin to more commonly mistake fool's gold with gold—the substance with atomic number 79—then the meaning of gold would not change. In other words, natural-kind terms like "gold" are not vague under Williamson's account—even if we are unsure whether something is gold or fool's gold—because they have fixed, natural boundaries that are not susceptible to deficiencies in our knowledge. Of course, many of our terms (like reasonable assistance) are not natural-kind terms, and it is these sorts of terms that illustrate genuine instances of vagueness for Williamson. But we shouldn't lose sight of the big picture for Williamson: There is a fact of the matter with respect to the proposition in the utterance, "Jane did not offer reasonable assistance," though we have no way of knowing the cut-off point in borderline cases involving such utterances (as illustrated in by the OS); still, we have reason not to be completely skeptical about our knowledge—inexact as it is—because we know enough about the meaning of "reasonable assistance" to know that taking two steps into a shallow lake is not in its extension, for example.[22]

---

[21] Ibid. p. 231.

[22] Based on what has been said so far, one might think that epistemicism is too rigid for any theory of legal interpretation. After all, the account suggests that there is a sharp cut-off with respect to all of our vague predicates, though we are ignorant of those cut-offs in borderline cases. It might seem incoherent—and impractical—to suggest that there is a single step at which Jane's offer of assistance changes from being unreasonable to reasonable. This is a fair concern, but the concern is perhaps assuaged to the extent that salvaging the principles of classical logic is important to one's theory of legal interpretation. The concern may also be assuaged somewhat by highlighting the compatibility of context with epistemicism. See Delia Graff, "Shifting Sands: An Interest-Relative Theory of Vagueness," *Philosophical Topics* 28.1 (2000), in which Graff proposes for classes of constraints with respect to our use of vague predicates. The constraints describe how vague predicates are relational properties between both the standards used for the predicate and the interests of the speaker using the predicate. A theory of legal interpretation might thus embrace the idea that vague expressions may involve different standards in the context of different cases, yet still correspond to epistemicism. While this point may not be profoundly revelatory, it does seem to temper the seeming austerity of epistemicism with respect to how it might relate to theories of legal interpretation.

2. Epistemicism within the Domain of Law

So far I have tried to highlight two fundamental features of epistemicism: (1) adherence to the principle of bivalence and the law of excluded middle, and (2) construal of vague statements as an epistemic phenomenon. In this section, I will first argue that any theory of legal interpretation that corresponds with epistemicism must be internally consistent with these features. I do this by examining two prominent theories of legal interpretation: Dworkin's *law as integrity* and *right answer thesis,* and Moore's theory that the law is a *functional kind*. I conclude by arguing that—even if a legal theory is internally consistent with these features—vagueness cannot be explained simply by our ignorance of the meaning and use of vague expressions within the domain of the law. Instead, epistemicism must be expanded to account for ignorance of normative considerations in legal theories with which it is otherwise consistent.

2.1. Bivalence and Excluded Middle must hold within the Law

Any analysis of the role of vagueness in Dworkin's *law as integrity* and *right answer thesis* would do well to begin with a disclaimer, namely, that these aspects of Dworkin's work are perhaps not concerned with vagueness at all. This might sound odd at first blush, particularly because Dworkin spent a fair amount of time talking about what many would consider paradigm cases of vague predicates (e.g., "cruel" or "fair"). But rather than think of these predicates as being vague, Dworkin viewed them as predicates expressing legal concepts that are simply contested.[23] Indeed, he went as far as likening vague predicates within the law as mere "semantic defects."[24] Timothy Endicott probably characterized this aspect of Dworkin's theory best when he suggested that "the central elements of his theory leave the vagueness of language

[23] See Ronald Dworkin, *Taking Rights Seriously* (Cambridge: Harvard University Press, 1978), p. 128.
[24] Ronald Dworkin, *Law's Empire* (Cambridge: Harvard University Press, 1986), p.17.

by the wayside."[25] I will say more about this in section 2.3—particularly the extent to which Dworkin's interpretative process justifies a single right answer in each hard case based upon, in part, moral and political virtues—but the important point for now is that the issue of vague legal predicates may be conceptually distinguished from the issue of contestable legal concepts. With this distinction in mind, I now turn to some heretofore unconsidered internal consistencies that Dworkin's theory shares with epistemicism—specifically, those that become manifest from an analogy between Dworkin's Hercules and Williamson's Omniscient Speaker.

Dworkin's theory of *law as integrity* may be stated succinctly: "According to law as integrity, propositions of law are true if they figure in or follow from the principles of justice, fairness, and procedural due process that provide the best constructive interpretation of the community's legal practice."[26] This means that the truth value of legal propositions is based upon a vast array of inextricably intertwined considerations, including both descriptive statements about what the law is and evaluative statement about what the law should be. More specifically, the process of determining the truth value of a legal proposition (e.g., "According to the law in Jane's case, Jane failed (or did not fail) to provide reasonable assistance.") requires an interpretive reasoning based on both facts and moral principles. A variety of legal theories in the natural law tradition are perhaps consistent with Dworkin's position that interpretative reasoning has this dual nature. However, Dworkin departs radically from most legal theories—natural law traditions or otherwise—in a way that makes his theory correspond uniquely to epistemicism: There is a unique, right answer to even the hardest of cases.

A dispositive concept in the law is one that imposes a duty upon a judge—if the concept is satisfied—to decide a particular legal case in a particular way; if the concept is not satisfied,

---

[25] Endicott, *Vagueness in Law* (Oxford: Oxford University Press, 1990), p. 70.
[26] Dworkin, *Law's Empire*, 225.

then the judge has a duty to decide the case in the opposite way (e.g., the concept of reasonable assistance). Dworkin raises what he calls the "bivalence thesis" about these concepts: "In every case *either* the positive claim, that the case falls under a dispositive concept, or the opposite claim, that it does not, must be true even when it is controversial which is true."[27] Two versions of a popular competing thesis ("no-right-answer" theses) take a contrary view, namely, that the law has no right answer in some cases. The first version of the no-right-answer thesis argues that the bivalence thesis does not exhaust logical space in a case because, for example, it might be false both that Jane is liable and that Jane is not liable for the degree of assistance she offered the drowning boy. Instead, there might be some third category such that a judge would not have a duty to find Jane liable, nor a duty to find her not liable—"vulnerability to liability," as Dworkin dismissively calls it. The second version of the no-right-answer thesis does not argue for a third category, but rather that it may not be true that either Jane provided reasonable assistance or that Jane did not provide reasonable assistance because Jane's act falls within a borderline case of reasonable and unreasonable assistance; in other words, it would be a mistake to say that either Jane did or did not provide reasonable assistance when her action was on the border of the two.[28]

Although Dworkin does not put it in these terms, both versions of the no-right-answer thesis fall prey to the problems of higher order vagueness to the extent that they deny bivalence and the law of excluded middle.[29] For example, at what point does Jane cease to be clearly liable for her action and become "vulnerable to liability," but not yet liable. As discussed in Part 1, there might be points at which it is unclear whether it is unclear whether Jane should be liable.

---

[27] Ronald Dworkin, *A Matter of Principle* (Cambridge: Harvard University Press, 1985), pp. 119-20.
[28] Ibid. p. 121.
[29] Again, note that one who embraces a semantic account of vagueness might argue that the fact that the semantic content of a legal sentence is indeterminate does not imply a denial of bivalence. Rather, one might argue that a semantic account of vagueness claims merely that it is unclear whether a statement such as "Jane did not offer reasonable assistance" is true. See supra note 11.

And if it is neither true nor false that Jane is liable (but perhaps vulnerable to liability), then two-valued truth tables of classical logic are no longer sufficient, even though there seems clearly to be propositions to know in the sentences "Jane is liable" and "Jane is not liable." The no-right-answer theses leave us where we began: Deny that classical logic can deal with borderline cases adequately or treat vague sentences as if they say nothing.

In addition to these problems, Dworkin rightly points out that the first version of the no-right-answer thesis fails on functional grounds. To be sure, the no-right-answer denies the key function of dispositive concepts, namely, providing a "bridge between certain sorts of events and the conclusory claims about rights and duties that hold if these events can be demonstrated to have occurred."[30] To put it another way, the point of dispositive concepts is to deny logical space between the conclusion that the concept holds or does not hold. With respect to the second version of the no-right-answer thesis (vagueness in legal language produces indeterminate propositions in the law), Dworkin foreshadows his theory's progression with a question that takes the following form: Provided that a legislature enacted a statute that assigns liability to those who fail to offer "reasonable assistance," provided whatever one might suppose about the state of mind of the legislators who wrote and enacted the law, provided whatever one might suppose about the attitude of the public toward imposing duties to rescue others, and provided whatever else might be relevant, *is Jane liable if the assistance she offered falls within a borderline zone of reasonableness, or is she not liable*? Dworkin's point is that the vagueness of terms like "reasonable assistance" and the vagueness in the other factors (state of mind of legislators, attitudes of pubic, etc.) are simply facts that must be taken into account—imperfectly due to incomplete knowledge, to be sure—in answering the question of Jane's liability. But the inevitable vagueness in such cases does not mean that there is no right answer to the question of

[30] Ibid. p. 125.

Jane's liability.[31] In short, Dworkin denies that there could be different answers in a vague case, each of which provides an equally good fit with respect to both the legal and moral facts involved in the interpretative reasoning of the judge. Still, a question remains: How could a judge know the one-right-answer to a vague case given that we have incomplete knowledge of legal and moral facts?

The answer is that a judge cannot know the unique, right way to decide every case. Therefore, if there is one right answer in every case and if it is impossible to know that one right answer in every case, then Dworkin's theory of legal interpretation is epistemic in nature. We can see this from the well-known hypothetical device Dworkin uses: *Hercules*, "the imaginary judge of superhuman intellectual power and patience who accepts law as integrity."[32] No mortal judge could fulfil the demands of law as integrity completely, for it requires the judge to grasp fully every interconnection and purpose of the law (including those that are descriptive with respect to current and prior legal documents, as well as those that are evaluative with respect to justice and fairness) and decide each case in a unique way that puts the law in its best light based upon these interconnections and purposes. For this reason, then, Dworkin asks us to imagine a Herculean judge—one whose full interpretation of the law no actual judge could approach, but one who an actual judge could imitate in a limited way.[33] Hercules knows that the law is built upon a coherent set of principles about justice and fairness and that law as integrity requires him

---

[31] Ibid. pp. 130-31. Dworkin argues that we are not limited to "hard facts" ("physical facts about behavior…of people") in deriving the answer in vague cases. He asks us to suppose that there are other things in the world—namely, moral facts—in virtue of which propositions of law may be true. By "moral facts" Dworkin means only that "a particular social institution like slavery might be unjust, not because people think it unjust, but just because slavery is unjust." Therefore, "a proposition of law might rationally be supposed to be true even if lawyers continue to disagree about the proposition after all the hard facts are known or stipulated. It might be true in virtue of a moral fact which is not known or stipulated." Ibid. 138. This claim is not defended explicitly in *A Matter of Principle*, though Dworkin does argue that some facts beside hard facts exist.

[32] Dworkin, *Law's Empire*, p. 239.

[33] Ibid. p. 245.

to satisfy these principles in each particular case.[34] Dworkin describes this in terms of reaching the proper interpretive *fit*, which includes those hard cases that are vague or that seem to lend themselves to multiple interpretations. The interpretive reasoning in such cases involves determining which answer "shows the community's structure of institutions and decisions—its public standards as a whole—in a better light from the standpoint of political morality." This is heady stuff, but Dworkin recognizes the difficulties that Hercules will face: The complexity of moral and political principles means that those principles will occasionally oppose each other; in reaching the one right answer in vague cases, Hercules must thus address higher-order convictions about how competing principles should be balanced when they compete.[35]

Hercules, then, illustrates that the impossible task of reaching a unique right answer in every vague case is not because a unique right answer does not exist, but rather because we are ignorant in a way that Hercules is not. Hercules is capable of examining every facet of a case, moral and non-moral, making his aim comprehensive. Actual judges are of course incapable of examining a case in such a comprehensive manner and are limited to partial aims. However, the answers reached by Hercules are not based upon "transcendental mysteries opaque to [actual judges]," but rather a complete grasp of the same material available to actual judges.[36] This allows us to draw three conclusions about Dworkin's argument. First, borderline cases involving utterances such as "Jane did not offer reasonable assistance and is thus liable" state meaningful propositions. For Dworkin, there is indeed a fact of the matter to be known. Second, if there is a fact of the matter in these sorts of borderline cases for Dworkin, then there must be a sharp cut-off point at which "Jane did not offer reasonable assistance." Third, it is impossible for actual judges to locate such cut-off points because actual judges are ignorant of them in borderline

[34] Ibid. p. 243.
[35] Ibid. p. 256.
[36] Ibid. p. 265.

cases, though they must model Hercules as closely as possible in their search for answers. The third point bears a striking resemblance to Williamson's "omniscient speakers" ("OS") thought experiment. Dworkin, like Williamson, assumes that ignorance is involved in all hard cases that come before actual judges. And the Hercules hypothetical, like the OS thought experiment, draws out this point by illustrating how borderline cases still exist even if there is nothing hidden because one is a superhuman judge. For both Hercules and an OS there would be a certain point at which the number of steps Jane took ceases to be unreasonable, but that point is a hidden boundary that actual judges are in no position to know. Of course, the key difference is that the account of vagueness in *law as integrity* is based upon more than our ignorance of meaning and use in borderline cases; it is also based upon ignorance of the relevant normative considerations (including balancing competing normative considerations, for instance). I will discuss this crucial difference and the upshot for epistemicism in part 2.3. First, however, I need to complete my account of how a legal theory is and isn't internally consistent with epistemicism. This task brings me to Moore's theory.

2.2. Vague Statements must be construed as an Epistemic Phenomenon within the Law

While we saw that Dworkin's theory possessed some elements of the natural law tradition (i.e., some notion that legal propositions depend in part on corresponding moral propositions), law as integrity is not clear about its metaphysical commitments.[37] The emphasis is rather epistemological, particularly regarding the extent that our knowledge is limited in comparison to an ideal judge. When we turn to Moore's theory of law as a functional kind, however, we see a theory that is centered squarely in the natural law tradition. Moore is thus

[37] Moore accuses Dworkin of refusing to make any serious metaphysical commitments, which Moore calls "the ostrich position because it refuses to look for the reference and extension of legal terms and because it refuses to look for any relation between legal and other, non-legal things." Moore adds, "In ethics and in law the dubious honor of occupying the ostrich position belongs most notably to Ronald Dworkin." Michael S. Moore, "Legal Reality: A Naturalist Approach to Legal Ontology", *Law and Philosophy* 21(6) (2002): p. 633.

obliged to provide a more explicit account of his metaphysical commitments regarding legal ontology, and it is his account of these commitments that drives a wedge between his theory and epistemicism.

First, consider Moore's position on vagueness and legal interpretation generally. In "The Semantics of Judging," Moore sets out to demonstrate the inadequacies of legal formalism, which he describes as a theory of adjudication asserting "that legal disputes can be, should be, and are resolved by recourse to legal rules and principles, and the facts of each particular dispute…decision[s] [are] to be logically deduced from these two items alone."[38] For Moore, the general problem with formalism is its inability to show how the meanings of words entail "some correct classification of particular items as either meant by those words or not."[39] Put another way, he argues that because language requires the characteristic *generality* (i.e., general predicates and classifying words are required of language, not just proper names), it follows that language includes certain other characteristics that make it impossible for formalism to say how *all* cases can be resolved—and perhaps impossible to say how *any* cases can be resolved.[40] One of these characteristics is vagueness, particularly what he calls *degree-vagueness*. With respect to this sort of vagueness, predicates might apply to a class of things that can be ordered along a series without natural boundaries, such as "is a heap" and other examples that yield a sorites paradox. Vagueness is thus a problem for formalism because legal rules governing, say, heaps, would result in cases in which deductions would be impossible inasmuch as the criteria for the way "heap" is normally used is itself vague.[41]

---

[38] Michael Moore, "The Semantics of Judging," *Southern California Law Review* 54 (1981): pp. 155-56. Moore describes Dworkin as taking the "extreme formalist position that there is a uniquely correct result in *every case*…" Ibid. p. 166.
[39] Ibid. p. 180.
[40] Ibid. p. 181.
[41] Ibid. p, 195. Importantly, Moore suggests that "natural kind words" present problems for formalism because: "Determining the extension of a natural kind word is thus not a matter of facing an item with a checklist of the

Building upon this critique of formalism, Moore stakes out an early formulation of his own view in "A Natural Law Theory of Interpretation."[42] He calls it a "natural law theory of legal interpretation" because it is characterized by the following two propositions: "(1) that there is a right answer to moral questions, a moral reality if you like; and (2) that the interpretive premises necessary to decide any case can and should be derived in part by recourse to the dictates of that moral reality."[43] Before attempting to develop this theory, he addresses certain "skepticisms" about the "moral case for relying on meanings in interpretation."[44] One of those skepticisms is of course vagueness. Moore's central response to the problem of vagueness in law is that it is truly a problem only if one endorses a conventionalist—rather than a realist—theory of meaning. In other words, conventionalist theories might suggest that a word is vague when we exhaust our conventions regarding correct usage of the word (or when our convictions conflict), with our only recourse being to resolve vague cases "as a matter of policy." However, Moore suggests that if we adopt a realist theory of meaning, then our conventions with respect to vague words are only placeholders for "the best theory we can muster."[45] While we may not know the best theory about a particular word, the realist is able to keep theorizing about what *the meaning of a word* (e.g., "death") *really is* based upon our best theories about what, say, death, *really is*. Our knowledge of the meaning may be (presently) deficient, but the word itself is not even vague if we view it as a natural kind, whose true nature we must uncover as our knowledge improves.[46] Under this rationale, the judge's task is not to change conventional meanings of

---

criteria for that word and checking off its properties to determine whether it is a piece of gold. Rather, it is assumed that gold, water, lemons and other natural kind words have a hidden nature, a basic and fundamental similarity to standard examples. It is this hidden nature that truly determines whether the particular item is to be classified as 'gold.'" Ibid. 205. I discuss how this point relates to epistemicism in the pages that follow.

[42] Michael Moore, "A Natural Law Theory of Interpretation", *Southern California Law Review* 58 (1985)p. 286.

[43] Ibid.

[44] Ibid. p. 302.

[45] Ibid. p. 308.

[46] Ibid.

vague words in hard cases, but to determine true meaning directly.[47] This view of law as a natural kind is developed further in Moore's later work, and I turn to that work and its relationship to epistemicism now.

In "Law as Functional Kind," Moore's realist commitments are stated in the following way: "(a) moral qualities such as justice exist (the existential condition); and (b) such qualities are mind- and convention-independent (that is, their existence does not depend on what any individual or group thinks—the independence condition)."[48] Now, here is the question that must be answered in determining the extent to which Moore's theory is inconsistent with epistemicism: In what way do legal propositions relate to the truth of corresponding moral propositions? He first dismisses the idea that the law is related to morality based upon either contingent necessity or analytic necessity.[49] Rather, for Moore, the proper connection between legal propositions and moral propositions should be thought of in terms of Saul Kripke's notion of "metaphysical necessity." Under this framework, a truth is metaphysically necessary if it depends upon how the world is only, not the conventions of human language use; for example: "'Water is H^2O' is (as far as we know) a metaphysically necessary truth because something wouldn't be water if it weren't H^2O."[50] The crux of Moore position is that the relationship between legal and moral propositions should be thought of in these terms. Regardless of conventional or social use, the nature of law is such that it necessarily includes things like justice, and, conversely, decisions that are unjust are necessarily not legal. The idea, then, is to describe law itself via direct reference, not some concept of what law is. While Moore

[47] Ibid. p. 326.
[48] Moore, "Law as a Functional Kind," p. 190.
[49] Ibid. pp. 198-99. Moore describes contingent necessity as "a generalization that is true of a finite sample of things but that is not necessarily true because of the nature of the kind of things making up the sample," while he describes analytic necessity as "part of the meaning of 'law.'" Ibid.
[50] Ibid. p. 200.

acknowledges that it is more controversial to think of terms like "law" as natural kinds than terms like "gold" (because terms like "law" might seem to lack a natural essence that dictates its meaning), he is resigned to the idea that law has a nature that gives meaning to the word "law" because law does in fact exist (unlike, say, unicorns, which do not exist and whose nature is thus determined by our concepts only).[51]

Moore assuages the controversial nature of his position by introducing the notion that law should also be thought of as a functional kind, which describes law's nature via function, not structure. To describe a thing (Moore uses the following analogy of sleep and a beating heart in the human body) as a functional kind means:

> "Sort[ing] through all the consequences of that activity in light of a hypothesis both about there being some larger system in which the activity occurs and about the system having an overall goal. The heart's beating and sleeping are both activities within (or of) the system we call the human body. We think such a system itself has a function or goal, namely, physical health. Such a system-wide goal is aided by some consequences of a heart's beating and by some consequences of sleeping, and not by others. We call the former the function of the heart or of sleep.[52]

Moore apples this framework to the law by again comparing the law to a natural kind like gold. We have knowledge of gold's essential nature (atomic structure) and examples of gold (a wedding ring, for instance). However, initially, we were ignorant of gold and we perhaps confused gold with fool's gold and labeled them both "gold." Our initial hypothesis about the class "gold" was thus inchoate. Ultimately, we acquired knowledge of gold's essential nature through more mature theories, amending the inchoate hypothesis accordingly. Nevertheless, it is still possible that our theory of gold may be improved should we acquire additional knowledge about the nature of the kind. For Moore, this is what the law is like: "The theory that law is a

[51] Ibid. p. 205.
[52] Ibid. p. 210.

kind and that its nature is given by such and such a goal is just that, a theory, falsifiable as is the atomic theory of gold."[53] The law is, in other words, like gold.

It is for this reason that Moore's theory does not correspond to epistemicism. To be sure, Moore's theory acknowledges that we have a great deal of ignorance about the law. But the nature of that ignorance is like the ignorance one might have about whether a particular mineral or chemical element is gold or fool's gold. Law does not depend upon language, usage, or a social conception, but rather it has a natural essence that is metaphysically necessary—like gold. Ignorance of such natural boundaries—like whether something is or is not gold—is not the kind of ignorance that is relevant to Williamson's epistemicism; indeed, such ignorance involves no vagueness at all. As discussed in Part 1, epistemic vagueness is about ignorance in borderline cases in which we have inexact knowledge—regarding meaning, not natural divisions. The extension of vague predicates like "reasonable assistance" is not what Williamson would call metaphysically necessary because the inexactness of "reasonable assistance" is based upon differences in meaning. It is true that Williamson's treatment of gold as a thing that is "stabilized by natural divisions" is the same as Moore's. However, natural-kind terms are not vague under Williamson's account—even if we are unsure whether something is gold or fool's gold—because they have fixed boundaries that make our ignorance of them categorically different than our ignorance of the semantics of a term.[54]

2.3. *Epistemicism Must Incorporate Normative Considerations within the Domain of the Law*

So far we have seen how Dworkin's theory is perhaps internally consistent with the key features of epistemicism and how Moore's theory is perhaps not. I will now examine whether epistemicism itself is adequate to account for vagueness in the legal domain—particularly in the

[53] Ibid. pp. 220-21.
[54] See supra note 19.

case of Dworkin's theory of law as integrity—ultimately concluding that epistemicism is an incomplete account of vagueness in the legal domain unless it accounts for ignorance of normative considerations.

If Dworkin's theory takes the position that there is a fact of the matter about legal propositions—for example, one unique answer regarding the utterance, "Jane did not offer reasonable assistance"—but we are incapable of knowing the cut-off point in such borderline cases because of our limited grasp of the meaning and use of expressions like "reasonable assistance," then Dworkin's theory seems roughly consistent with epistemicism. Like epistemicism, law as integrity does not require us to be completely skeptical about our knowledge because an actual judge surely knows enough about the meaning of "reasonable assistance" to know that taking two steps into a shallow lake is not in extension of "reasonable assistance," for example. However, actual judges simply do not have the ability to grasp the use and meaning of "reasonable assistance" in borderline cases in the way Hercules—or an omniscient speaker—would. This is all well in good. We have located a theory of legal interpretation that is internally consistent with the major features of epistemicism. But, unlike epistemicism, the conclusion of Dworkin's theory—that actual judges are incapable of deriving unique answers in borderline cases as Hercules would—is based upon more than an actual judge's limited grasp of the use and meaning of vague expressions in borderline cases. In addition, it is based upon normativity. This seems inevitable for any legal theory that is internally consistent with the key features of epistemicism I have discussed.

Here is brief account of why this might follow in light of Dworkin's theory. First, if the extent to which one has a moral duty to reasonably assist others is a normative question, and the extent to which one has a legal duty to reasonably assist others is a question that must be

answered in the law, then the law overlaps with this normativity *prima facie*.[55] Second, if a theory of legal interpretation holds that there is a fact of the matter about legal propositions—including with respect to the any relevant overlapping normative considerations therein—then any ignorance regarding the unique fact of the matter in legal propositions is based upon ignorance of normativity in addition to ignorance regarding use and meaning of any vague expressions within legal propositions. Third, epistemicism holds that we are incapable of knowing the cut-off point in borderline cases based upon our ignorance of the meaning and use of vague expressions only—due to ignorance of semantic laws. Finally, then, epistemicism does not account for vagueness in the legal domain because it does not account for the requisite normative considerations in legal theories that are otherwise consistent with epistemicism.

What is the upshot? Well, we should acknowledge that there are theories of legal interpretation—chiefly, law as integrity—that are internally consistent with the fundamental features of epistemicism. However, within the domain of law—and specifically in the case of legal theories that are internally consistent with epistemicism—epistemicism itself accounts for vagueness incompletely. While an actual judge's ignorance may be based upon ignorance of the use and meaning of vague expressions, it is also based upon normative considerations. Unlike the ideal judge Hercules, it would be impossible for an actual judge to have complete knowledge all the competing normative (and non-normative) considerations relevant to a particular borderline case. But such normative considerations play a role in the actual judge's ignorance, and epistemicism does not account for that. Accordingly, epistemicism does not appear to be an account of vagueness that may be applied generally across all domains. Rather, it must be expanded to include normative considerations in certain domains like law.

[55] Of course, the respective moral and legal duties may very well differ regarding the extent to which one is required to reasonably assist others, but this does not change the fact that the law is engaged in questions of normativity.

To be sure, this raises more questions about both epistemicism and about legal theories like law as integrity. First among these questions might be the following; If the law—specifically, the unique answer in a particular case—is based upon both moral and non-moral facts, then how exactly does the law relate to such moral facts and what does that mean for the law's ontology?[56] While Dworkin's theory commits him to some form of realist metaphysics, he does not elaborate upon his ontological commitments with respect to the law in this way. The possible answers to this question are familiar (e.g., moral naturalism, moral non-naturalism, and so on), but it is not clear which of the possible answers is the most coherent with respect to a legal theory that is consistent with epistemicism. Alternatively, one might object that much of the law simply has nothing to do with normativity, making epistemicism in its traditional form more or less able to account for vagueness in the legal domain. For instance, does a law prohibiting vehicles in a public park apply to a military truck mounted on a pedestal, intended to serve as a war memorial? H.L.A. Hart's well-known hypothetical demonstrates—it might be argued—that the "penumbra" of a law's meaning has nothing to do with normativity.[57] There are perhaps two responses to this objection: (1) So what. Even if there are some laws for which normativity does not have a direct bearing, many laws do (and the ones that don't simply have one less factor to consider in reaching the right answer); or (2) Even the laws that seem to be the most morally irrelevant have bearing on justice, including the overall justness of the legal system; otherwise, they are not laws at all, but more akin to words on a piece of paper. There is obviously much more that could be said regarding both this objection and the metaethical

[56] See Moore, "Legal Reality," for an overview of the various answers to this question.
[57] H.L.A. Hart, "Positivism and the Separation of Law and Morals", *Harvard Law Review* 71 (1958): pp. 606-15; see Frederick Schauer, "A Critical Guide to Vehicles in the Park", *New York University Law Review* 83 (2008): pp. 1109-34, for a thorough analysis of Hart's hypothetical and the responses to the hypothetical, including Lon Fuller's vehicle-on-a-pedestal counterexample. By "penumbra," Hart meant the vague area around a law's edge. So while we are certain that the law's core means that standard automobiles are clearly vehicles and a man jogging in the park is clearly not a vehicle, it is less clear how the law's penumbra relates to things like toy vehicles, bicycles, and so on. Ibid. p. 1126.

questions that my argument raises. It is possible that these objections and questions mean that epistemicism is wrong, or else simply incoherent with respect to theories of legal interpretation. But there is no reason to think this simply because difficult questions and objections may be raised. Depending upon one's ontological commitments, there exist equally good answers to these questions and objections. I am content to leave those questions aside for the time being, with the hope that I have answered two different questions: (1) What is required of a particular legal theory in order for it to be internally consistent with epistemicism, and (2) what is required of epistemicism in order for it to account for such legal theories.

To take stock more broadly, it is important to again acknowledge the vast amounts of excellent work that is being conducted on vagueness and the law. While much of this work focuses upon practical concerns regarding problems that vagueness creates for legislation, or, alternatively, the extent to which vagueness might be a useful in tool in the law, there has been very little theoretical work regarding the issue explored here: Whether a particular theory of legal interpretation is consistent with a particular account of vagueness. This is an important question in philosophy of law, and I have modestly tried to begin filling the void with this paper. As a first stab at this task, I have explored two canonical legal theories that seem most aligned with epistemicism. Ultimately, I think that there is an interesting payoff from this exploration, namely: That there are legal theories that are internally consistent with the central precepts of epistemicism; however, with respect to the law—and particularly with respect to legal theories that are internally consistent with epistemicism—vagueness cannot be explained simply by our ignorance of the meaning and use of vague expressions. Moving forward, it would be of great interest to explore the extent to which epistemicism might be expanded to account for ignorance of the requisite normative considerations in legal theories with which it is otherwise aligned.

ACKNOWLEDGMENTS

For their helpful and extensive comments, I am grateful to Trenton Merricks and two anonymous reviewers at *Law and Philosophy*. I also thank the participants at the 2015 Pacific Meeting of the American Philosophical Association, in Vancouver, British Columbia.

University of Virginia

Corcoran Department of Philosophy

Charlottesville, Virginia