

# A PUZZLE ABOUT VAGUENESS, REASONS, AND JUDICIAL DISCRETION

Hrafn Asgeirsson\* 

*School of Law, University of Surrey, Guildford, UK*

## ABSTRACT

The following two theses seem both plausible and consistent: in cases where it is indeterminate whether the relevant legal language applies to the relevant set of facts, officials are not bound to decide the case one way rather than the other, but may reason either way; all reasons for action are—in some relevant sense—knowable. In this paper, I point out what I take to be a robust but unacknowledged tension between these two claims. The tension requires some careful teasing out, but the basic idea is that given certain further plausible assumptions concerning law, language, and normativity, the two claims turn out to be inconsistent. In addition to examining the sources of the tension in some detail, I also address several possible objections to my argument and discuss which of the many theses should be rejected.

---

The following two theses seem both plausible and consistent: first, that in cases where it is indeterminate whether the relevant legal language applies to the relevant set of facts (and, crucially, no other legal norms conclusively apply), officials are not bound to decide the case one way rather than the other, but may reason either way; second, that all reasons for action are—in some relevant sense—knowable. In this paper, I point out what I take to be a robust but unacknowledged tension between these two claims. The tension requires some careful teasing out, but the basic idea is that

\* Thanks to everyone who took time to talk to me about the ideas in this paper (and related issues). And in addition to the people acknowledged in various footnotes, special thanks also to Mark Schroeder, Lawrence Solum, John Mikhail, Ezequiel Monti, and Teresa Marques, as well as audiences at the University of Girona Legal Philosophy Research Group, the Pompeu Fabra Law and Philosophy Research Group, the Georgetown University Law Center, the 2016 University of Oslo CSMN Workshop on Law and Language, the 2017 EUI International Workshop on Reason-Giving in Law in Florence, the Special Workshop on the Normativity of Law at the 2015 IVR in Washington, DC, the 2015 Legal Philosophy Workshop at the University of Edinburgh, and the 2016 Aristotelian Society Open Sessions at Cardiff University.

given certain further plausible theses/assumptions concerning law, language, and normativity, the two claims turn out to be inconsistent. In addition to examining the sources of the tension in some detail, I also address several possible objections to my argument and discuss which of the many theses should be rejected. The aim, however, is not to survey all the options—a paper like this is by necessity tentative to some extent. Rather, the point is, first, to make clear why much more needs to be said about the way vagueness affects the reasons for action that (the enactments of) legal norms give rise to and, second, to identify some promising ways to relieve the tension between the two theses.

I will anchor my discussion in Joseph Raz's work, since he is the most explicit proponent of the two focal theses (as well as some of the other theses involved). The paper is intended to be much more than a note on Raz's views, however. For one, nearly every legal theorist joins him in holding some version of the judicial discretion thesis. And although far fewer explicitly endorse the knowability thesis, they certainly should—or so I argue. Both theses are eminently plausible, I believe, as are the other claims on which the puzzle relies; hence this paper. There must of course be a way out, but each will carry with it deep commitments in philosophy of language and/or metaethics—or so I hope to show.

## I. THE TENSION BETWEEN KNOWABILITY AND DISCRETION

In *From Normativity to Responsibility*, Raz proposes an account of what it takes for something to be a reason for action.<sup>1</sup> For a fact to constitute a reason, he says, it is not enough that it counts in favor of some action; it also has to be true that the agent for which it is a reason (if it is a reason) can respond to that fact using her rational powers (i.e., the faculty of Reason). Or else it would not be able to play a role in explaining the agent's action. This is Raz's *explanatory/normative nexus* (henceforth, "Raz's Nexus").<sup>2</sup>

Among other things, Raz's Nexus implies two substantial constraints on what it takes for something to be a reason—knowability and availability. If a fact is unknowable to the relevant agent, she could not—even given her best efforts—respond to it using her rational powers. As a result, that fact does not constitute a reason for her, even if it counts in favor of some action that is available to her. Same goes if the relevant action is not available to the agent—in that case, she cannot rationally respond to the fact in question either, even if it is a knowable one. Consequently, it does not constitute a reason for the agent to perform that action.

1. Note that the concern here is what it takes to *be* a (normative) reason, as opposed to what it takes to *have* such a reason. For a discussion of this distinction, see, e.g., Mark Schroeder, *Having Reasons*, 139 *PHIL. STUD.* 57 (2008).

2. See JOSEPH RAZ, *FROM NORMATIVITY TO RESPONSIBILITY* (2011), at 87.

In this paper, we will be concerned only with the knowability condition:

KNOWABILITY: If  $p$  is a reason for  $A$  to  $\phi$ , then  $p$  is knowable to  $A$ .<sup>3</sup>

Raz's agent-relative version of the knowability constraint has particular appeal in the context of the law, of course, evidenced for example by the significance of promulgation in modern legal systems. Indeed, arguably, the primary function of promulgation is to ensure that the law is knowable to those subject to it, and thus able to guide their behavior.<sup>4</sup> I hope to show, though, that the puzzle arises for anyone who thinks—much more generally—that in order for a fact to count as a reason for action, it (at a minimum) has to be knowable to a highly idealized agent. This is an extremely weak—and overwhelmingly plausible—accessibility thesis, arguably consistent with both internalism and externalism about reasons. And, as it turns out, it is all we need to fuel the puzzle.<sup>5</sup>

My aim, then, is to show that—at least *prima facie*—any version of KNOWABILITY will be in deep tension with the way we tend to think about judicial discretion:

DISCRETION: If it is indeterminate whether the relevant legal language applies to the relevant set of facts, officials are—assuming no other legal norms conclusively apply—not bound to decide the case one way rather than the other, but may reason either way.<sup>6</sup>

As stated, DISCRETION focuses only on specifically relevant necessary elements of judicial discretion—the principle is decidedly not meant to capture the full nature of the phenomenon. And while there is of course

3. *Id.* at 87, 110, 122.

4. Which is not to say that promulgation is not otherwise valuable; for a discussion, see, e.g., Andrei Marmor, *The Ideal of the Rule of Law*, in *A COMPANION TO THE PHILOSOPHY OF LAW AND LEGAL THEORY* 666, 670–671 (Dennis Patterson ed., 2010).

5. I say more about the plausibility of this thesis in [Section IIA](#) below.

6. See, e.g., JOSEPH RAZ, *THE AUTHORITY OF LAW: ESSAYS ON LAW AND MORALITY* (1979/2002), at 96, 113–114. I should note that, for ease of exposition only, I don't distinguish here between a norm (or set thereof) *conclusively* applying with respect to a given case and the norm (or set thereof) *requiring* a course of action, in virtue of so applying. Like John Gardner, *Concerning Permissive Sources and Gaps*, 8 *OXFORD J. LEGAL STUD.* 457 (1988), and many others, I think that there can be cases where a norm (or set thereof) conclusively applies, in the sense that the relevant officials have most reason to decide accordingly, yet they are not required to do so. In such cases, the sources of the norms that settle the balance of reasons are *permissive*, as it's often called, though, as Leslie Green, *Positivism, Realism and Sources of Law*, in *THE CAMBRIDGE COMPANION TO LEGAL POSITIVISM* 33 (Torben Spaak & Patricia Mindus eds., 2021), rightly points out, what's involved goes beyond a "bare" permission, given that such sources generally provide positive reasons as well, i.e., reasons in favor of the relevant course of action. For discussion of *ought* versus *must* more generally, see, e.g., Justin Snedegar, *Reasons, Oughts, and Requirements*, in *OXFORD STUDIES IN METAETHICS: VOLUME 11* 155 (Russ Shafer-Landau ed., 2016). Substantively, my own view is that permissive sources are considerably more prevalent than many take them to be, certainly in relation to borderline cases, but for those who think otherwise, the puzzle can simply be reformulated (either generally or with respect the relevant subset of cases); see note 8 below.

much to be said about the power that judges have to resolve cases of vagueness (and other problem cases),<sup>7</sup> the puzzle discussed here concerns the conceptual coherence of the foundations of judicial discretion, rather than its substantive contours. Accordingly, I want—at least for the purposes of this paper—to remain as neutral as possible on substantive issues and take for granted only a minimal characterization of what such discretion (among other things) consists in; formulating the thesis merely in terms of the absence of bindingness, and the correlated presence of permission, is distinctly meant to accommodate a maximal range of positions on the nature of judicial discretion. What form the relevant freedom takes is an important matter, but outside the scope of the paper.<sup>8</sup>

That said, there are limits to how neutral I can be. Most importantly, while minimal, DISCRETION is still couched in normative terms. Assuming that normative concepts/terms are best analyzed in terms of reasons, the thesis therefore presupposes a reasons-based account of legal norms. Indeed, without such an account, the puzzle does not arise. For some, this will mean that they get off the bus already at the outset. However, rejecting a reasons-based analysis of judicial discretion does more than simply deflate the concept, I think. I would go so far as to say that it amounts to rejecting the notion altogether. Thus, insofar as we take judicial discretion seriously, we are already assuming a reasons-based analysis of legal norms.<sup>9</sup> And on such an account, the most straightforward basis for judicial discretion is that—relative to each case—such discretion depends (partly, but crucially) on how the relevant legal norms apply to the subject. In cases where it is indeterminate whether a defendant has violated the relevant legal norm (and no other norms conclusively apply), officials are not bound to decide the case one way rather than the other, and a subject violates a legal norm only if she acts contrary to a reason provided by that norm. Together, this entails that in cases where it is indeterminate whether a defendant has acted contrary to a reason provided by the relevant norm, officials are not bound to decide the case one way rather than the other.

7. See, e.g., H.L.A. HART, *THE CONCEPT OF LAW* (1994); H.L.A. Hart, *Discretion*, 127 HARV. L. REV. 652 (2013); Ronald Dworkin, *Judicial Discretion*, 60 J. PHIL. 624 (1963) [hereinafter Dworkin, *Judicial Discretion*]; Ronald Dworkin, *The Model of Rules*, 35 U. CHI. L. REV. 14 (1967); Joseph Raz, *Legal Principles and the Limits of Law*, 81 YALE L.J. 823 (1972) [hereinafter Dworkin, *Model of Rules*]; RAZ, *supra* note 6; and Gardner, *supra* note 6. Note that in this paper, I am only concerned with cases involving vagueness. I take it that what I say applies equally to other types of linguistic indeterminacy (given its general relation to unknowability), but on this occasion I do not have space to thoroughly discuss the extent to which the puzzle generalizes. Thanks to an anonymous referee for pressing me to address this.

8. It is worth noting, though, that the puzzle can be reproduced without too much difficulty for those who hold that borderline cases (necessarily or contingently but commonly) trigger some relevant mandatory “tie-breaking” second-order norms (e.g., relating to burden of proof or other procedural concerns, broadly understood) and that judges therefore never really have said discretion; the puzzle then simply relates to the conditions under which such triggering is supposed to occur. Thanks to an anonymous referee, and to David Prendergast as well, for pressing me to clarify this.

9. I say a bit more about this at the end of [Section II.E](#) below.

Whatever substantive form it takes, the normative power that lies at the foundation of judicial discretion, then, is—at least on what I take to be the most straightforward account—a function of the normative status of (the actions of) the relevant subject.

I should note some further reasons I will not seriously consider rejecting DISCRETION here. In addition to finding it overwhelmingly plausible, it is a core commitment of almost all theories of law, whether positivist or nonpositivist. Almost everyone holds that, for at least some robust subset of borderline cases, judges may reason either way (in some significant sense)—although the nature and size of the relevant set may vary, and people have different ideas about what kind of reasoning is appropriate.<sup>10</sup> I do not mean to entirely rule out that the tension I identify might ultimately provide some reason against the thesis, but I think that once we consider some of the available moves, that reason (if any) would be a very weak one.

*The puzzle explained:*

I will try to make as explicit as possible how the puzzle arises. There are many theses involved—some more controversial than others—but each offers a point of pressure, a possible way out. Hopefully, at the end of the paper, we will have some idea which of all the many moving parts are the most movable.

Now let us see how the tension between DISCRETION and KNOWABILITY arises, step-by-step—first by appeal to abstract principles and then illustrated by example. DISCRETION typically rests on a number of claims, one of which concerns the nature of vague language:

VAGUE LANGUAGE: If an object (broadly construed) is a borderline case of a predicate (relative to a context of utterance), then it is indeterminate whether the predicate applies to that object (in that context).<sup>11</sup>

Coupled with the thesis that—either generally or in a significant set of relevant cases—the legal content of a legal provision directly corresponds to its linguistic content, VAGUE LANGUAGE produces a further thesis:

VAGUE LAW: If an object (broadly construed) is a borderline case of the (usually complex) operative predicate of a legal provision, then it is indeterminate whether the corresponding legal norm applies to that object.<sup>12</sup>

10. To my knowledge, Ronald Dworkin is the only exception; see Dworkin, *Judicial Discretion*, *supra* note 7; Dworkin, *Model of Rules*, *supra* note 7. Thanks to an anonymous referee for helpful comments here.

11. For the sake of constructing the puzzle, I assume that there are determinate borderline cases. This is a widely held assumption (and accepted, for example, by Raz), but I discuss the possibility of rejecting it in [Section II.G](#).

12. See RAZ, *supra* note 6, at 72. The term “operative predicate” here just means the predicate determining the facts on which the relevant legal norm operates. Thanks to an anonymous referee for helpful comments about how to formulate this thesis.

Now, for ease of exposition, I here adopt an unrestricted version of the above thesis relating linguistic content and legal content, but the argument will go through on any restricted version thereof, too (and most everyone would accept *some* such version of it, irrespective of their other jurisprudential commitments): as long as it is true for some significant set of cases that the legal content of a legal provision directly corresponds to its linguistic content, we have a robust puzzle (holding other theses constant).<sup>13</sup>

On the assumption that legal norms are best analyzed in terms of the legal reasons of which they consist, or to which they give rise (depending on our favored account of the relationship between norms and reasons), VAGUE LAW entails the following:

VAGUE LEGAL REASONS: If it is indeterminate whether a legal norm applies to an object (broadly construed), then it is indeterminate whether it gives the relevant subject(s) a legal reason for action, vis-à-vis that object.<sup>14</sup>

(This—I take it—is just a limited instance of a more general principle concerning vagueness and normativity: if it is indeterminate whether a norm applies to an object, then it is indeterminate whether it gives the relevant agent a reason for action, vis-à-vis that object.)

The view under consideration, then, entails that if a subject acts in a way that is borderline legal (relative to some provision), it is indeterminate whether she acts contrary to a legal reason (provided by that provision), because it is indeterminate whether such a reason is present. On the assumption that a subject violates a legal norm only if she acts contrary to a legal reason provided by that norm, this further entails that it is indeterminate whether she violates the relevant legal norm. As a result, assuming no other legal norms conclusively apply to the case (such as a rule of lenity), if such a case goes to court, the relevant officials (jury or judges) are therefore not bound to decide the case in one way rather than the other, but may reason either way (in line with one's favored theory of adjudication). Or so this line of reasoning goes.

VAGUE LANGUAGE, however, has other implications, which ultimately get the view under consideration into trouble, insofar as it is supposed to include KNOWABILITY. The root of the problem is that in the literature on vagueness it is widely accepted that borderline propositions are unknowable—either because it is considered a consequence of the

13. Note also that in order to generate the puzzle, we technically do not need *complete* correspondence between linguistic content and legal content—we just need it to be the case that, for a significant set of cases, the legal content in question inherits the linguistic content's "grey area," even if the former were to fail to correspond to the latter in other respects (e.g., by including, say, unstated offense elements and/or justified exceptions, etc.).

14. See RAZ, *supra* note 6, at 74.

particular account in question or because it is taken to be independently motivated:<sup>15</sup>

IGNORANCE OF LANGUAGE: If it is indeterminate whether a predicate applies to an object (relative to a context of utterance), then it is unknowable whether the predicate applies to that object (in that context).<sup>16,17</sup>

Given KNOWABILITY, IGNORANCE OF LANGUAGE entails that in cases of vagueness it is *not* the case that the relevant legal provision gives subjects a legal reason for action. And since both these are—by assumption—determinately true, this further entails that in cases of vagueness the relevant legal provision *determinately* does not give subjects a legal reason for action,<sup>18</sup> which

15. For examples of the former, *see, e.g.*, TIMOTHY WILLIAMSON, *VAGUENESS* (1994), at 185 ff., and SCOTT SOAMES, *The Possibility of Partial Definition*, in *PHILOSOPHICAL ESSAYS, VOLUME 2: THE PHILOSOPHICAL SIGNIFICANCE OF LANGUAGE* 362, 370 (2009); for an example of the latter, *see, e.g.*, Stephen Schiffer, *Quandary and Intuitionism: Crispin Wright on Vagueness*, in *ESSAYS FOR CRISPIN WRIGHT: LOGIC, LANGUAGE, AND MATHEMATICS* 154 (Alexander Miller ed., 2020). Although not explicitly addressed, this is arguably also a consequence of supervaluationism, at least as presented in Kit Fine, *Vagueness, Truth and Logic*, 30 *SYNTHESE* 265 (1975). On Fine's view, when one asserts a vague proposition, one asserts the set of all its admissible precisifications (*id.* at 282); in borderline cases, since at least one member of that set is false, one determinately ought not assert the relevant proposition. It seems to me that similar reasoning would apply to belief (and thus to knowledge). I take it that—similarly—when one believes a vague proposition, one believes the set of all its admissible precisifications. As before, in borderline cases, at least one member of that set is false. Presumably, however, one determinately ought not believe a set of propositions some member of which is false; in other words, determinately, one is not justified in believing borderline propositions, in which case, determinately, one cannot know them. Thanks to Alex Kaisermann for pressing me to clarify this.

16. Note that I take this thesis to entail a corresponding thesis about propositions/facts: if it is indeterminate whether *o* is P, then it is unknowable whether *o* is P. For the purposes of this paper, albeit not generally, these theses can be used interchangeably.

17. IGNORANCE OF LANGUAGE immediately produces a parallel with VAGUE LAW (holding other assumptions constant): if an object (broadly construed) is a borderline case of the (usually complex) operative predicate of a legal provision, then it is unknowable whether the corresponding legal norm applies to that object; this, in turn, entails a parallel with VAGUE LEGAL REASONS: if it is unknowable whether a legal norm applies to an object (broadly construed), then it is unknowable whether it gives the relevant subject(s) a legal reason for action, *vis-à-vis* that object. These two theses help to provide a fuller picture of Raz's treatment of borderline cases—in particular, his commitment to IGNORANCE OF LEGAL REASONS allows us to reliably infer that he accepts IGNORANCE OF LANGUAGE (*see* RAZ, *supra* note 6, at 74)—but they are not necessary to produce the puzzle, since the problematic ignorance concerns the borderline cases themselves, rather than the favoring relation between them and the relevant agent's course of action. Thanks to an anonymous referee for prompting me to clarify this.

18. The assumption here—nearly universally accepted in the logic of vagueness—is that the “determinately” operator is closed under logical consequence. As Elia Zardini, *Higher-Order Sorites Paradox*, 42 *J. PHIL. LOGIC* 25 (2013) points out, arguably, the only author that has seriously questioned this assumption is HARTRY FIELD, *SAVING TRUTH FROM PARADOX* (2008). Field's motivation for rejecting closure, however, is to defuse the semantic paradoxes (such as the Liar Paradox) and so it is unclear whether we should take the rejection to extend to the case of vagueness. Zardini doubts that such an extension could be motivated, *but see* Andrew Bacon, *Non-Classical Metatheory for Non-Classical Logics*, 42 *J. PHIL. LOGIC* 335 (2013), for a discussion indicating that Field is indeed interested in a unified interpretation of the “determinately”

produces an immediate conflict with VAGUE LEGAL REASONS. VAGUE LEGAL REASONS, IGNORANCE OF LANGUAGE, and KNOWABILITY are thus inconsistent.

Moreover, KNOWABILITY entails that if a subject acts in a way that is borderline legal (relative to some provision), then she determinately does not act contrary to a legal reason (provided by that provision). She therefore determinately does not violate the relevant legal norm, assuming (as before) that a subject violates a legal norm only if she acts contrary to a legal reason provided by that norm. Insofar as the matter depends on the subject having violated the norm in question, the relevant officials are thus bound to decide in her favor, assuming (safely, I think) that—as a matter of legal fact—officials are bound to decide in favor of the defendant if she determinately did not violate the relevant legal norm. KNOWABILITY thus also conflicts with DISCRETION.<sup>19</sup>

All this has been rather abstract, so to make things more concrete, let me borrow and extend an example from Raz.<sup>20</sup> Council tax regulations give people whose properties count as dwellings a legal reason to make payment to the council. We can say that in virtue of the relevant regulations, the fact that some property, P, is a dwelling is a legal reason for the owner, A, to make payment to the council. But let's then say that some property, P\*, owned by B, is a borderline case of the predicate "dwelling." This entails that it is indeterminate whether the predicate applies, which, in turn, entails that it is indeterminate whether council tax regulations apply to P\*. As a result, it is indeterminate whether council tax regulations give B a legal reason to make payment to the council. Consequently, if B decides not to pay, it is indeterminate whether she violates the relevant legal norm. If B's case goes to court, the relevant officials therefore are not bound to decide the matter one way rather than the other, but may reason either way, assuming no other legal norms conclusively apply to the case.<sup>21</sup>

operator. Still, rejecting closure under logical consequence remains highly exceptional. Thanks to an anonymous referee for pointing out the need to make this assumption explicit.

19. I should note that although the results of this line of reasoning are similar to Dworkin's reply to what he calls "the argument from vagueness" (see RONALD DWORKIN, *A MATTER OF PRINCIPLE* (1985), at 128–131), the reasoning here is quite different—in particular, it does not depend, as Dworkin's reasoning does, on the disputable universal capacity of external considerations to close what would otherwise be gaps in the law, but rather appeals only to considerations internal to the analysis of normative reasons.

20. See RAZ, *supra* note 6, at 72–73.

21. Note that this conclusion remains even if we think that the legal reason in question must consist not only in P\*'s borderline status (call this "*p*") but also in the relevant legal regulation (call this "*q*"), i.e., if these together constitute the agent's "complete" reason. (See, e.g., JOSEPH RAZ, *PRACTICAL REASON AND NORMS* (1975/1990), at ch. 1 [hereinafter RAZ, *PRACTICAL REASON*], and JOSEPH RAZ, *ENGAGING REASON: ON THE THEORY OF VALUE AND ACTION* (1999) [hereinafter RAZ, *ENGAGING REASON*], though I'm simplifying here for the sake of space.) The reason is that on almost any logic of vagueness, if either conjunct is indeterminate—as *p* is here—then the conjunction is considered indeterminate as well (though some take it to be false, in which case we don't even need the knowability requirement to generate the puzzle); see, e.g., Fine, *supra* note 15. Thanks to Ezequiel Monti, as well as an anonymous reviewer, for pressing me to clarify this.



However, if  $P^*$  is a borderline case of the predicate “dwelling,” then it is unknowable whether  $P^*$  is a dwelling. Given KNOWABILITY, this entails that council tax regulations determinately do not give B a reason to make payment to the council.<sup>22</sup> Here, then, is the first inconsistency. Further, however, KNOWABILITY also entails that if B decides not to pay, then it is *not* indeterminate whether she violates the relevant legal norm: she determinately does not. Thus, assuming no other legal norms conclusively apply to the case, officials are—as a matter of legal fact—bound to decide in her favor, which conflicts with the conclusion about discretion in the previous paragraph. The conclusion from all this is that the view under consideration is inconsistent.

## II. RELIEVING THE TENSION—POSSIBLE WAYS OUT

I hope it has been made clear both how the tension between KNOWABILITY and DISCRETION arises and how many theses are involved. With so many moving parts, it is not feasible to examine each of them here. Instead, I will begin with a few clarifications and then turn to considering some suggestions for which theses to reject. Some of these will aim to preserve as much as possible of the general framework on which the puzzle relies, while others will propose that we give up some of its core tenets—although none of them will consider rejecting DISCRETION, as I have indicated. I hope to show that, although there must of course be a solution to the puzzle, there is no easy way out—rejecting any of the theses involved carries with it deep commitments in philosophy of language and/or metaethics. If what I say is correct, it is clear that much more needs to be said about the way in which vagueness affects the reasons for action that (the enactment of) legal norms give rise to.

### A. Clarifying KNOWABILITY and IGNORANCE OF LANGUAGE

Let me begin by clarifying KNOWABILITY and IGNORANCE OF LANGUAGE a bit, to avoid certain misunderstandings. First, it is important to distinguish KNOWABILITY from the significantly stronger, internalist thesis that in order for something to count as a reason for action, it has to be knowable that the *favoring relation* holds between the relevant fact and some course of action. Some externalists, although by no means all,<sup>23</sup> would presumably reject this stronger accessibility thesis, even in its weakest form,<sup>24</sup> but I take it that there is very little—if any—reason for even the staunchest

22. Here, as before, closure of determinacy under logical consequence is assumed.

23. See, e.g., John McDowell, *Might There Be External Reasons?*, in *WORLD, MIND, AND ETHICS: ESSAYS ON THE ETHICAL PHILOSOPHY OF BERNARD WILLIAMS* 68 (J.E.J. Altham & Ross Hartison eds., 1995).

24. See, e.g., David Copp, *Four Epistemological Challenges to Ethical Naturalism*, 30 *CAN. J. PHIL.* 31 (2000).

externalist to resist the weaker thesis on which the puzzle relies, i.e., that the relevant fact itself has to be knowable.

Second, it helps to get clear about the modality involved in KNOWABILITY. On Raz's account, KNOWABILITY is a rather strong condition: for something to count as a reason, someone at the time would have found out about it had they put in their best efforts.<sup>25</sup> Most naturally, perhaps, this is read to apply to the agent herself, for which the relevant consideration is or is not a reason—namely that had *she* put in her best effort, she would have come to know the relevant fact.

Now, even if we go along with the idea that there are epistemic accessibility constraints on what it takes for something to count as a reason for action we might worry that Raz's constraint is too strong.<sup>26</sup> For the purposes of our discussion here, however, the strength of the constraint doesn't really matter, since my argument—if it goes through—will go through even on extremely weak readings of KNOWABILITY. The tension with which I'm concerned still arises, for example, given the weaker thesis that something counts as a reason (for a particular agent) only if it can in principle be known by an agent with normal human epistemic capacities under idealized conditions.

Although there are even weaker versions of KNOWABILITY, this naturalistically specified level of constraint ensures that the tension arises even if we adopt an orthodox epistemicist account of vagueness (on which an omniscient being, for example, would know all the true borderline propositions). Of course, if we instead adopt a semantic or ontological account (or certain unorthodox epistemicist accounts), then an even weaker constraint suffices, since—on such accounts—borderline propositions cannot even be known by cognizers with perfect epistemic capacities. The bottom line is that as long as there is *any* epistemic accessibility constraint on what it takes to be a reason, this will cause serious misalignment in the set of theses relevant to judicial discretion.

I take these weak versions of KNOWABILITY to be overwhelmingly plausible, and consistent with both internalism and externalism about reasons. Whatever reasons are, there is surely *some* link between the relevant facts and our access to them, however weak and counterfactual. I'm not sure I would even know how to conceive of a reason that an idealized agent, under idealized conditions, cannot even in principle be responsive to—at least under ideal conditions, the favoring relation between some fact *r* and some action *A* seems to go hand in hand with the ability of an ideal agent to do *A* on the basis of *r*. Hence, we ought to embrace at least some version of KNOWABILITY, which is all we need to get the puzzle going.

25. See RAZ, *supra* note 2, at 110 n.6.

26. See, e.g., Ruth Chang, *Raz on Reasons, Reason, and Rationality: On Raz's from Normativity to Responsibility*, 8 JERUSALEM REV. LEGAL STUD. 1, 10–11 (2013).

Third, it is important to note in relation to IGNORANCE OF LANGUAGE that when we are considering whether a proposition is knowable (relative to a cognizer), we hold fixed everything that is relevant to its truth-value in the world of evaluation. As Wright (2003) points out, the modality involved in knowability claims—at least as they feature in the debate about the nature of vagueness—is “constrained by the distribution of truth values in the actual world.”<sup>27</sup> Among other things, this allows us to fend off the immediate worry that some borderline propositions are in fact knowable. Consider, for example, Joe, who is (actually) borderline tall. Now, he might of course have been determinately tall, in which case the proposition *that Joe is tall* is—in some sense—knowable, despite being indeterminate in the actual world. This might be due to John’s being taller in some possible world, or it might be due to a different standard of evaluation for “is tall.” It is far beyond the scope of this paper to discuss further ways in which borderline propositions might be knowable in this unconstrained sense. This much is clear, however: it is important to keep the constrained modality in mind throughout the paper, or else my argument will be unduly taken to fail for (at a minimum) any contingent singular borderline proposition, i.e., for any case in which a contingent property is predicated of a particular object (such as when “dwelling” is predicated of P\*).

#### B. “No Reason Either Way”?

It might be suggested that—contrary to what has been assumed so far—cases of vagueness in the law should be analyzed not as cases in which it is indeterminate whether the law provides a reason to do as directed (call this the *standard analysis*), but as cases in which the law fails to provide reason either way. If that is correct, then the fact that KNOWABILITY entails the absence of reasons is not inconsistent with DISCRETION. We just have to be careful to keep in mind that this absence “goes both ways,” as it were.<sup>28</sup>

It may seem reasonable, for example, to say that—due to the borderline status of P\*—council tax regulations neither require B to pay nor permit her not to pay. And that, as a result, the law neither requires officials to treat B as having had reason to act as prescribed nor requires the contrary. Perhaps an “absence of reason either way” analysis dissolves the problem, by allowing us to block the move from VAGUE LAW to VAGUE LEGAL REASONS?

Despite the *prima facie* appeal of the suggested analysis, the answer—I think—is no. Consider, first, a reading of it on which the relevant notion of permission is that of *weak* permission, characterized as nothing more than the absence of prohibition. On such a reading, we can simply reproduce the puzzle, since it will be trivially true that whatever is not prohibited

27. Crispin Wright, *Vagueness: A Fifth Column Approach*, in *LIARS AND HEAPS* 101 (J.C. Beall ed., 2003).

28. Something like this analysis is arguably suggested, for example, by the treatment of borderline cases in Timothy Endicott, *Interpretation and Indeterminacy*, 10 *JERUSALEM REV. LEGAL STUD.* 46 (2014).

is permitted (call this the *permissive closure principle*).<sup>29</sup> Given permissive closure, the proposed analysis entails a contradiction, in which case borderline cases are analytically impossible. The permissive closure principle thus functions much like the knowability constraint does in the puzzle as originally presented.

If, on the other hand, we interpret permission *strongly*, then the proposed analysis says—roughly—that, in the relevant legal system, there is neither a norm requiring  $\phi$ -ing nor a norm permitting not  $\phi$ -ing. On this reading, the permissive closure principle is far from trivial, but also highly controversial. On some theories of law, the principle will be necessarily true, in which case there are no gaps of this sort;<sup>30</sup> weak and strong permission are equivalent and the puzzle can therefore—as before—be reproduced. On other theories of law, it is contingent whether the principle is true, relative to particular systems of law.<sup>31</sup> Thus, if the principle is *false* relative to a given legal system, then that system will contain gaps of the proposed sort—so at least for those systems, the puzzle is not reproduced. However, in the remaining systems, the fact that the principle is true entails that there is no discretion in borderline cases, which is problematic, since—intuitively—borderline cases (and thus judicial discretion) remain even when a system permits whatever is not prohibited.<sup>32</sup>

This last point leads me to a more fundamental worry about the suggested analysis, namely that it makes no distinction between borderline cases and cases about which the law is genuinely silent. This is a distinction that we have reason to preserve, and one that Raz—in particular—has been careful to make.<sup>33</sup> Council tax regulations, for example, neither require me

29. See, e.g., RAZ, *supra* note 6, at 76–77. On the weak reading, the permissive closure principle is true simply in virtue of the fact that the notions of *requirement* and *permission* are interdefinable, i.e.,  $Pp =_{df} \sim O\sim p$  and  $\sim O\sim p =_{df} Pp$ . For a discussion of strong versus weak permission, see, e.g., G.H. VON WRIGHT, *NORM AND ACTION* (1963), at 86–87; CARLOS ALCHOURRÓN & EUGENIO BULGIN, *NORMATIVE SYSTEMS* (1971), at ch. 7; Eugenio Bulygin, *On Legal Gaps*, 2002–2003 *ANALISI E DIRITTO* 21 (2002); PABLO NAVARRO & JORGE RODRÍGUEZ, *DEONTIC LOGIC AND LEGAL SYSTEMS* (2014), at 78–80; Jorge Rodríguez, *Normative Systems, Legal Gaps, and Logical Closure* (unpublished manuscript). *But see* RAZ, *PRACTICAL REASON*, *supra* note 21, at 85 ff., for doubts about whether the distinction really tracks a difference in permission, as such.

30. See RAZ, *supra* note 6, at 77 n.19; see also THOMAS HOBBS, *LEVITHAN* (1651/1994), at ch. 21, and HANS Kelsen, *GENERAL THEORY OF NORMS* (Michael Hartney trans., 1979/1991), at 131–132. According to Raz, the principle may be contingent with respect to normative systems generally, but it is true of systems in which the so-called Sources Thesis is true, which—in his view—is necessarily the case with law. The Sources Thesis—associated with exclusive legal positivism—says, roughly, that all norms of a given system have social sources; see RAZ, *supra* note 6, at 47.

31. See, e.g., JULIUS STONE, *LEGAL SYSTEM AND LAWYERS' REASONINGS* (1964); ALCHOURRÓN & BULGIN, *supra* note 29, at ch. 7; NAVARRO & RODRÍGUEZ, *supra* note 29, at 158–166.

32. Of course, closure rules for borderline cases are possible—for antipositivist and positivist suggestions, see, e.g., DWORKIN, *supra* note 19, at 128–131; and Gardner, *supra* note 6, respectively. As many have pointed out, however, higher-order vagueness would prevent us from fully eliminating vagueness from the law. The sources of such closure rules would also have to be binding, as opposed to permissive, as discussed in note 6, *supra*.

33. Most authors writing on the nature and possibility of gaps in normative systems embrace this distinction, but for Raz's own discussion, see RAZ, *supra* note 6, at 70–77.

to eat healthily nor permit me not to eat healthily. And they neither require me to wash my car nor permit me not to wash my car. But to say that officials have any form of discretion with respect to these circumstances would be extremely odd, absurd even; there seems to be a clear jurisdictional difference between borderline cases and cases about which the law is silent, which is unaccounted for by the proposed analysis. Granted, there is (at least on certain views of vagueness) a sense in which both types of cases are undecided and the analysis perhaps manages to capture this, but, on the assumption that judicial discretion is a function of the reasons that applied to the subject at the time of acting, collapsing the distinction between borderline cases and silences severely overgenerates such discretion.

I do not mean to indicate that it is impossible to meet this worry, but it would require an account of judicial jurisdiction robust enough to adequately distinguish between those cases about which the law is completely silent and those about which it is silent but still in some relevant sense “purports” to regulate, despite the fact that—on the suggested analysis—both types of cases are supposed to have the same basic normative structure (characterized by an absence of reasons). This is a tall order, I think—in part because I doubt that there are any cases of the latter sort—but if such an account could be had, this might allow the proponent of the “absence of reason either way” analysis to say that judicial discretion is a function not of the absence of reasons *tout court* but of the absence of reasons within the scope of the relevant court’s jurisdiction, and that cases of vagueness are in fundamental respects like the ones that the law purports to regulate.

Although I am skeptical that the relevant distinction can be adequately made, one reason for thinking that it *can* might come from considering laws that incorporate reference to evaluative notions. Some might want to claim, for example, that due to the unconscionability doctrine in contract law, parties to a contract are—prior to adjudication—neither conclusively required to act in accordance with the relevant agreement nor conclusively permitted to act contrary to it. And that this is clearly different from how the law treats people’s diets and car maintenance: the law in some relevant sense purports to regulate the former but not the latter, or so the thought goes. Of course, to raise the mere possibility of such cases is not to provide an account of judicial jurisdiction (i.e., of what it consists in and what follows from that), but thinking along these lines perhaps suffices to show that there may be principled ways to address the overgeneration worry.

This may be so. Still, it remains the case that, in addition to the need for a suitable and robust enough account of judicial jurisdiction, we would both have to adopt a theory of law on which the strong permissive closure principle is contingent and concede that the proposed “solution” to the puzzle would apply only to those systems relative to which the principle is false. And we would also be stuck with the counterintuitive result that in systems

in which there is no discretion in intrajurisdictional cases about which the law is silent (i.e., systems relative to which the permissive closure principle is true), there is—by the same token—no discretion in borderline cases; whereas, intuitively, we could have the latter without the former, and vice versa. The standard analysis, by contrast, does not suffer from any of these problems. All in all, then, vagueness in the law seems better analyzed in terms of it being indeterminate whether the relevant subject has a reason to act as the relevant provision prescribes.<sup>34</sup>

### C. Retroactivity to the Rescue?

Most legal systems allow law to be applied retroactively in certain cases, either by legislation or by judicial decision-making. In the United States, for example, retroactive tax legislation is not unheard of. And in the United Kingdom, it seems fair to say that judges have significant leeway to introduce changes into the law, based on extralegal considerations. When a common law rule produces an unjust outcome, for example, the court may decide that certain features of the relevant case allow it to be distinguished from other cases to which the rule applies; in this way, conduct that was determinately legal at the time is sometimes made illegal retroactively. This complicates the picture I have presented, since every actual legal system allows for at least some exceptions to the rule that cases in which the defendant determinately did not violate the legal norm in question ought to be decided in her favor. This will not change the fact that, given *IGNORANCE OF LANGUAGE*, *VAGUE LEGAL REASONS* and *KNOWABILITY* are inconsistent, but it certainly makes it less clear that the same goes for *KNOWABILITY* and *DISCRETION*.

Intuitively, defendants have a weighty complaint about being found against if they determinately did not violate the legal norm in question; call this the *No violation!* complaint. However, if the adjudication of borderline cases involves (legally sanctioned) strong retroactivity, then such a complaint will—as a general matter—not be available in such cases. On the assumption that a *No violation!* complaint is indeed unavailable in borderline cases—supported both by intuition and by actual legal practice—positing strong retroactivity thus seems provide a reasonable explanation for this “datum,” or in any case preserves consistency with it. Perhaps the issue I am pointing out is less consequential than I make it out to be?

The answer, I think, is no. Although the resolution of borderline cases does involve retroactivity, there is—at least intuitively—a clear difference between a case in which a person is, say, convicted of an offense that was determinately not part of the law at the time of action (strong retroactivity) and a case in which a person is convicted for doing something that is borderline illegal (weak retroactivity). All else equal, the former seems worse than the latter, which is—I submit—explained in part by a difference in

34. Thanks to Andrei Marmor for pressing me to address this point.

the reasons involved. In the former case, the defendant did not—at the time of her action—act contrary to the reasons provided by the norm on the basis of which she is being convicted, while in the latter case it is indeterminate whether she has done so. Thus, intuitively, DISCRETION employs a weak notion of retroactivity/discretion, rather than a strong one. So although we *could* understand DISCRETION in terms of strong retroactivity and doing so would suffice to explain why the *No violation!* complaint is not available to defendants in cases of vagueness, weak retroactivity seems to me to provide a more straightforward explanation of the same phenomenon.

It is worth pointing out, though, that there are also problems downstream from a strong analysis of DISCRETION. Actual legal systems neither confer unlimited lawmaking power to judges nor contain norms that specifically permit strongly retroactive judicial decisions in borderline cases.<sup>35</sup> It might of course be suggested that, in each actual legal system, borderline cases somehow form a subset of the set of cases in which such decisions are permitted; perhaps, some more general—but still limited—norm guarantees strong judicial discretion in borderline cases. It is difficult to see, however, how this holds true of any actual legal system, let alone of all (which would be required in order to block the puzzle this way). Consider, for example, the US legal system. Under what general norm permitting judges to apply law retroactively could borderline cases—as a unified class—conceivably fall?

#### D. Expanding the Domain of Legal Reasons (Conservatively)?

It might be suggested that the way I have construed the domain of legal reasons is too limited. Just as we can—at least on some accounts—sometimes convert partial beliefs based on evidential probabilities into outright beliefs about chances/epistemic probabilities,<sup>36</sup> or approximate truth into full truth,<sup>37</sup> we can also sometimes convert borderline cases into knowable propositions: if P\* is a borderline case of the predicate “dwelling,” for example, then presumably it is true and, perhaps even knowable, that this is so.<sup>38</sup> Maybe this fact is appropriately regarded as a legal reason, and one that B acts contrary to in some significant sense if she does not pay council tax?

Some authors put (something like) this idea in terms of subjects being “put on notice”—that if they do not do as directed in borderline circumstances, then they may be subject to official action, such as prosecution

35. One fundamental reason why no actual legal system contains a specific norm of this sort is that it is actually part of the problem, as I have described it, that officials do not generally conceive of judicial decisions in borderline cases as blatantly/determinately retroactive, because they do not generally think that this is the normative effect that KNOWABILITY has.

36. See John Hawthorne & Jason Stanley, *Knowledge and Action*, 105 J. PHIL. 571, 581–585 (2008).

37. See Alexander Bird, *What Is Scientific Progress?*, 41 NOÛS 64, 76–78 (2007).

38. Thanks to Finnur Dellsén for valuable discussion about this issue.

and punishment.<sup>39</sup> Usually, albeit not universally, this provides subjects with a reason to err on the side of caution in such circumstances, a reason that seems grounded in the law in a way that is not totally dissimilar to the case of “normal” legal reasons—especially if they conform better to the reasons that apply to them by so erring. Maybe the solution is to say that acting contrary to a “notice-based” reason in a borderline case is tantamount to it being indeterminate whether one acts contrary to a legal reason provided by the relevant legal norm, in which case it is indeterminate whether one violates it?

I am not prepared to rule this option out entirely, but I think we have good reason to think that it doesn’t work. I do think that—at least in some cases—the fact that it is indeterminate whether  $p$  is a reason for  $A$  to  $\phi$ . I want to resist, however, the idea that it is a *legal* reason.

Now, for practical reasoning in general, I think it is true that the fact *that it is indeterminate whether  $p$*  is one of the things that can “kick in” in circumstances in which the case for a course of action depends on  $p$ . Say, for example, that there is a drizzle outside, such that it is indeterminate whether it is raining. If it were clearly raining, that fact would be a reason for me to bring an umbrella, should I venture outside. But since it is indeterminate whether it is raining, I cannot—assuming a standard account of vagueness—know that it is raining (or, for that matter, know that it is not raining). Hence, if we assume KNOWABILITY, *that it is raining* is not a reason for me to bring an umbrella. Usually, however, this is of limited practical significance. The reason is that other considerations tend to kick in and determine what we ought to do. The same facts about the world that make it the case that it is indeterminate whether it is raining can presumably constitute a reason for me to bring an umbrella, for example because they make it sufficiently probable that it will rain. The general thought here is that, in determining what to do, we can usually replace things that we don’t or can’t know with things that we do know—and this includes cases of vagueness.

I say “usually” because there do seem to be cases of practical reasoning in which it is much less clear that any other considerations kick in in this way. Consider, for example, the difference between the father who promises to take his family to the park on Sunday if it isn’t raining and the person who randomly promises to give a stranger ten bucks if the next guy she sees is bald. In the former case, if on Sunday it is indeterminate whether it is raining, presumably the father ought still to take his family to the park. Granted, *that it is not raining* may not be a reason for him to do so, but other reasons kick in, making the overall balance of reasons count in favor of taking them to the park (although it could go either way, depending on the circumstances). In the random promise case, if the next guy that the promisor

39. See, e.g., Scott Soames, *Toward a Theory of Legal Interpretation*, 6 N.Y.U. J.L. & LIBERTY 231 (2011); Jeremy Waldron, *Vagueness and the Guidance of Action*, in PHILOSOPHICAL FOUNDATIONS OF LANGUAGE IN THE LAW 58 (Andrei Marmor & Scott Soames eds., 2011).



sees is borderline bald, then *that the guy is bald* is not a reason for her to give a stranger ten bucks (assuming KNOWABILITY). In this case, however, it is far from clear that any other reasons kick in. For one thing, no stranger reasonably expects to get ten bucks.

Suppose, however, that there were in fact a reason for the promisor to err on the side of caution and give the stranger the money, perhaps due to threat of violence, or something of that sort. In this case, it seems to me that even if we counted the fact that it is indeterminate whether the man is bald as being among the promisor's reasons, it would be very odd to call it a *promissory* reason. We may have all sorts of reasons "downstream" from having made a promise, but only some of them are properly labeled promissory. The same, I think, applies to the case of legal reasons. People may have reasons to err on the side of caution when it comes to borderline cases of illegal conduct, or even reasons to run the risk, but I hesitate to call these *legal* reasons. For one thing, to keep the comparison to promissory reasons, it would seem odd to say that if the promisor did not give the stranger the money, she would be breaking her promise. And it doesn't seem any more palatable to say that, in virtue of whatever is the relevant downstream reason, it is indeterminate whether she would be breaking it. Similarly, it seems very strange to say that if a subject acts contrary to a notice-based reason associated with the enactment of a legal norm, it is thereby indeterminate whether she violates that norm.

It is worth adding that even if we were to expand the domain of legal reasons in some relevant way, the inconsistency between VAGUE LEGAL REASONS and KNOWABILITY (given IGNORANCE OF LANGUAGE) would still remain. And, further, if what I said in the paragraph above is correct, we would also be stuck with the *No violation!* complaint, mentioned in the previous subsection, which would be problematic, since neither intuition nor legal practice indicates that such a weighty complaint is available in cases of vagueness. All in all, then, this suggestion does not seem to provide a robust way out.

#### E. Whose Reasons?

Another suggestion might be to apply something like the previous response to officials, rather than subjects. Perhaps we can "disconnect" DISCRETION from the reasons that apply to subjects and focus instead on the reasons that apply to officials in borderline cases, one of which is *that it is indeterminate whether A's case is one in which the operative predicate applies?* The worry here is that I have been focusing on the wrong set of agents.<sup>40</sup>

The best support for a move like this presumably comes from Hart's collective acceptance theory of norms. On such an approach, it is at least conceptually possible that the powers of discretion conferred on judges via the rule(s) of recognition are constitutively independent of anything to do with the subject's reasons. Although contrary to what I think is actually the case,

40. Thanks to Kenneth Ehrenberg for this suggestion.

it could be claimed that the accepted norm in relation to adjudication is just that in borderline cases judges have certain normative powers, and that borderline cases are defined simply by reference to the operative predicate, in which case judicial discretion is independent of the normative situation of the subject. DISCRETION, instead, would be a function of VAGUE LAW, rather than of VAGUE LEGAL REASONS (as the puzzle assumes).

While I think it is not so counterintuitive to take propositions about indeterminacy (as opposed to indeterminate propositions) to be among the legal reasons that officials have in borderline cases, it is quite problematic to disconnect DISCRETION in this way. In addition to the factual doubts already mentioned, it would require giving up the most straightforward and robust basis for the notion of judicial discretion in borderline cases, i.e., that it comes about as a consequence of it being indeterminate whether a norm has been violated. Severing this link thus seems both factually problematic and theoretically costly.

Further, on a noneliminative collective acceptance theory of norms like Hart's,<sup>41</sup> the relevant group's acceptance can be partly—but crucially—characterized in terms of a “normative perspective,” which can help relieve an equally significant part of the tension without requiring us to revise our understanding of the basis of judicial discretion in the way suggested. The bulk of my argument on this point is set out in Section II.H below, but to anticipate the basis of it, consider how the acceptance in question grounds a “point of view.” With respect to judicial discretion, the collective acceptance of the relevant officials for example makes it the case that “from the point of view of the law,” in cases where it is indeterminate whether the subject has violated the norm in question, judges are—assuming no other norms conclusively apply—not bound to decide the case one way rather than the other, but may reason either way. A similar analysis applies to first-order norms that apply to subjects: if the law prohibits, say,  $\phi$ -ing, then, for any subject S, S has—from the point of view of the law—a reason not to  $\phi$ . And the principle that a subject violates a legal norm only if she acts contrary to a legal reason provided by that norm here presumably takes the following form: a subject S violates a legal norm N iff, from the perspective of the law, A acts contrary to a reason provided by N. As I explain in Section II.H, we can—in a significant enough sense—reject KNOWABILITY for such perspectival reasons (despite accepting the thesis for “real” reasons), in which case the above suggestion to disconnect judicial discretion from the subject's normative situation becomes entirely unnecessary. Of course, none of the concerns I have raised here by itself constitutes a knockdown argument against this suggestion, but together I think they give us good reason to seek other ways to address the puzzle.

41. For the purposes of this paper, I will take eliminative individualism about social norms to be a nonstarter, since I do not see how such a view could really account for judicial discretion—on such an account, there is really no such norm.

Of course, some might respond by embracing the link between judicial discretion and the notion of violating a norm, but stipulate that for a subject *S* to violate a legal norm *just is* for *S*'s actions to fail to satisfy the operative predicate in the relevant circumstances (rather than to have acted contrary to a legal reason provided by the relevant norm, however conceived). This is, in effect, to reject a reasons-based account of legal norms altogether. On such a suggestion, the very foundation of the puzzle is removed, as the move from VAGUE LAW to VAGUE LEGAL REASONS is blocked, in which case the inconsistency with KNOWABILITY is not generated. The *No violation!* complaint would of course also go away.

I don't want to claim that this more radical move would not offer a possible way out, but giving up a reasons-based account of legal norms is ultimately too costly, I think. Not only does this strategy lead to an impoverished analysis of judicial discretion—so impoverished that one would have to reject not just the puzzle but the subject matter itself—it would actually no longer make sense to talk about legal *norms*, since normative notions like *obligation*, *power*, *permission*, and so on, are arguably best analyzed in terms of reasons.<sup>42</sup> Indeed, norms are—I think—at least partly) constituted by the reasons to which they give rise. This is of course only a problem if we want to be able to talk sincerely about law as a normative system, but then—at least absent a wholesale rejection of normativity<sup>43</sup>—I do not see how we could take law not to be a normative phenomenon. So this “solution” comes at a price that I, for one, am not prepared to pay.<sup>44</sup>

#### F. Does the Argument Equivocate?

So far, I have—in line with most standard theories of vagueness—been taking it for granted that the unknowability involved in IGNORANCE OF LANGUAGE and KNOWABILITY is of the same kind. But what if the former employs a specific vagueness-related notion of ignorance while the latter employs an ordinary factual notion? In that case, the argument equivocates and the problematic inference doesn't go through.

The idea behind this strategy is that we can distinguish between two notions of unknowability—one factual and one vagueness-related—and that these have different implications vis-à-vis the epistemic accessibility constraint on what it takes to be a reason. On this suggestion, the factual

42. See, in addition to RAZ, *ENGAGING REASON*, *supra* note 21, e.g., JEAN HAMPTON, *THE AUTHORITY OF REASON* (1998); T.M. SCANLON, *WHAT WE OWE EACH OTHER* (1998); MARK SCHROEDER, *SLAVES OF THE PASSIONS* (2007); JOHN SKORUPSKI, *THE DOMAIN OF REASONS* (2010); and DEREK PARFIT, *ON WHAT MATTERS* (2011). It is important to note, though, that my argument does not strictly require a “reasons-first” approach to normativity (according to which reasons are metaphysically or conceptually/explanatorily prior to other normative notions); for my purposes here, it suffices that (in a robust set of cases) the existence of a norm entails the existence of “corresponding” reasons, in the relevant sense, which is a significantly weaker assumption, and one compatible with other approaches to the analysis of normativity.

43. See, arguably, Brian Leiter, *Normativity for Naturalists*, 25 *PHIL. ISSUES* 64 (2015).

44. Thanks to an anonymous referee for prompting a clarification here.

unknowability of  $p$  entails that  $p$  is determinately not a reason, but the same does not go for vagueness-related unknowability; instead, in case  $p$  is a borderline proposition, the respective ignorance entails indeterminacy regarding  $p$ 's being a reason.<sup>45</sup>

This is a promising strategy, I think, but it comes with very deep commitments in philosophy of language and metaethics (as well as, arguably, in epistemology and philosophy of mind). The main issue is that it would require a theory of vagueness on which vagueness-related ignorance is analytically distinct from ordinary, factual ignorance, or else we can simply reproduce the problem. Of course, all standard theories of vagueness—even Timothy Williamson's epistemicism<sup>46</sup>—take indeterminacy to be a special source of ignorance, but it is far less clear that there is supposed to be any difference in the actual ignorance produced. In both cases, the agent does not know whether  $p$ , which makes it difficult to claim equivocation on the notion of ignorance at play. It seems clear, then, that in order to make something like this suggestion work, we will have to appeal to something beyond mere ignorance to properly distinguish the factual cases and the vagueness-related cases.

I think we can make some headway here by reminding ourselves that KNOWABILITY is a cognitive constraint generated by Raz's Nexus, as a result of what it takes to respond to something using one's rational powers. Perhaps, when  $p$  is indeterminate, the Nexus has different implications for what is cognitively required in order for an agent to rationally respond to  $p$ ? This is not an entirely unexplored question, but—as Robert Williams notes—there is very little consensus about what the cognitive role of indeterminacy is, i.e., about what attitude(s) one ought to take toward  $p$ , when (one knows, or believes, that)  $p$  is indeterminate.<sup>47</sup> Still, I think the bare-bones idea provides enough grounds to distinguish between the following two principles (using "X" as a placeholder for the relevant attitude(s)), which is all we need to make the strategy under consideration palatable:

If  $p$  is determinate, then A is able to respond to  $p$  using her rational powers only if A can come to know that  $p$

If  $p$  is indeterminate, then A is able to respond to  $p$  using her rational powers only if A can come to X that  $p$

The strategy here, then, is to claim not that the argument equivocates, but—rather—that it mistakenly applies KNOWABILITY to cases of vagueness, when in fact it only applies in determinate cases.

45. Thanks to Stefan Sciaraffa for valuable discussion about this strategy.

46. WILLIAMSON, *supra* note 15.

47. Robert Williams, *Decision Making Under Indeterminacy*, 14 PHILOSOPHERS' IMPRINT 1 (2014).

This strategy resonates particularly well with theories on which vagueness is to be explained primarily in psychological terms, in part—but crucially—by elucidating what *sui generis* attitude we ought to take toward borderline cases.<sup>48</sup> Note, though, that the suggestion here is importantly—albeit subtly—different from simply embracing one of the relevant psychological accounts already on offer. Granted, on such accounts, IGNORANCE OF LANGUAGE is not derivable from the distinctive attitude(s) one ought to take toward *p*, but these views still generally accept this thesis, either as independently motivated or as a consequence of other features of the particular account,<sup>49</sup> and so we are stuck with the puzzle unless we reject or modify at least one of the other theses involved; hence, the focus here on KNOWABILITY and the above suggestion to restrict it to determinate cases. I hope to have shown, however, that paying attention to the attitude(s) one ought to take toward borderline cases can help sever the link between KNOWABILITY and IGNORANCE OF LANGUAGE, by prompting us to take a closer look at Raz's Nexus. But it is by no means an easy way out—in order for it to work, a lot more needs to be said about the cognitive role of indeterminacy.

### G. Rejecting IGNORANCE OF LANGUAGE, or at Least Its Consequences

One of the assumptions of both the puzzle and the previous suggestions—and of much of the debate about vagueness, both in philosophy and in law—is that there are in fact clear borderline cases. Over the past two decades or so, however, this assumption has been met with increasing resistance and some authors now favor views that seem to better respect the actual difficulty of forming reliable judgments about borderline cases.<sup>50</sup> As it is still a widely held assumption, I go along with it for the sake of argument, but it is worth pointing out what happens once we give it up.

Not everyone will agree, and—as we will see below—some authors prefer to provide other grounds for the same result, but rejecting the assumption that there are any clear borderline cases arguably limits the effect of IGNORANCE OF LANGUAGE in the following way: if, for the relevant range of cases, it is never clear whether it is indeterminate whether a predicate applies to an object (relative to a context of utterance), then it is indeterminate whether it is unknowable whether the predicate applies to that object (in that context). IGNORANCE OF LANGUAGE might of course still hold as a

48. See, e.g., Harty Field, *Indeterminacy, Degree of Belief, and Excluded Middle*, 34 *NOÛS* 1 (2000); Stephen Schiffer, *Vagueness and Partial Belief*, 10 *PHIL. ISSUES* 220 (2000); Crispin Wright, *On Being in a Quandary: Relativism, Vagueness, Logical Revisionism*, 60 *MIND* 45 (2001); Robert Williams, *Degree Supervaluational Logic*, 4 *REV. SYMBOLIC LOGIC* 130 (2011). On Field's account, for example, one ought to both refuse to accept *p* and refuse to accept not-*p*, while on Williams's proposal, one ought to assign to *p* a credence of 0.5. To illustrate the stark difference between these two views, as an indicator of the lack of consensus regarding the cognitive role of indeterminacy, it suffices to point out that the former rejects classical probability for vague propositions, while the latter does not.

49. See note 15, *supra*.

50. See, e.g., DIANA RAFFMAN, *UNRULY WORDS* (2014), at 67.

penumbral connection, but it would no longer have the problematic consequences that fuel the puzzle.<sup>51</sup>

Some authors have also proposed grounds for rejecting not only the consequences of IGNORANCE OF LANGUAGE but also the thesis itself, although these accounts remain nonstandard. For example, David Barnett joins Cian Dorr in arguing that it might be *indeterminate* whether a fully informed person knows *whether p*, in case *p* is indeterminate.<sup>52</sup> There are some differences to their arguments, but basically they agree that it might be the case that if it is indeterminate whether *p* and some suitably specified cognizer knows all the relevant facts, then it is indeterminate whether she knows *whether p*. There will simply be no fact of the matter regarding which epistemic state she is in, given full knowledge of the relevant facts.

To be sure, any suggestion to revise IGNORANCE OF LANGUAGE will be highly controversial, and the thesis is still representative of any standard account of vagueness, but if the revised version of it turns out to capture the correct understanding of the relationship between vagueness and (un)knowability, then the tension between KNOWABILITY and DISCRETION is indeed relieved. Consequently, this option would allow us to retain all of the puzzle's legal and metaethical theses, but our commitments in philosophy of language would run pretty deep.

Still, it is probably the "simplest" strategy on offer. And—fittingly—such accounts would, I think, entail that it is indeterminate whether borderline propositions can guide action. That is, if it is indeterminate whether *p*, then it is—on these views—indeterminate whether *p* is knowable to the agent; consequently, it is indeterminate whether she would be able to respond to *p* using her rational powers.

#### H. Rejecting, or at Least Modifying, KNOWABILITY for Legal Reasons?

As alluded to in Section II.E, there is another promising way to relieve the tension between KNOWABILITY and DISCRETION: to accept KNOWABILITY for "real" reasons but reject it for "pure" legal reasons. Although restricting KNOWABILITY and revising IGNORANCE OF LANGUAGE are in some sense easier options, I think we may learn more about the normativity of law by examining whether, or at least in what way, KNOWABILITY applies to legal reasons.

As Scott Shapiro has usefully pointed out—drawing heavily on Raz's influential notion of a *detached legal statement*—we can distinguish between (at least) two senses of "legal reason."<sup>53</sup> On one reading, "legal" functions as an adjective—on the other, as a qualifier. On the adjectival reading, a legal reason to  $\phi$  is a real reason to  $\phi$ , the existence of which depends on

51. Thanks to an anonymous referee for helpful comments about this option.

52. David Barnett, *Does Vagueness Exclude Knowledge*, 82 PHIL. & PHENOMENOLOGICAL RSCH. 22 (2011); Cian Dorr, *Vagueness Without Ignorance*, 17 PHIL. PERSPS. 83 (2003).

53. Scott Shapiro, *What Is the Internal Point of View?*, 75 FORDHAM L. REV. 1157 (2006); see RAZ, PRACTICAL REASON, *supra* note 21, at 171–177; RAZ, *supra* note 6, at 153–157. Raz himself drew heavily on HANS Kelsen, *PURE THEORY OF LAW* (Max Knight trans., 1960/1967).

the operations of legal institutions (suitably understood). A real reason here is just a fact or consideration that in fact counts in favor of some action (subject to any further existence conditions, such as knowability). On the qualified reading, one has a legal reason to  $\phi$ —roughly—just in case, from the point of view of the law, one has a real reason to  $\phi$  (at least partly in virtue of the law having directed one to  $\phi$ ). Given the distinction between these two understandings of legal reasons, it makes sense to ask whether KNOWABILITY applies equally to both. I think a very good case can be made that it applies—if it applies—to the former, but not the latter.

Now, on the adjectival reading, legal reasons are real reasons that people have in virtue of the operations of legal institutions. Accordingly, if KNOWABILITY applies to real reasons, then it applies to legal reasons (in this sense). On the qualified reading, however, legal reasons *purport* to be real reasons, and so the way in which KNOWABILITY applies—if it applies—is less clear. It seems reasonable to say, though, that if legal reasons in a robust enough sense purport or aim to be real reasons, then the epistemic constraint on what it takes to *be* a real reason is a standard of evaluation for legal reasons, rather than an existence condition—much like, say, an effort to create something of a kind *K* will be subject to evaluation at least partly by reference to *K*'s essential features (if *K* has such features). If this is the sense in which KNOWABILITY applies to legal reasons on the qualified reading, then *p*'s being unknowable doesn't outright disqualify *p* from being a legal reason, although it will be a defective one, lacking—as it does—one of the essential features of real reasons (i.e., being knowable).

On this understanding of KNOWABILITY, there is no conflict with DISCRETION, and the tension is relieved. Nevertheless, there is an uneasy asymmetry here. If what I have said is correct, we can evaluate vague legislation partly by reference to the essential features of real reasons—in the case of KNOWABILITY, for example, the extent to which individual cases will measure up will vary depending on what it would take to come to know them given one's best efforts. But if we also assume a standard account of vagueness, then the failure—as we have seen—will be total in borderline cases, which seems to predict that we should have a fairly negative view of judicial discretion, more negative than we in fact seem to have. Why would it make sense for officials to have discretion in borderline cases, if such cases constitute a complete failure, *vis-à-vis* what the law purports to do (i.e., to provide real reasons for action)?

The answer is far from clear to me. To be sure, there is an inclination within most legal systems to treat cases of vagueness “cautiously”—to decide in favor of the defendant, or at least not to be neutral about the fact that they are borderline cases. The common law rule of lenity in criminal law is a prime example. And, sometimes, this is justified in terms of a failure to guide action (adequate notice, due process, etc.). However, as commentators point out, the rule of lenity is less common in criminal law than one might expect, and when it is present, it is often not applied as frequently as

one would think. In addition, it is extremely rare to find similar rules outside criminal law. This suggests a significantly less negative attitude regarding judicial discretion than seems predicted by retaining KNOWABILITY for real reasons (and thus as a standard of evaluation for legal reasons, in the qualified sense) along with a standard account of vagueness (on which borderline propositions are unknowable).

This is by no means a clear problem, and it may well be that there are many possible explanations of this apparent asymmetry that concern none of the theses that we have been examining here. There tend to be so many considerations at work in deciding legal cases that it is difficult to make firm judgments about how different considerations are weighted against each other, whether in particular cases or—especially—across entire legal systems. It is worth considering, however, whether the asymmetry goes away if we combine the suggestion here with the suggestion from the previous subsection. Perhaps we should both revise KNOWABILITY and adopt an account of vagueness on which it is indeterminate whether borderline propositions are knowable? At the very least, what this does—at least ideally—is to align our evaluative attitude toward judicial discretion with the extent to which the law succeeds in its aspirations to provide subjects with real reasons for action.

Here is how the story would go on such an account. Say that it is indeterminate whether some object *o* (broadly construed) is a borderline case of a predicate *P* occurring in a legal provision *L* and that, consequently, it is indeterminate whether, from the point of view of the law, *Po* is a real reason for *A* to  $\phi$ . In case *A* doesn't  $\phi$ , this entails that it is indeterminate whether *A* violates the legal norm to which *L* corresponds. If *A*'s case goes to court, then the relevant officials are—assuming no other legal norms conclusively apply to the case (such as a rule of lenity)—neither bound to decide against *A* nor to decide in her favor. As a result, they may (in some relevant sense) reason either way.

Furthermore, it will—on this account—be indeterminate how such cases measure up to the knowability standard (against which we can partly, but crucially, evaluate legislative directives). This marks a difference, for example, in how we should weigh the law's action-guiding failure in such cases, which is then weighed against other considerations in determining the outcome of the case. I submit that, holding everything else fixed, if we adopt a standard account of vagueness, we ought to give considerably more weight to action-guiding considerations in such deliberations than if we adopt a nonstandard account (i.e., one on which indeterminacy entails indeterminate knowability), due to their different implications for the (im)possibility of knowledge. The latter, I think, is more in line with our intuitions about proper practical reasoning; as concerns action-guidance, there is, intuitively, something worse about retroactive and secret law than vague law. This suggests that although either strategy—i.e., revising IGNORANCE OF LANGUAGE or rejecting KNOWABILITY for legal reasons—will make the puzzle go away, combining the two might be a particularly attractive option.



### III. CONCLUSION

I hope to have made it clear how the (at least *prima facie*) tension between KNOWABILITY and DISCRETION arises and that this should motivate us to say much more about the way vagueness affects the reasons for action that (the enactment of) legal norms give rise to. I also hope to have indicated which of the possible ways out I take to be most fruitful. I haven't been able to consider all the options, of course, but have tried to make palatable three strategies: first, to argue that the argument mistakenly applies KNOWABILITY to cases of vagueness, when in fact it only applies to determinate cases; second, to modify IGNORANCE OF LANGUAGE in such a way that indeterminacy entails indeterminate knowability/unknowability, rather than determinate unknowability; and, third, to reject KNOWABILITY for legal reasons (in the qualified sense) but accept it for real reasons. These suggestions, however, all generate deep commitments in philosophy of language and metaethics. And they do so not in virtue of any specific features of the relevant proposals but rather because they all involve rejecting one (or more) of the eminently plausible theses on which the puzzle relies. Thus, it seems safe to generalize and say that although there must of course be a solution to the puzzle, there is no easy way out.