Legal Theory, **19** (2013), 242–281.
© Cambridge University Press 2014 0361-6843/13 $15.00 + 00
doi:10.1017/S1352325213000165

DWORKIN'S INTERPRETIVISM AND THE PRAGMATICS OF LEGAL DISPUTES*

David Plunkett
Dartmouth College
Timothy Sundell
University of Kentucky

One of Ronald Dworkin's most distinctive claims in legal philosophy is that LAW is an *interpretative concept*, a special kind of concept whose correct application depends neither on fixed criteria nor on an instance-identifying decision procedure but rather on the normative or evaluative facts that best justify the total set of practices in which that concept is used. The main argument that Dworkin gives for interpretivism about some concept—LAW, among many others—is a disagreement-based argument. We argue here that Dworkin's disagreement-based argument relies on a mistaken premise about the nature of disagreement. We propose an alternative analysis of the type of dispute—what we call "seeming variation cases"—that Dworkin uses to motivate the idea of interpretative concepts. We begin by observing that genuine disagreements can be expressed via a range of linguistic mechanisms, many of which do not require that speakers literally assert and deny one and the same proposition. We focus in particular on what we call "metalinguistic negotiations," disputes in which speakers do not express the same concepts by their words but rather negotiate how words should be used and thereby negotiate which of a range of competing concepts should be used in that context. We claim that this view has quite general theoretical advantages over Dworkin's interpretivism about seeming variation cases and about the relevant class of legal disputes in particular. This paper thus has two interlocking goals: (1) to undermine one of Dworkin's core arguments for interpretivism, and (2) to provide the foundations for a noninterpretivist alternative account of an important class of legal disputes.

*Thanks to David Braddon-Mitchell, Alexi Burgess, Ronald Dworkin, Andy Egan, Triantafyllos Gkouvas, Sandy Goldberg, Mark Greenberg, Scott Hershovitz, AJ Julius, Frank Jackson, Chris Kennedy, Brian Leiter, Peter Ludlow, Kate Manne, Tristram McPherson, Eliot Michaelson, Daniel Nolan, George Pavlakos, Karl Schafer, François Schroeter, Laura Schroeter, Scott Shapiro, Seana Shiffrin, Alex Silk, Dale Smith, Nicos Stavrapolous, Kevin Toh, Kevin Walton, and Pekka Väyrynen. Versions of this paper were presented to audiences at the Julius Stone Institute of Jurisprudence at the University of Sydney School of Law, the Legal Theory Workshop at the University of Melbourne School of Law, the Centre for Law and Cosmopolitan Values at the University of Antwerp, the Jurisprudence Discussion Group at Oxford University, and the 2012 Inland Northwest Philosophy Conference on Pragmatism, Law, and Language. Thanks to everyone who participated in those sessions for their helpful questions and comments.

https://doi.org/10.1017/S1352325213000165 Published online by Cambridge University Press

I. INTRODUCTION

One of Ronald Dworkin's most distinctive claims in legal philosophy is that LAW is an interpretative concept.[1] According to Dworkin, *interpretative concepts* are a special kind of concept whose correct application depends not on fixed criteria or an instance-identifying decision procedure but rather on the normative or evaluative facts that best justify the total set of practices in which that concept is used.[2] Dworkin argues for the thesis that LAW is such a concept by appealing to the intelligibility of a certain kind of disagreement. This disagreement-based argument is central to Dworkin's discussion of legal philosophy in *Justice for Hedgehogs* and is arguably at the core of Dworkin's earlier "semantic sting" argument in *Law's Empire*.

Dworkin's disagreement-based argument can be glossed as follows. He begins by drawing our attention to a distinction between two types of legal dispute. Both types of dispute seem to be about *what the law is* in a given jurisdiction at a given time. In the first type, it seems that the parties agree on the conditions for something's *being a law*—at least the conditions for the relevant jurisdiction at the relevant time—but disagree about whether those conditions are met in the case at hand.[3] Such disputes evince a disagreement over a purely descriptive question—what the relevant empirical facts actually are.

In the second type, however, this model does not seem to capture what is going on. In the second type of legal dispute, legal actors such as lawyers or judges express different views about "what the law is" *despite* full agreement about all of the relevant empirical facts. In this second type of dispute, these different views are held and expressed consistently; they are not expressed merely as one-off judgments, possible artifacts of abnormal conditions. Neither does it seem that the parties to the dispute are *unaware* that there is background agreement on the relevant empirical facts. In these disputes, the speakers thus have systematically different dispositions about what to count as "the law," despite full, mutually known agreement about all of the relevant empirical facts. So, it would seem, the disagreement expressed in such a dispute cannot be about the empirical facts. It has to be about something else, something deeper and more fundamental. We call disputes of

1. RONALD DWORKIN, LAW'S EMPIRE (1986); and RONALD DWORKIN, JUSTICE FOR HEDGEHOGS (2011). In this paper we use small capitals to designate concepts. On this convention, *see* ERIC MARGOLIS & STEPHEN LAURENCE, CONCEPTS: CORE READINGS (1999).

2. *See* DWORKIN, HEDGEHOGS, *supra* note 1.

3. Or, to put things in a more holistic register (one that is perhaps better to use in developing a Dworkinian view of law): in some legal disputes, it seems that the parties agree on the conditions for something's *being a part of the law*—at least those conditions for the relevant jurisdiction at the relevant time—but disagree about whether those conditions are met in the case at hand. For expository purposes, in this paper we often slide freely between more or less holistic terminology. Nothing of substance hangs on this in the context of this paper. For a discussion of some of the things that *do* hang on it more generally, *see* Mark Greenberg, *The Standard Picture and Its Discontents, in* 1 OXFORD STUDIES IN PHILOSOPHY OF LAW (L Green & B Leiter eds., 2011).

https://doi.org/10.1017/S1352325213000165 Published online by Cambridge University Press

this second type—disputes where speakers persist in the dispute about what to count as "the law" despite full agreement on the relevant empirical facts and awareness of that background agreement—*bedrock legal disputes.* (In this terminology, we diverge from Dworkin, for reasons described below.)

We grant here that Dworkin is right that bedrock legal disputes sometimes occur in legal practice, or at least that we can imagine legal actors engaging in them. So, what should we think is going on in bedrock legal disputes? Dworkin observes that if we take the meaning of the word "law" (or, similarly, "the law") to be determined by the rules that legal actors use in applying it, then we are forced to the conclusion that the parties to bedrock legal disputes mean different things by the word and thus that they do not genuinely disagree with one another. He writes:

> If two lawyers are actually following *different* rules in using the word "law," using different factual criteria to decide when a proposition of law is true or false, then each must mean something different from the other when he says what the law is. . . . So the two judges are not really disagreeing about anything when one denies and the other asserts this proposition. They are only talking past one another. Their arguments are pointless in the most trivial and irritating way, like an argument about banks when one person has in mind savings banks and the other riverbanks.[4]

Dworkin thinks that it is clear, however, that not all bedrock legal disputes should be analyzed as "pointless in the most trivial and irritating way." He therefore reasons that however differently the speakers may use the relevant terms, it must be the case that they nevertheless mean the same things by them.

Dworkin claims on this basis that the meaning of "law" *cannot* be determined by some set of common rules speakers follow in applying it. Thus, he concludes that the concept the term "law" picks out must be different from ordinary concepts. In particular, the content of this concept cannot consist in fixed application criteria. Dworkin offers his notion of *interpretive* concepts as an alternative. If the correct application of the concept LAW (and thus the correct usage of the word "law"), is fixed by the normative or evaluative facts that best justify a set of practices, then the speakers can share the *same* concept (the concept LAW) after all, despite their persistent and systematic disagreement about what is to count as "the law."

This leads Dworkin to the following account of what is going on in bedrock legal disputes. Speakers in bedrock legal disputes are (1) expressing a disagreement about what sorts of facts need to obtain in order for something to count as part of the law; and (2) able to express this disagreement in virtue of sharing—despite their significant differences—a common concept LAW. Dworkin calls disputes of this sort *theoretical disagreements*—and thus his

4. DWORKIN, EMPIRE, *supra* note 1, at 43–44.

https://doi.org/10.1017/S1352325213000165 Published online by Cambridge University Press

theory of what we are calling "bedrock legal disputes" is that they are best analyzed as theoretical disagreements.[5] Crucially, in a theoretical disagreement it is possible for speakers to mean the same thing by a word—that is, express the same concept by that word—despite widespread divergence in patterns of usage and views about when to apply that word. For Dworkin, interpretative concepts explain how this can happen. And thus they are a crucial part of his positive account of what is going on in bedrock legal disputes.

So far we are glossing Dworkin's disagreement-based argument as one that is primarily about the concept LAW. It is precisely here, after all, that Dworkin's disagreement-based argument for interpretivism has had the most impact within legal philosophy. However, in *Justice for Hedgehogs*, Dworkin makes clear that this disagreement-based argument is meant to apply much more generally. In *Hedgehogs*, Dworkin endorses interpretivism about a wide range of legal, moral, and political concepts, including FREEDOM, DEMOCRACY, and EQUALITY. Indeed, Dworkin also does so for many seemingly descriptive concepts, including BALDNESS, BOOK, and LION—at least when they are employed under certain special circumstances. The core of the disagreement-based argument itself is no different for each of these different concepts. In this paper, we address Dworkin's more general argument schema, the disagreement-based argument type that he applies not just to LAW, but to FREEDOM, DEMOCRACY, JUSTICE, and, under the right conditions, even BALDNESS, BOOK, and LION.

What is it that unites the full range of disputes to which Dworkin applies his argument? In other words, what must a dispute involving some expression other than "law" or "the law" have in common with bedrock legal disputes in order for a parallel argument to be run? Roughly, the relevant class corresponds to those disputes where, even in the presence of "very great and entirely intractable differences of opinion about instances,"[6] we nevertheless are reluctant to conclude that the disagreement consists in a mere talking past. In such a case, Dworkin writes:

> the question always remains, in spite of even very radical disagreement, whether the pattern of that disagreement is better explained by the hypothesis that those who disagree share a single interpretive concept and disagree about its character, or by the alternative hypothesis that the disagreement is illusory

5. As this way of putting things makes clear, we do not use the term "theoretical disagreement" as a description of the second kind of dispute that Dworkin draws our attention to. This is because calling such disputes "theoretical disagreements" already involves a theory of what is *in fact* going on in these cases—or at least it does so if one follows how Dworkin himself uses the term "theoretical disagreement." We want to reject Dworkin's account of what is going on the relevant disputes, so we need a more theoretically neutral term for picking them out. We have introduced our term "bedrock legal disputes" in order to do so. One could, of course, just stipulate that the term "theoretical disagreement" is more theoretically neutral than we take it to be. But we think that is less helpful and more likely to lead to confusion for readers already familiar with Dworkin's more theoretically loaded use of the term.

6. Dworkin, HEDGEHOGS, *supra* note 1, at 161.

https://doi.org/10.1017/S1352325213000165 Published online by Cambridge University Press

> like our agreement to meet at the bank. [Where by "bank" one speaker means savings bank, while the other means side of a river.][7]

In the case of many of our most important moral and legal disputes, Dworkin argues that the conclusion that the disagreement is illusory is unpalatable. To accept it, "we would have to accept what seems ludicrous: that the most fervent and passionate of our political arguments are just silly misunderstandings."[8]

For Dworkin, members of the relevant class of disputes are characterized by a couple of special features. First, participants in such a dispute persist in their dispute even as they come to be aware not only of all of the relevant nonlinguistic facts but even of the significant and systematic divergence in their applications of the relevant term. Second, despite those differences in usage of the relevant terms, we theorists (along with the speakers themselves) take the dispute to express a genuine disagreement. In this class of disputes, normal considerations—factors such as patterns of usage and dispositions—would arguably point us toward the presence of semantic variation, toward the conclusion that the speakers mean different things by their words. But the fact that in these cases the speakers nevertheless genuinely disagree with one another seemingly prevents us from going that route. For ease of exposition, it will be helpful to have a label for the class of disputes that meet this description, something Dworkin himself does not provide. Let us refer to disputes of this kind as *seeming conceptual variation cases*, or *seeming variation cases* for short.

We think that there is a serious flaw in Dworkin's disagreement-based argument for interpretivism about the concepts that speakers employ in seeming variation cases. Dworkin's argument rests, we claim, on a crucial and largely unargued-for premise about disagreement. The premise is that the best way to explain how an exchange between two speakers serves to express a genuine disagreement is, in almost all circumstances, to suppose that those speakers mean the same thing—that is, express the same concepts—with the words they use in that exchange. Indeed—as can be seen in the passages we quote above—much of what Dworkin writes suggests that he thinks the positing of shared word meanings is not merely the *best* explanation of why an exchange reflects a genuine disagreement but indeed the *only possible* explanation. We argue that Dworkin's position here is a mistake in both its stronger and weaker formulations. And we claim that there is a better way of accounting for what is going on in seeming variation cases, a way that does not require positing a new kind of concept and which is more smoothly integrated into an overall account of thought and talk. Our goal in this paper is to develop this alternative to Dworkin's account of seeming variation cases—including cases of bedrock legal disputes—and to show

7. *Id.*
8. *Id.* at 162.

https://doi.org/10.1017/S1352325213000165 Published online by Cambridge University Press

how it underscores a central weakness in Dworkin's disagreement-based argument.

Our account rests on an appreciation of the variety of ways in which the disagreements between speakers can be expressed in linguistic exchanges. To see how this is so, first distinguish two things: (1) linguistic exchanges that seem to evince or reflect a disagreement—what we call *disputes*; and (2) *disagreements* themselves, which we take to involve a rational conflict in the attitudes (for example, the beliefs) actually held by the subjects, irrespective of whether those subjects enter into a linguistic exchange with one another. Within *disputes,* we make a further distinction. Some disputes are *canonical*: they reflect disagreements over the truth or correctness of literally expressed semantic content.[9] In canonical disputes there is a single proposition (or proposition-like entity) that is part of the literally expressed content asserted by one speaker and denied by the other. In these unexotic exchanges, the fact that speakers mean the same things by the words they use is a crucial component in an explanation of how there could be a coherent subject matter under dispute and how, therefore, their disagreement could be genuine. Not all genuine disagreements, however, are expressed via canonical disputes. Sometimes the disagreement concerns information that happens to be communicated pragmatically rather than semantically. Such disputes are *noncanonical*; they do not reflect a disagreement about the literally expressed semantic content. Yet these disagreements have the potential to be "genuine" in any sense that matters.

Our proposal is that many seeming variation cases express, noncanonically, a disagreement about which among some candidate class of *non* interpretative concepts should be deployed in the context at hand. This is what we call a disagreement in *conceptual ethics*: it concerns a normative question about thought and talk and, in particular, about concept use.[10] Some

9. Why "truth *or correctness*" rather than simply "truth"? Why "content" rather than the more specific "proposition"? We intend for our notion of canonical disagreement to leave room for expressivist accounts *as* accounts in terms of canonical disagreement. Given that certain expressivist views eschew talk of truth and propositions, we want to make clear that these notions are not crucial for articulating what canonicalness amounts to. Dworkin's arguments are parallel in important ways to some of the standard arguments for expressivist accounts of normative concepts (such as in Allan Gibbard, Wise Choices, Apt Feelings: A Theory of Normative Judgment (1990); and Allan Gibbard, Thinking How to Live (2003)). While we do not press the point in this paper, we take the general arguments we present in this paper to have import for expressivist accounts of normative thought and talk. We pursue this line further in other work and, having noted the connection, set aside the more general terminology and often speak loosely in terms of propositions and their truth for the remainder of this paper. For more discussion of these issues, *see* David Plunkett & Timothy Sundell, *Disagreement and the Semantics of Normative and Evaluative Terms*, 13 Philosophers' Imprint 1–37 (2013). It is also not entirely clear what differentiates Dworkin's interpretivism about concepts from recent forms of expressivism such as that developed in Gibbard, Thinking. We do not take any stand on that important issue in this paper.

10. The terminology of "conceptual ethics" is taken from Alexis Burgess & David Plunkett, *Conceptual Ethics I*, Phil. Compass (2013); and Alexis Burgess & David Plunkett, *Conceptual Ethics II*, Phil. Compass (2013). These papers distinguish between a broad use of the term "conceptual ethics," where it is used to refer to normative and evaluative issues about thought and talk,

https://doi.org/10.1017/S1352325213000165 Published online by Cambridge University Press

disagreements in conceptual ethics happen overtly, of course. In such cases, the literal semantic content of the relevant expressions concerns a topic within conceptual ethics. Other disagreements in conceptual ethics, however, are expressed via pragmatic mechanisms and are thus expressed in *noncanonical* disputes. In this paper, we focus on one mechanism in particular that can be used to express a disagreement in conceptual ethics: *metalinguistic* uses of a term, cases in which a speaker uses a term (rather than mentions it) in order to communicate something about the proper usage of that very term. We call the specific kind of dispute that we are interested in—one in which metalinguistic usages of a term serve to communicate a disagreement in *conceptual ethics*—a *metalinguistic negotiation*. We claim that this "metalinguistic analysis" of seeming variation cases—including bedrock legal disputes—has decisive advantages over Dworkin's own interpretivist analysis.

Our paper thus has two goals. First, we aim to demonstrate that Dworkin's argument for interpretivism about a range of concepts—including LAW—is unsound in virtue of its reliance on a false premise. Second, we aim to advance a positive alternative to his analysis of seeming variation cases, an alternative that does justice to the considerations Dworkin raises in support of interpretivism but that has further theoretical advantages. As we emphasize in Section III, we do not here argue that the metalinguistic analysis is preferable to *other* alternatives to Dworkin's interpretivism. We do not, for example, argue here that our analysis is preferable to expressivist treatments of bedrock legal disputes or in other seeming variation cases. Nevertheless, we hope to make an initial positive case for the metalinguistic analysis. And, given our first goal, if it turns out that important arguments for these other alternatives have in common with Dworkin a reliance on faulty assumptions about disagreement—as they do in at least some cases (for example, disagreement-based arguments for expressivist accounts of bedrock legal disputes)—then this lends further support to our metalinguistic account of seeming variation cases, relative to the competition. For purposes of this work, however, we limit ourselves to demonstrating (1) the flaw in Dworkin's argument; and (2) the strength of the metalinguistic analysis over Dworkin's interpretivist account in particular.

and a narrower use of the term, where it concerns issues about concept choice and usage in particular. The metalinguistic negotiations that we focus on—the ones that matter most for our argument against Dworkin—revolve around conceptual ethics in the narrower sense, since they revolve around the normative issue of which of a range of competing concepts should be used. But as our longer discussion of metalinguistic negotiations in Section IV underscores, not all metalinguistic negotiations involve that particular issue in conceptual ethics—some, for instance, concern how a threshold should be set for a given relative gradable adjective in the context. See our discussion of metalinguistic negotiations in Plunkett & Sundell, *supra* note 9.

https://doi.org/10.1017/S1352325213000165 Published online by Cambridge University Press

II. INTERPRETIVISM IN *JUSTICE FOR HEDGEHOGS*

In this section, we start by clarifying Dworkin's understanding of concepts and of interpretative concepts in particular. This allows us to get clearer on some important features of Dworkin's disagreement-based argument—including its scope—and puts us in a better position then to evaluate it.

In *Justice for Hedgehogs*, Dworkin argues for interpretivism about a wide range of legal, political, and moral concepts. His main argument for doing so is as follows: noninterpretive concepts are inadequate to the task of understanding seeming variation cases. What does he think noninterpretive concepts are like? In *Hedgehogs*, Dworkin divides noninterpretative concepts into two categories: *criterial* concepts and *natural-kind* concepts. He writes, "[s]ome of our concepts are *criterial* in this sense: we share the concept when, but only so far as, we use the same criteria in identifying instances."[11] He illustrates by pointing to the case of BOOK: "We share some concepts because we agree, except in cases we all regard as borderline, about what criteria to use in identifying examples. We mainly agree about how many books there are on a table, for example, because we use the same tests in answering that question."[12]

Dworkin describes the other type of noninterpretive concept—natural-kind concepts—by starting with the idea of natural kinds themselves. "Natural kinds are things that have a fixed identity in nature, such as a chemical compound or an animal species."[13] In turn, he understands natural-kind *concepts* in terms of their ability to pick out the relevant natural kind, irrespective of the particular test that is used to do so: "people share a natural-kind concept when they use that concept to refer to the same natural kind."[14]

Dworkin holds that unlike criterial concepts, natural-kind concepts can be shared by speakers who employ different criteria in identifying instances. However, he claims that they must nevertheless be ready to accept a common decision procedure. He writes:

> You and I disagree, say, about whether an animal we encounter in Piccadilly is a lion, and it turns out that I identify lions by their size and shape and you only by what you believe to be their distinctive behavior [for example, roaring rather than talking]. . . . You and I assume that "lion" names a distinct biological kind and that the beast we met is a lion if it has a lion's biological essence, whatever that is, whether or not it meets the criteria either of us normally uses to identify lions. If you understand DNA, and if tests showed that the creature we saw had the DNA of a lion, you would likely change your opinion to recognize talking lions. Criterial concepts do not work that way:

11. DWORKIN, HEDGEHOGS, *supra* note 1, at 158.
12. *Id.* at 6.
13. *Id.* at 159.
14. *Id.*

https://doi.org/10.1017/S1352325213000165 Published online by Cambridge University Press

> nothing you discovered about the molecular structure of my copy of *Moby Dick* could convince you that it was not a book.[15]

Thus natural kinds allow for a greater degree of variation between speakers before they can no longer count as sharing a concept. Provided we would in the end accept a common essence-revealing decision procedure such as a DNA test, we count as sharing a concept such as LION despite our distinct criteria for identifying instances.

Dworkin claims that neither criterial concepts nor natural-kind concepts, however, can serve as the basis for a satisfying analysis of the disputes that interest him in *Hedgehogs*: seeming variation cases and bedrock legal disputes among them. Consider criterial concepts. If a concept such as JUSTICE were criterial, then speakers who did not apply the same criteria for identifying instances could not be said to have the same concept in mind when they used the word "justice." Yet when a utilitarian and a Kantian, for example, engage in an dispute involving the expression "justice," it is clear, Dworkin suggests, that they do not apply the same criteria for identifying instances.[16] Supposing JUSTICE to be criterial, then, those speakers do not share the relevant concept and thus they are not genuinely disagreeing.

Switching to the example of DEMOCRACY, Dworkin writes, "if you and I mean something entirely different by 'democracy,' then our discussion about whether [for example] democracy requires that citizens have an equal stake is pointless: we are simply talking past one another."[17] When it comes to criterial concepts, "it is the identity of our criteria that makes disagreement genuine when it is genuine."[18] In seeming variation cases, speakers do *not* share criteria for identifying instances, for example, of what counts as "justice" or "democracy." Thus, if JUSTICE or DEMOCRACY were criterial, then seeming variation cases involving the terms "justice" or "democracy" would not express genuine disagreements. But some seeming variation cases involving the terms "justice" or "democracy" do express genuine disagreements. So, Dworkin concludes, JUSTICE and DEMOCRACY are not criterial.

Dworkin's story is slightly more complicated, but much the same in spirit, when it comes to natural-kind concepts. Dworkin holds that unlike with criterial concepts, "people can refer to the same natural kind even when they use, and know they do, different criteria to identify instances."[19] As noted above, however, Dworkin holds that natural-kind concepts have in common with criterial concepts that "people do not share a concept of either kind unless they would accept a decisive test—a kind of decision procedure—for finally deciding when to apply the concept (except in cases

15. *Id*
16. *Id.* at 162.
17. *Id.* at 6. It should be noted that the view that "citizens have an equal stake" is Dworkin's own view about democracy in *id.* However, the details of the specific debate that Dworkin is engaged in about what democracy requires are not relevant for our discussion here.
18. *Id.* at 159.
19. *Id.* at 160.

https://doi.org/10.1017/S1352325213000165 Published online by Cambridge University Press

they agree are marginal)."[20] Because of this, Dworkin claims that natural-kind concepts are no better suited than criterial concepts to play a role in an analysis of seeming variation cases about DEMOCRACY, JUSTICE, LAW, and the like.

In bedrock legal disputes, for example, speakers do not merely fail to apply the same criteria in identifying instances falling under the concept LAW. They do not even accept a common decisive test for how to apply some purportedly shared concept. When all the facts are in about, for example, the statute books and judicial decisions, there is no equivalent to a DNA test that could force parties to a bedrock legal dispute finally to converge on their opinion of what the law is. Precisely the distinguishing feature of bedrock legal disputes is that even as speakers become mutually aware of the relevant descriptive facts and even as it becomes clear that they do not employ the same criteria or accept the same decision procedures for applying the word "law," they nevertheless persist in their disagreement. If LAW were a natural-kind concept, then in virtue of their failure to accept a common decision procedure, the parties to bedrock legal disputes would not share the concept and thus would not genuinely disagree. But parties to at least some bedrock legal disputes do genuinely disagree. So, Dworkin concludes, LAW is not a natural-kind concept.[21]

Since parties to seeming variation cases do not share a criterial concept of DEMOCRACY or JUSTICE or LAW and since they do not share a natural-kind concept of DEMOCRACY or JUSTICE or LAW, then either there is some further, alternative kind of concept or there is no shared concept at the center of their disputes. Dworkin, as noted, finds the second option unpalatable. So he offers his own alternative, *the interpretive concept*, as a solution:

> We must therefore recognize that we share some of our concepts, including the political concepts, in a different way: they function for us as *interpretive* concepts. We share them because we share social practices and experiences in which these concepts figure. We take the concepts to describe values, but we disagree, sometimes to a marked degree, about what these values are and how they should be expressed. We disagree because we interpret the practices we share rather differently: we hold somewhat different theories about which values best justify what we accept as central or paradigm features of that practice. That structure makes our conceptual disagreements about liberty, equality, and the rest genuine.[22]

By understanding concepts like DEMOCRACY, JUSTICE, or LAW in this way, we can take speakers to share those concepts even when they differ not only in the criteria they employ for identifying instances but in the decision procedures they endorse for settling difficult cases about applications of the

20. *Id.*
21. *See id.* at 168–170.
22. *Id.* at 6.

https://doi.org/10.1017/S1352325213000165 Published online by Cambridge University Press

concept. Because they employ the concept as part of a shared practice, and because speakers share an interpretative concept just insofar as they share certain social practices in which that concept figures, speakers can share the concept despite these other differences. Because they share the concept, their disputes expressing thoughts involving this concept can evince genuine disagreement.

This type of solution immediately raises a difficult issue. How can we tell which concepts are interpretive and which are not? Although he focuses on certain central legal, political, and moral concepts, Dworkin's suggestion seems to be that almost *any* concept can be interpretive (or "function interpretively") if it appears in the right kind of dispute. Indeed, he takes the question of which category a given concept falls into (on a particular occasion of use) to itself be an interpretive question. Criterial or natural-kind concepts can, on occasion, be employed in contexts in which the criteria or decision procedures for identifying instances are up for grabs. In such circumstances, it may be that the correct application of the concepts in question is best interpreted as depending on the normative or evaluative facts justifying the relevant set of practices.

Dworkin asks us to imagine, for example, a statute granting a tax exemption to bald men:

> This silly statute would convert the question of baldness into a genuine interpretive issue: officials, lawyers, and judges would have to contrive some highly artificial definition of baldness (not necessarily a hair-counting definition) by asking which such definition would make most political sense of the exemption.[23]

In such a situation, two lawyers, for example, might disagree about whether a particular person was bald, and they might persist in their dispute even as they come to be aware of all of the relevant facts about that person and about the differences in their willingness to apply the term "bald." And yet, under special circumstances like these, their disagreement would be entirely genuine. The dispute thus qualifies as a seeming variation case, and Dworkin's argument can apply even to the ordinarily noninterpretive concept BALD.

The same goes for natural-kind concepts. Imagine, for example, that some kind of radiation can change an animal's cells, "not just randomly but into those produced by the DNA of a different animal."[24] In particular, if the radiation is applied to a lion, it systematically transforms the lion's cells into the cells that would have been produced had the animal's parents been house cats. Two decision procedures for determining the application of the relevant natural-kind concepts are now in conflict, and we can imagine a

23. *Id.* at 164.
24. *Id.* at 165.

https://doi.org/10.1017/S1352325213000165 Published online by Cambridge University Press

corresponding dispute between two biologists. One biologist might argue that the animal is a lion, even post-transformation, in virtue of being the offspring of lions. The other biologist might argue that the animal is a cat, in virtue of presently having the DNA of a cat:

> If their arguments then took the form of a debate about the most useful way of continuing the established classificatory practices of zoology, we might well say that the concept of a lion had become for a time something like an interpretive rather than a natural-kind concept.[25]

Dworkin's case is constructed in just such a way as to make the biologists' dispute qualify as a seeming variation case. That it is possible to construct such a case even with natural-kind terms shows that they, too, can pass through what Dworkin calls an "interpretive phase."[26]

Dworkin's disagreement-based argument can thus be seen to have a strikingly large range of applications. It is tempting, but not quite right, to say that certain special expressions—the legal, moral, or political expressions, for example—express interpretive concepts, while most other expressions do not. In fact, for Dworkin, almost any expression can express an interpretive concept, provided it is used in the context of a seeming variation case. Or alternatively, almost any concept—criterial, natural kind, or presumably others if there are others—can pass through an interpretive phase if that concept is employed in a seeming variation case. Whether such an analysis is actually right for a particular concept under a particular set of circumstances is, according to Dworkin, an interpretive question itself. But if the dispute looks to qualify as a seeming variation case, and if we are indeed reluctant to describe the parties as engaged in a "silly misunderstanding," then, according to Dworkin, there are powerful reasons for interpreting the concept as interpretive.

III. DWORKIN'S TAXONOMY OF CONCEPTS

One serious problem with Dworkin's disagreement-based argument for interpretative concepts is with the account of interpretative concepts as such. There are a number of important features of interpretative concepts that Dworkin highlights in his work. They are different from the sorts of *prosaic* concepts we use to understand the meanings of many of our ordinary descriptive terms, such as "table" and "book." They are not natural-kind concepts; their correct application depends on normative and evaluative matters in a way that is different from those concepts. We can have

25. *Id.*
26. These examples are, in different ways, rather unrealistic. But that is not crucial to them. Dworkin offers the example of tax policies concerning "books" (criterial) and scientific debates about Pluto's status as a "planet" (natural kind) as more realistic alternatives.

https://doi.org/10.1017/S1352325213000165 Published online by Cambridge University Press

mistaken views about what these concepts involve yet still correctly apply them. What the correct conditions are for applying those concepts can coherently be put into doubt by those yielding the concepts. And the content of these concepts is not given in individualist terms—that is, some form of anti-individualist content-externalism seems to be true of them. These are all important features of interpretative concepts, and yet it is not at all clear which of them Dworkin intends to be the *essential* properties of interpretative concepts. We think this is a significant issue for Dworkin, especially since many other philosophers grant that certain concepts have all (or at least many of) these features but do not take themselves to be interpretivists about those concepts at all. So what more precisely makes something a genuinely interpretative concept as opposed to not? Dworkin's answer is far from clear.

This ties into another serious problem for Dworkin's disagreement-based argument for interpretive concepts. As our discussion in the last section underscores, Dworkin has quite specific views about the nature of *noninterpretive* concepts in *Justice for Hedgehogs*, and those views play a prominent role in his argument there and elsewhere for interpretive concepts. Yet Dworkin's views about noninterpretive concepts—his notion of a *criterial* concept and his account of *natural-kind* concepts—are hardly consensus views. Dworkin's criterial concepts, for example, build in the existence of mechanical tests for proper application as well as reliable first-person access to those tests. Neither of these is widely thought to be a feature of ordinary concepts, including Dworkin's own examples, such as BOOK.

Dworkin's reliance on his own idiosyncratic theories of noninterpretative concepts is a serious problem for his disagreement-based argument for interpretive concepts. In arguing for the inadequacy of noninterpretive concepts, he canvasses only a tiny set of the available theories of those concepts. Expressivist theories of concepts, for example, are often motivated by disagreement-based arguments very much like Dworkin's and thus are built from the ground up to address the possibility of disagreement in cases much like those Dworkin considers.[27] Sophisticated versions of neodescriptivism and inferentialist theories likewise have resources for explaining how speakers who differ systematically in their usage might nevertheless share concepts.[28]

27. *See, e.g.*, GIBBARD, WISE CHOICES, *supra* note 9; and GIBBARD, THINKING, *supra* note 9. It is worth noting here that Kevin Toh, the philosopher who has done the most of anyone to develop a form of expressivism in the context of legal philosophy, draws on disagreement-based arguments to argue in favor of expressivism about what he calls "internal legal statements"—roughly, statements *of* law. *See* Kevin Toh, *Hart's Expressivism and His Benthamite Project*, 11 LEGAL THEORY 75–123 (2005); and Kevin Toh, *Legal Judgments as Plural Acceptances of Norms, in* OXFORD STUDIES IN PHILOSOPHY OF LAW (L. Green & B. Leiter, eds., 2011). In both papers Toh rightly underscores how expressivism about internal legal statements (whether of his form in particular or the nearby form that he attributes to Hart) can be motivated by much the same disagreement-based arguments that Dworkin gives for his interpretivism about concepts.

28. *See, e.g.*, MICHAEL SMITH, THE MORAL PROBLEM (1994); and RALPH WEDGWOOD, THE NATURE OF NORMATIVITY (2007).

https://doi.org/10.1017/S1352325213000165 Published online by Cambridge University Press

Seen in this light, Dworkin's argumentative strategy is deeply problematic: ordinary concepts (as Dworkin sees them) do not do the trick; natural-kind concepts (as Dworkin sees them) do not do the trick; so we need interpretive concepts. But why think that ordinary concepts (as Dworkin conceives them) and natural-kind concepts (as Dworkin conceives them) are the only other kinds of concepts there can be?[29] Maybe the problem is with how Dworkin sees noninterpretive concepts!

To reiterate: Dworkin has not given us a good theory of exactly what interpretive concepts are. And worse, Dworkin considers only a highly restricted set of theoretical resources for understanding what *noninterpretative* concepts are. These are both very serious problems. One might therefore be tempted to stop the discussion of Dworkin's argument there.

We are sympathetic to this response. However, there is an important sense in which stopping here does not do justice to Dworkin's disagreement-based argument. By focusing exclusively on his own theories of noninterpretive concepts, Dworkin can give the impression that it is features peculiar to his own theories of noninterpretative concepts that allow his disagreement-based argument to go forward. This, however, is not the case. To see this, consider a kind of straightforward neodescriptivism—a version, say, modeled on the view described in Frank Jackson's *From Metaphysics to Ethics*.[30] On such a view, the correct application of a concept is determined by the dispositions of the speaker to apply the term.[31] Nothing is required in the way of mechanical application tests, and there is no reason to think the speaker would have reliable first-person access to the application conditions of her concepts. Those application conditions are reflected in her usage and dispositions but not necessarily in her own beliefs about the concept. Such a view—suitably clarified and filled out—is well known among philosophers and applies happily to both categories of Dworkin's noninterpretive concepts.

By eliminating reference to mechanical tests, decision procedures, and any particularly strong form of first-person introspective epistemic access, this type of view eschews the most controversial features of Dworkin's own theories of noninterpretive concepts. Moreover, by rejecting mechanical tests for application and first-person epistemic access, our neodescriptivist has a much-augmented set of resources for explaining shared meaning in the presence of differing beliefs about meaning. For our neodescriptivist, speakers who accept differing criteria might nevertheless share a concept. After all, one (or both) of the speakers might simply be mistaken in her

29. For a helpful starting point on this literature, *see* MARGOLIS & LAURENCE, *supra* note 1.
30. FRANK JACKSON, FROM METAPHYSICS TO ETHICS: A DEFENCE OF CONCEPTUAL ANALYSIS (1998).
31. In his full developed story of thought and talk, Jackson himself also leaves room for Lewisian naturalness—of the rough sort outlined in David Lewis, *New Work for a Theory of Universals,* 61 AUSTRALASIAN J. PHIL. 343–377 (1983)—playing a role here. But we leave that out here for purposes of exposition, as does Jackson in much of JACKSON, *supra* note 30. For a good overview of some of the main arguments on behalf of giving Lewisian naturalness a role in the determination of the content of thought and talk, *see* THEODORE SIDER, WRITING THE BOOK OF THE WORLD (2012).

https://doi.org/10.1017/S1352325213000165 Published online by Cambridge University Press

beliefs about the application criteria that in fact correspond to her dispositions to use the term.

But, crucially, eliminating these things does not by itself equip the theory to handle seeming variation cases. By the very description of seeming variation cases, the speakers involved *do not merely* differ in the criteria or tests they explicitly *avow* for the applications of their terms. They also in fact apply those terms in very different ways. They do so sincerely and consistently and under a wide range of circumstances. In other words, they differ systematically in their *dispositions* to use the term. Every theory of meaning should allow that dispositions to apply a term provide at least prima facie evidence regarding the concept that term picks out. And the simple form of neodescriptivism considered here looks to such dispositions exclusively.

The point here is that while Dworkin can give the impression that his disagreement-based argument relies crucially on certain idiosyncratic views of ordinary concepts, it is not those views that generate the worry Dworkin's interpretivism is meant to address. Eliminating the most controversial aspects of those views does nothing by itself to eliminate seeming variation cases or to give rise to some account of them. In other words, Dworkin could, if he had had the inclination, have stated his argument in more theoretically neutral terms. Other theories of concepts may or may not have the resources to explain seeming variation cases or to explain them away. But rejecting the more controversial features of Dworkin's views about noninterpretative concepts is not sufficient to do so. Everyone has to have a story to tell.

Dworkin's own positive story here is, as noted, mysterious in certain crucial respects. But Dworkin's work in this area is widely read, and his positive account plays an important role in wider debates in philosophy of law, including debates about positivism and antipositivism. Since his argument for that account does not, despite appearances, rely on his perhaps faulty or easily dismissed views about ordinary concepts, it remains a worthwhile task to address the core of his argument head-on.

IV. GENUINE DISAGREEMENT AND METALINGUISTIC NEGOTIATION

In philosophical discussions of disagreement, a distinction is usually drawn between two very different types of "disputes," broadly construed. (Reminder: we use "dispute" to mean linguistic exchanges that seem to express a real disagreement between the speakers, whether in fact they do.) On the one hand, there are those disputes that evince "genuine disagreement" between the speakers, and on the other, there are those disputes that constitute what are often termed "mere talkings past."[32] In the first category,

32. For a helpful discussion of this distinction, *see* David J. Chalmers, *Verbal Disputes,* 120 PHIL. REV. 515–566 (2011).

https://doi.org/10.1017/S1352325213000165 Published online by Cambridge University Press

we find uncontroversial cases of "genuine disagreement," such as that in dialogue (1):

(1a) The cat is on the mat.
(1b) No, the cat is not on the mat.

In the second category, we find cases where, owing to some linguistic misunderstanding (hence "*talking past*"), the speakers fail to express a genuine disagreement. There is a range of linguistic misunderstandings that can underlie this second type of case, including undetected ambiguity, differences among mutually comprehensible dialects or idiolects, or even variation in the resolution of context-sensitive linguistic features. The exchanges in (2) to (4) should serve to give an idea of the relevant phenomenon.

(2a) Nell is at the $\text{bank}_{\text{[financial]}}$.
(2b) No, Nell is not at the $\text{bank}_{\text{[river]}}$.

(3a) Burgers come with $\text{chips}_{\text{[crisps]}}$.
(3b) No, burgers do not come $\text{chips}_{\text{[french fries]}}$.

(4a) Ivan is $\text{tall}_{\text{[for a philosopher]}}$.
(4b) No, Ivan is not $\text{tall}_{\text{[for a basketball player]}}$.

In each of (2) to (4), speakers fail, in one way or another, to mean the same things by their words (at least relative to their context).[33] In (2), the speakers "mean" different things because of a language-internal ambiguity; the sound-form [bank] can pick out (at least) two different concepts in English. In (3), the speakers "mean" different things because of a difference between American and British English. In American English, the sound-form [chips] picks out what a British speaker would call "crisps," while in British English the sound-form [chips] picks out what an American would call "french fries." Finally, in (4) the speakers "mean" different things in that they use the context-sensitive term "tall" to denote different properties. By setting the threshold for height differently, they refer to different things with the word, even though in a different sense it "means" the same thing

33. A quick but crucial point of clarification: Do the speakers in (2), for example, "mean different things by their words"? Or do they "use different words—homophonous, but distinct in virtue of differences in meaning"? We submit that this distinction is terminological (and—with a nod to our broader thesis—that it is not just terminological but, at least in this context, *merely* terminological). Disputes that we analyze below as metalinguistic—in which speakers negotiate how a word shall be used or which concept it shall be used to express—can, with no important theoretical changes, be redescribed as disputes in which speakers negotiate which of two competing, homophonous words shall be used. In fact, we think that in more general discussions of concepts and word meaning, the latter form of description (Which of the competing homophonous words should we choose?) is more apt. However, the former form of description (How should this very word be used in the present circumstances?) has certain expository advantages, and we stick to that way of talking for the duration of the paper.

https://doi.org/10.1017/S1352325213000165 Published online by Cambridge University Press

for them: something like "having a maximal degree of height greater than the contextually supplied threshold."[34]

One way or another, in each of these three cases, speakers use their words to *refer* to different things—different locations, different foodstuffs, different properties. Owing to that fact, they assert and deny propositions that in fact are entirely consistent. The seeming disagreements in (2) to (4) are thus not disagreements at all; or, if you prefer, they are disagreements, but not *genuine* disagreements.

Where does the real difference between (1), on the one hand, and (2) to (4), on the other, lie? The dispute in (1) reflects a disagreement that not only is genuine but is also, in our sense, *canonical*. That is, the literal semantic content of the speaker of (1a)'s assertion is logically inconsistent with the literal semantic content of the speaker of (1b)'s assertion.

The disputes in (2) to (4) reflect disagreements that, if they are disagreements at all, fail to be genuine. But not only do they fail to qualify as genuine, they also fail to qualify as canonical. That is, the content literally expressed by, for example, the speaker of (2a) is logically *consistent* with the content literally expressed by the speaker of (2b). Likewise for (3) and (4).

In considering the contrast between exchanges like (1) on the one hand and exchanges like (2) to (4) on the other, it is easy to think that this overlap between the categories of *genuineness* and *canonicalness* is no coincidence. One might naturally come to think that the dispute in (1) expresses a genuine disagreement precisely *in virtue* of being canonical, and that the disputes in (2) to (4) *fail* to express genuine disagreement precisely in virtue of being noncanonical. One might therefore come to think that, quite generally, if an exchange expresses a genuine disagreement, then it must be the case that the speakers literally express logically inconsistent propositions. (In other words, and simplifying a bit, in a genuine disagreement at least one speaker must say something false.)

That principle, however, cannot be right. The reason is that in virtue of its focus on literal semantic content, the principle artificially ties the question of whether a dispute is *genuine* to the question of the *linguistic mechanism* via which competing claims about that subject matter are transmitted. But, intuitively, questions about whether there is a stable subject of disagreement have *nothing to do* with the empirical linguistic question of whether competing claims about that subject happen to be communicated semantically or pragmatically. This intuition is precisely correct, and indeed there are many perfectly clear instances of disputes that express genuine disagreement but which are not canonical—in which the speakers do not literally express inconsistent propositions.

There is a wide range of quite unexotic disagreements that demonstrate this point, but to save time, we consider just one before moving on to the

34. *See* Chris Barker, *The Dynamics of Vagueness*, 25 LINGUISTICS & PHIL. 1–36 (2002).

https://doi.org/10.1017/S1352325213000165 Published online by Cambridge University Press

cases most relevant to Dworkin's argument. So consider the exchange in (5):

(5a) There are forty-nine states in the United States.
(5b) No, there are fifty states in the United States.

There is a widely held if not entirely uncontroversial view among linguists and philosophers of language according to which number words of the kind that appear in (5) literally express only a "lower-bound" reading. On such a view, an utterance of the expression in (5a) literally expresses the proposition that there are *at least* forty-nine states in the United States, while an utterance of the expression of (5b) literally expresses that the proposition there are *at least* fifty states in the United States. But of course those two propositions are consistent. Indeed, they are both true.

That there are *exactly* forty-nine states is communicated by the speaker of (5a), but it is part of the pragmatic upshot of the utterance in context, not the semantic content of the speaker's assertion. Nevertheless, we submit that no plausible account of genuineness could exclude disagreements like that expressed in (5). The speakers of (5a) and (5b) *take themselves* to be engaged in a substantive dispute. All of the usual linguistic markers of disagreement (such as the felicity of phrases like "nope," "nuh uh," and the like) are present. And of course, the question of how many states comprise the United States is a substantive one of great practical significance if little controversy. Indeed, the question of whether the dispute in (5) is *substantive* seems entirely orthogonal to the question of the linguistic mechanism by which the relevant information is communicated. This is hardly surprising: what matters to whether a dispute is substantive is its *topic*—Is it something addressed by both speakers? Is it worth arguing about? Is it plausibly resolvable? and so on—and not the linguistic means by which the competing claims are advanced.

A. Metalinguistic Disputes

We observe above that the strong "at least one falsehood" generalization drawn from the distinction between disputes like that in (1) and those in (2) to (4) is false. Disagreements over information communicated via implicature or other pragmatic mechanisms provide clear and unexotic counterexamples. However, they do not yet speak directly to the types of cases Dworkin addresses. After all, Dworkin argues that if we suppose criterialism to be true, we find ourselves committed to the claim that parties to theoretical disagreements *express different concepts* with their terms. He could happily grant that the dispute in (5) reflects a genuine disagreement because while the speakers do assert mutually consistent propositions, there is no actual variation in word meaning or concept choice. By contrast, in the cases

https://doi.org/10.1017/S1352325213000165 Published online by Cambridge University Press

Dworkin considers, the criterialist must conclude not just that the speakers each express true propositions but that they do so *because* they employ distinct concepts expressed with the same word. Even if we allow—based on examples like (5)—that some genuine disagreements are noncanonical, we might thus remain skeptical that a disagreement could be both genuine and noncanonical *in virtue of a difference in word meaning or concept choice*. We think such skepticism is unwarranted, however. Having opened the door to noncanonical, genuine disagreements in general, we see no reason not to allow for this particular variety of noncanonical genuine disagreement.

The point is easiest to see in the case of context-sensitive terminology, so we begin there and then move to cases more like the disagreements Dworkin addresses. Consider, in particular, linguistic expressions that are context-sensitive in virtue of being *gradable*.[35] This includes expressions such as "tall," "big," "cold," and so on that denote a specific property only once some parameter—a threshold for height, size, or temperature, say—has been settled by the conversational context or by the parties to the discourse in which the expression is used. If we know how that parameter is set—if, for example, we can hold fixed the contextual threshold for "coldness"—then expressions involving gradable adjectives such as "cold" can provide us with useful information about the temperature.

There is no reason at all, however, that things cannot work in precisely the reverse direction. If we know what the *temperature* is—if we can hold fixed the relevant heat-or-lack-thereof facts—then expressions involving gradable adjectives such as "cold" can provide us with useful information about the *context*. This latter kind of usage is described by Chris Barker, who calls it a "sharpening" or "metalinguistic" usage of a term.[36]

To imagine the type of usage we have in mind, suppose that Oscar has just arrived at a research station in Antarctica. Shivering from his walk to the shelter, he glances at the outdoor thermometer and asks his new colleague Jill, "Is this cold?" Jill regretfully replies, "Nope, I'm afraid this isn't cold." Oscar and Jill's exchange is not one in which they attempt to get clearer on what the temperature is. They already know what the temperature is. Rather, Jill helps Oscar augment his knowledge about certain features of their shared conversational context. In particular, she informs him that the local threshold for "coldness" is lower on the temperature scale than the mutually known current temperature. Why would Oscar care about a "merely linguistic" fact like this? *Because how we use words matters*. Oscar wants to be able to communicate smoothly with his new colleagues. But in addition

35. *See* Christopher Kennedy, *Vagueness and Grammar: The Semantics of Relative and Absolute Gradable Predicates*, 30 LINGUISTICS & PHIL. 1–45 (2007).

36. Barker, *supra* note 34. Barker does not extend his discussion to cases of disagreement over information communicated metalinguistically. His cases are expanded to include disagreements in Timothy Sundell, *Disagreements about Taste*, 155 PHIL. STUD. 267–288 (2011). Barker considers disagreements of this type in later work. *See* Chris Barker, *Negotiating Taste*, 56 INQUIRY 240–257 (2013).

https://doi.org/10.1017/S1352325213000165 Published online by Cambridge University Press

to that, learning how they use the word "cold" reveals useful information about the range of temperatures he can expect to be typical in his new environment.

Within this category of metalinguistic usage, there is a further distinction to be drawn.[37] In some cases of metalinguistic usage, it is quite natural to think that there are antecedently settled facts about what the linguistically relevant features of the conversational context are like, facts that are at least partially independent of the intentions—or at least the very local intentions—of the parties to a conversation. Oscar's and Jill's exchange may be like this: independently of their intentions or their usage in this particular context, there are facts about how "cold" is used at the research station. Oscar wants to know what those facts are, and that is what Jill informs him about. In such cases, it is natural to think of the speakers as exchanging information that is, in some (perhaps quite loose) sense, objective—information about what the context is actually like. If a disagreement should arise over that information, then the disagreement is a factual one about which of two or more competing characterizations of the shared conversational context is most accurate. However, not all cases of metalinguistic usage fit this profile.

Suppose that Oscar and Jill are not colleagues in an Antarctic research station but rather office mates in Chicago. Oscar is a "freezy cat"—he often feels chilly. Jill is not. While they look together at their shared, known-to-be-accurate thermostat, Oscar utters (6a) and Jill utters (6b).

(6a) It's cold in here.
(6b) No, it's not cold in here.

In this case, it is much less natural to think that there is some antecedently settled, objective fact of the matter about the contextually salient threshold for "coldness." Rather than advancing competing factual claims about some independently determined threshold, Oscar and Jill here *negotiate* over what that threshold shall be. Why would they consider it worth their time to engage in such a negotiation when they already agree what the temperature is? Why engage in a dispute over how to use a word? The answer is the same as before. It is worth engaging in such a dispute because it often matters a great deal how we use our words. For Oscar and Jill, as for many of us, the word "cold" plays a crucial functional role in our discourse and coordinated decision-making with respect to climate control. In particular, an agreement among all parties who share an indoor space that the space can aptly be described as "cold" leads swiftly to action—more specifically, to

37. While Barker does not make reference to the distinction we draw below, it is observed in both Andy Egan, *Disputing about Taste*, *in* DISAGREEMENT (R. Feldman & T. Warfield eds., 2010); and Sundell, *supra* note 36, though the discussions there both avoid making commitments with respect to it.

https://doi.org/10.1017/S1352325213000165 Published online by Cambridge University Press

thermostat-turning-up action. Thresholds matter. Why should Jill have to turn up the heat if the office cannot even be described as "cold"?

We think that metalinguistic usages of this latter type—usages where speakers do not simply exchange factual information about language but rather negotiate its appropriate use and the associated choice of concepts—are common. (At least as common as arguments about climate control!) Indeed we think such usages extend well beyond climate control to disagreements about what should count as "tall" for our basketball team, or "spicy" for our chili, or "rich" for our tax base. In any such case, the speakers may both assert true propositions, but they express those true propositions by virtue of the fact that they set the relevant contextual parameters in different ways.

Why are such exchanges perceived as disputes, when the speakers fail to assert inconsistent propositions? Because in addition to asserting those propositions—in fact *via* their assertion of those propositions—they also pragmatically *advocate* for the parameter settings by which those propositions are asserted. The claim that one parameter setting is preferable to some incompatible alternative parameter setting is very much the kind of thing over which two speakers can disagree. And given the right context—for example, a context where we must coordinate our thermostat adjustings, or our basketball player picks, or our chili preparation, or our progressive taxation brackets—such disagreements can be very much worth having. And to emphasize our earlier point, the question of whether the disagreement is worth having is entirely independent of whether the competing claims are advanced via semantic or pragmatic mechanisms.

In the cases of metalinguistic negotiation we are considering so far, the negotiations concern how to fix parameter settings for bits of context-sensitive terminology. In these cases, the speakers use their words to pick out the same function from a context to a property—in this sense, Oscar and Jill "mean" the same thing by "cold"—but they nevertheless use those terms to refer to different properties because they set the contextual parameters differently. It is in this latter sense that they do not mean the same things by their words: they use them, in the context, to denote different things.[38] So it goes with context-sensitive expressions such as gradable adjectives.

But metalinguistic negotiation is not confined to gradable adjectives or other context-sensitive expressions. It can concern even words that are seemingly quite fixed in their meaning. Consider the following example, discussed by Peter Ludlow in his paper "Cheap Contextualism."[39] The example concerns a debate that Ludlow heard on sports radio. The debate concerned the greatest athletes of the twentieth century and the question

38. The distinction we are drawing here is what Kaplan aims to capture with his distinction between an expression's *character*—its contextually invariant meaning—and its *content*—what it picks out relative to the context. *See* David Kaplan, *Demonstratives: An Essay on the Semantics, Logic, Metaphysics, and Epistemology of Demonstratives*, *in* THEMES FROM KAPLAN (Joseph Almog, John Perry & Howard Wettstein eds., 1989).

39. *See* Peter Ludlow, *Cheap Contextualism*, 18 PHIL. ISSUES 104–129 (2008).

https://doi.org/10.1017/S1352325213000165 Published online by Cambridge University Press

of whether that list should include the racehorse Secretariat. Simplifying a bit, we can imagine the following exchange as part of that debate:

(7a) Secretariat is an athlete.
(7b) No, Secretariat is not an athlete.

Unlike the cases of metalinguistic sharpening involving gradable adjectives, there is little reason to think that the relevant linguistic expression here—"athlete"—is semantically context-sensitive. But just as in those cases, there is also little reason to think that the dispute in (7) concerns straightforward factual matters about the topic at hand. After all, the facts about Secretariat's speed, strength, and so on are mutually known by the speakers of (7a) and (7b), just as the facts about the temperature are mutually known by Oscar and Jill. It thus seems much more natural to conclude—as Ludlow does—that the dispute in (7) in fact reflects a disagreement about which of two competing concepts is more appropriate to the conversation. Each speaker literally expresses a true proposition given the concept they in fact express with the word "athlete." But beyond that, they pragmatically advocate for the use of the concept in virtue of which they assert those propositions.[40] Thus, despite the relevant expression's not being context-sensitive in the traditional sense, this dispute is metalinguistic in precisely the same sense as the dispute between the office mates Oscar and Jill. In each case, speakers disagree about how the relevant expression should be used—what it should be used to refer to—under the circumstances.

Ludlow's Secretariat case is particularly vivid, but we see no reason to think that this type of metalinguistic negotiation is at all uncommon. Many of us are familiar with disputes about whether Missouri is in the "Midwest," or whether Pluto is a "planet," or whether the federal antidrug effort constitutes a "war." In each case, most of the relevant facts—the location of Missouri, the size and orbit of Pluto, the contents of the relevant antidrug policies—are mutually known among parties to the dispute. And yet it seems that the disputes are, or at least have the potential to be, genuine disagreements in any plausible sense. It may not matter very much which states we choose to include in the Midwest. But it can matter a great deal whether a policy is meant to address a social ill or advance our cause in a war. As

40. It might at this point be objected that certain forms of externalism—context externalism in the Oscar and Jill case, and content externalism in the Secretariat case—force an alternative analysis of all of these cases. As we argue at length in Plunkett & Sundell, *supra* note 9, only the most implausibly strong forms of externalism could force reanalysis of *all* seemingly metalinguistic disputes. (Cases of metalinguistic negotiation across mutually comprehensible dialects make the point nicely vivid.) Thus everyone is stuck with metalinguistic dispute, externalist or not; the question is merely how far the phenomenon extends. Once it is in the picture, we find it quite unnatural to think that it does not extend to cases of the kind we consider here. In any case, externalism by itself is not going to help Dworkin, because *many* philosophers (most of whom are not interpretivists) also accept externalism. Whatever exactly Dworkin's interpretivism amounts to, in order for it to be an interesting and unique view, it cannot simply be a way of stating that some form of externalism is true.

https://doi.org/10.1017/S1352325213000165 Published online by Cambridge University Press

in the case of Oscar's and Jill's debate about the "coldness" of their office, metalinguistic negotiations influence and advance more general processes of coordinated decision-making and action.

To see this point especially clearly, consider one last example. Suppose that two speakers engage in the dialogue in (8):

(8a) Waterboarding is torture.
(8b) No, waterboarding is not torture.[41]

Suppose further that the speaker of (8a) quite generally follows the U.N. General Assembly in defining torture as any act inflicting severe suffering, physical or mental, in order to obtain information or to punish, while the speaker of (8b) quite generally follows former U.S. Justice Department practice in defining torture as any such act inflicting pain rising to the level of death, organ failure, or the permanent impairment of a significant body function.[42]

Perhaps some form of externalism or expressivism—or even Dworkinian interpretivism—could have it come out that despite compelling evidence to the contrary, the speakers of (8a) and (8b) mean the same thing by "torture." But suppose that they do not. Even if we suppose that these speakers express different concepts with the word "torture" and thus that they each literally express a true proposition, we are in no way committed to the conclusion that the exchange in (8) is confused, nongenuine, or a *mere* talking past. As in the Secretariat case and as in the office thermostat case, speakers do more than literally assert the propositions that are the semantic contents of the expressions they utter. They also presuppose, connote, and implicate. And they also, by virtue of their choices of parameter settings and concepts, pragmatically advocate for those very parameter settings and concepts.

Does the dispute in (8) express a disagreement worth having? Of course. The matter of which interrogation techniques are aptly described as "torture" matters a great deal in our moral, ethical, and legal discourse and in the coordinated actions that result from those discourses. Just as the word "cold" plays a certain functional role in our debates and decisions about climate control—a role that is not built into the *semantics* of that term but which is an important feature of how it is used in these debates and decisions—so, too, does the word "torture" play a (much higher-stakes)

41. The example of torture is mentioned in Chalmers, *supra* note 32; and discussed in some detail in Timothy Sundell, *Disagreement, Error, and an Alternative to Reference Magnetism*, 90 AUSTRALASIAN J. PHIL. 743–759 (2011). Both make points similar to the one we make here. We also discuss this case in Plunkett & Sundell, *supra* note 9.

42. *See* United Nations Convention against Torture and Other Cruel, Inhuman or Degrading Treatment or Punishment, *opened for signature* Dec. 10, 1984, 1465 U.N.T.S. 85 (entered into force June 26, 1987); and Memorandum for Alberto R. Gonzales, Counsel to the President, Re: Standards of Conduct for Interrogation under 18 U.S.C. secs. 2340–2340A, U.S. Dep't Just. Off. Assistant Att'y Gen., August 1, 2002.

https://doi.org/10.1017/S1352325213000165 Published online by Cambridge University Press

role in our debates and decisions about the treatment of prisoners. Our choice among concepts competing to play that role is fraught with moral, ethical, and legal considerations and consequences.

If we choose to analyze the dispute in (8) as a metalinguistic negotiation, then we shall conclude that it concerns not the truth-value of some agreed-upon proposition but rather the question of which of two competing concepts is better. That of course raises a difficult question: What makes one concept better suited to the circumstances than another? We of course do not have the answer to that, but we do not need one for our argument to go through. Whatever substantive legal or metaethical theory one might point to to claim that one of (8a) or (8b) is *false* or that one of (8a) or (8b) is *incorrect* can be adapted to claim that one of the competing concepts employed in (8a) or (8b) is *worse* than the other. There are more ways to err than simply to express a false proposition. When we speak, we aim to succeed along a range of dimensions, and thus we are liable to fail along that same range of dimensions. Our claims are evaluable not simply by the truth of what we literally say but also by our choices of how to carve up the world with the concepts we employ.

So where, now, do we find ourselves vis-à-vis Dworkin's argument for the need for interpretative concepts? We argue above, first, that a dispute's expression of *genuine disagreement* is entirely consistent with both parties to that dispute literally expressing true propositions. (This follows from the fact that some genuine disagreements are expressed via noncanonical disputes.) Second, we argue that among those noncanonical disputes that express genuine disagreement, some are noncanonical in virtue of the parties to those disputes employing distinct concepts. (This follows from the fact that some genuine disagreements are expressed via noncanonical, *metalinguistic* disputes.) The fact that in a metalinguistic dispute the information at issue is communicated pragmatically and not semantically does not preclude the disagreement's being genuine, because the linguistic mechanism of communication is *always* independent of whether the disagreement is worth having. The fact that in a metalinguistic dispute, the information at issue is *about language* also does not preclude the disagreement's being genuine, because *how we use language matters.*

These conclusions fatally undermine Dworkin's quite general assumption that if parties to a dispute cannot be described as meaning the same things by their words in a given linguistic exchange (i.e., using their words to express the same concepts), then we are forced to conclude that their disagreement amounts to a "silly misunderstanding." Without that assumption, Dworkin's argument for the existence of interpretive concepts fails, as does the general argument *form* he uses to demonstrate that specific candidate concepts (e.g., LAW or DEMOCRACY) are indeed interpretive. Of course, it is still open to Dworkin (or anyone else) to argue that certain disputes (e.g., those involving "law" or "democracy") are best explained by thinking of the relevant concepts as interpretive. But one needs to do so *without* relying

https://doi.org/10.1017/S1352325213000165 Published online by Cambridge University Press

on the idea that the only disputes that express genuine disagreements are those in which the speakers mean the same things by the words that they use. Moreover, one needs to explicitly take on the idea that the disputes in question (e.g., bedrock legal disputes) are best analyzed in metalinguistic terms.

If what we have said in this section is on the right track, there is good reason to believe that significant parts of communication involve a metalinguistic dimension. Other things being equal, we should therefore expect that we also use metalinguistic negotiation when communicating in legal contexts. Indeed, if metalinguistic negotiation occurs widely across nonlegal contexts, it would be bizarre if it were not used in legal contexts. So, do we engage in metalinguistic negotiation in the particular part of legal thought and talk that Dworkin draws our attention to: bedrock legal disputes and legal seeming variation cases more generally? We argue that the answer is yes. In the next section, we argue for the following view: at least *some* bedrock legal disputes are best analyzed as metalinguistic negotiations, as are other important seeming variation cases that come up in law, including at least some cases that Dworkin and other legal philosophers refer to as "hard cases."

V. BEDROCK LEGAL DISPUTES AND HARD CASES

Is it open to Dworkin to allow that metalinguistic negotiations exist but to object that the type of analysis we have sketched here is implausible as applied to the specific type of legal disputes he has in mind? Of course. Perhaps metalinguistic negotiation is quite common indeed, but, for reasons we have failed to discuss so far, it is bad (perhaps absurd!) as a diagnosis of the class of legal disputes Dworkin is interested in—or, perhaps, for *any* important legal disputes. In order to make such an argument, one would need to identify what is different about the relevant legal disputes. But here it is worth recalling the breadth of Dworkin's target with his disagreement-based argument. He argues that not only are certain legal, moral, and political concepts interpretive but also concepts of other kinds, provided they appear in the context of a seeming variation case. Given the scope of that conclusion, it is not obvious what resources Dworkin could draw on in quarantining his preferred cases from more general considerations about linguistic communication. Whatever those resources might be, they are not given by the disagreement-based argument itself.

So what, then, should we make of the metalinguistic analysis in the legal context? Let us start with bedrock legal disputes—the sorts of cases that are at the core of Dworkin's discussion of interpretivism in the legal context. To see very generally how a metalinguistic analysis of a given bedrock legal dispute might go, consider an exchange of the form represented in (9):

https://doi.org/10.1017/S1352325213000165 Published online by Cambridge University Press

(9a) The law requires that ϕ.
(9b) No, the law requires that Ψ.

Suppose that ϕ and Ψ represent incompatible understandings of what the law is in a given jurisdiction (at a given time). To draw on Dworkin's own discussion of *TVA v Hill*—a case that he uses in *Law's Empire* to motivate his interpretivism about LAW—suppose further that (9a) is true only if the concept expressed by this speaker (let us call this concept LAW1) is such that what counts as "law" is fixed only by the meaning of the relevant legal texts, at least when that meaning is clear. And suppose that (9b) is true only if the concept expressed by this speaker (let us call this concept LAW2) requires that the relevant legal text be read in such a way that it "accords with some modicum of common sense and public weal."[43] Finally, suppose that the speaker of (9a) uses "law" quite generally in such a way that expressions containing it are true only if the relevant concept is sensitive only to the clear meaning of the relevant texts. And suppose that the speaker of (9b) uses "law" quite generally in such a way that common sense and public weal are relevant in addition to the meanings of the relevant texts. This type of dispute is precisely the kind of dispute that Dworkin thinks should be understood as what we are calling a "bedrock legal dispute." So, for now, let us suppose that Dworkin is right, and this *is* a bedrock legal dispute.

Suppose we want to analyze a dispute like this as metalinguistic. We then suppose that, just as their systematically differing usage would seem to indicate, the speakers of (9a) and (9b) do indeed "mean something entirely different"[44] by "law" and thus that they do indeed express distinct concepts by that word. This is precisely the view that Dworkin accuses noninterpretivists of being committed to. But it does not have the implication he claims it does. After all, on the metalinguistic analysis, in addition to asserting the propositions that are the literal content of the expressions they utter, the speakers of (9a) and (9b) implicate, presuppose, and connote a whole range of further information. Beyond that, of course, they pragmatically advocate for the concepts they choose to express with their words. When all the relevant descriptive facts on the ground are known, as they are in plausible cases of bedrock legal disputes, then—in precise parallel to the thermostat case, the Secretariat case, and the torture case—we can say that the parties to the dispute disagree about how "law" should be used and thus about which concept—LAW1 or LAW2—is better suited to the circumstances. That is a disagreement that is not only comprehensible but very much worth having.

How viable a view is this of a given bedrock legal dispute? First, recall the core of what bedrock legal disputes are as we are defining them: they are legal disputes in which speakers persist in engaging in a dispute about

43. This phrase is taken from Justice Powell's dissent in the case of TVA v. Hill, which Dworkin discusses in DWORKIN, EMPIRE, *supra* note 1.
44. DWORKIN, HEDGEHOGS, *supra* note 1, at 6.

https://doi.org/10.1017/S1352325213000165 Published online by Cambridge University Press

what to count as "the law" despite full agreement on the relevant empirical facts *and* awareness of that background agreement. Second, recall what Dworkin thinks must be the case with such disputes: they involve a disagreement about a topic that is worth arguing about. Our account vindicates this thought, and does so while claiming that the disputes in question *really are* bedrock legal disputes (as we have defined them). In other words, insofar as one analyzes a bedrock legal dispute as a metalinguistic negotiation, one can vindicate (1) the idea that the bedrock legal dispute involves a genuine disagreement, and (2) that this disagreement concerns something important that is well worth arguing about. And, moreover, one can do so while taking the dispute to be simply an instance of a *much more general* form of linguistic exchange, the existence of which is motivated by less philosophically loaded cases from outside the legal context. This, we submit, puts our metalinguistic account in a strong position. In other words, it lends support to the following claim: bedrock legal disputes are best analyzed as metalinguistic negotiations.

One point about this thesis is worth emphasizing. In accepting the claim that bedrock legal disputes are best analyzed as metalinguistic negotiations, one is *not* accepting the claim that every instance of something that *seems* to be a bedrock legal dispute is in fact a metalinguistic negotiation. One might argue that many seeming bedrock legal disputes reflect empirical disagreements expressed with concepts whose application conditions are surprisingly complex.[45] Or one might argue that most seemingly bedrock legal disputes in fact reflect opportunistic usages that disqualify them from being genuine bedrock legal disputes, in our sense.[46] Or, to take another example, one might argue that in the context of overall theory choice among accounts of legal thought and talk, bedrock legal disputes are simply not that important a class of disputes to worry about.[47]

Our view is entirely consistent with all of these ideas, and in fact we are sympathetic to a more general view in legal philosophy that combines them. It is important to keep in mind that the view on the table is this: *if a dispute is in fact a bedrock legal dispute, then there is very good reason to analyze it as a metalinguistic negotiation.* How many disputes in fact qualify is an open question, as is how important these disputes ultimately are for developing an account of legal thought and talk. Those questions are both beyond the scope of this paper.

What we now want to argue is this: our proposed metalinguistic account is naturally extended to many other parts of legal thought and talk beyond bedrock legal disputes. These include many so-called "hard cases." On this front, consider the following sorts of well-known cases from the law,

45. For a helpful illustration of an account of legal thought and talk that pursues this strategy, *see* SCOTT SHAPIRO, LEGALITY (2011).

46. *See* Brian Leiter, *Explaining Theoretical Disagreement,* 76 U. CHI. L. REV. 1215–1250 (2009).

47. *See id.*

https://doi.org/10.1017/S1352325213000165 Published online by Cambridge University Press

drawn from Scott Shapiro's discussion of hard cases in *Legality*.[48] When a drug trafficker trades a firearm for illegal drugs, is he *using* the gun? Are representatives who are elected through the use of electronic voting machines *chosen by written votes*? Does the use of a riding lawn mower violate the rule that there be no *vehicles* in the park? An individual takes an action that causes an event that unexpectedly causes an injury. Is the individual's action a *proximate* cause of the injury?

Disputes about such issues have exactly the sorts of features that make them ripe for analysis as a metalinguistic negotiation. The speakers involved are by and large mutually aware of all the relevant nonlinguistic facts (including, as in many cases of bedrock legal disputes, facts gleaned from the study of legal materials). And it is at least plausible that there is no antecedently settled matter of fact about the meaning. Everyone involved knows what a riding lawn mower is like, after all. (It's smaller than a car; it has an internal cumbustion engine; and so on.) Descriptive nonlinguistic facts are not at issue. The question is whether it should count as a "vehicle." But "vehicle" is clearly a vague term, and lawn mowers are a classic borderline case—so settled facts about meaning are not at issue either. What matters is a decision about how, for present purposes, to make the vague term precise. Thus, we claim, there is good reason to think that many such legal disputes are metalinguistic negotiations.[49]

Whether in fact any given legal dispute is a metalinguistic negotiation depends, of course, on further details of the individual cases. But well-known cases of the kind listed above are so similar in structure to core cases of metalinguistic negotiation that at least *some* of the relevant legal disputes are very likely metalinguistic negotiations. Thus, we think, awareness of the phenomenon of metalinguistic negotiation can help illuminate what is going on not only in bedrock legal disputes but in legal thought and talk more generally.

Despite these advantages, one might worry that a metalinguistic analysis of important parts of legal thought and talk falters as a theory because it conflicts with speakers' own intuitions about their activities. Very few legal practitioners—even those engaged in paradigm hard cases or what at least seem to be instances of bedrock legal disputes—take what they are doing to involve, in the first instance, matters about language and thought. Furthermore, when presented with the thesis that they are in fact engaged in a metalinguistic negotiation, many of them might resist the idea that this is a good account of their activities. So our account seems to involve attributing a type of serious *error* to legal actors. The worry is that this attribution of error counts significantly against the plausibility of our account of legal thought and talk.

48. Shapiro, *supra* note 45, at 235–237.

49. For further discussion of our proposed metalinguistic account of such cases, *see* David Plunkett & Timothy Sundell, *Antipositivist Arguments from Legal Thought and Talk: The Metalinguistic Response*, *in* Pragmatism, Law, and Language (G. Hubbs & D. Lind eds., 2014).

https://doi.org/10.1017/S1352325213000165 Published online by Cambridge University Press

In other work, we develop a detailed response to worries of this type.[50] But we sketch here the core of our argument. First, we take there to be quite general reasons to think that legal actors might not have perfectly reliable intuitions about some aspects of their own practice. Quite generally, practitioners of almost any type of activity can be bad theorists of their own activity. On this front, consider that—as historians and philosophers of science have shown us—scientists are, in certain respects, not reliable sources about what they are doing when they engage in scientific activity. There is no reason to think that the case is different in the case of law.

Second, we think there is good reason to deny that ordinary people—ordinary legal practitioners among them—actually have fine-grained intuitions of the sort that could track distinctions between subtly differing modes of linguistic communication. In general, folk theories of linguistic communication are unlikely to track reliably the distinctions drawn on in philosophy of language and empirical linguistics. Those distinctions are made only with the use of technical terms—"semantics," "pragmatics," and even the relevant sense of "meaning" among them—that are not designed to correspond perfectly to ordinary language terms, or to express concepts that correspond perfectly to folk-linguistic concepts. For this reason, we are skeptical that folk-linguistic theories are sufficiently fine-grained to deny explicitly that our proposed account is correct. If our skepticism is warranted, then there is good reason to be skeptical that legal practitioners have views that in fact conflict with the metalinguistic analysis. This line of reasoning would thus suggest that our account does not in fact qualify as an error theory.

That being said, we grant for the sake of argument that our account involves attributing *some* error to the self-understanding of legal practitioners. But the account has a good explanation for why legal practitioners are prone to making the sort of errors that they do. For instance, suppose we grant that legal practitioners see themselves as being primarily involved in canonical moral disputes. This means that they see themselves as engaging in a canonical dispute about normative and evaluative facts. Not everyone will think that legal practitioners often see themselves this way. But Dworkin (and Dworkinians) will.

Start by considering the relationship between a canonical moral dispute (e.g., about what falls under the shared concept DEMOCRACY) and a corresponding metalinguistic negotiation over the relevant morally important terms (e.g., "democracy"). As we argue above, both kinds of dispute can express genuine disagreements well worth having, and the two types of dispute ultimately turn on similar (if not identical) normative or evaluative facts. A debate in conceptual ethics about which concept to express by the

50. *See* Plunkett & Sundell, *supra* note 9; and Plunkett & Sundell, *supra* note 49. The latter in particular addresses in detail the worry that the metalinguistic analysis constitutes—especially in the legal context—a kind of error theory.

https://doi.org/10.1017/S1352325213000165 Published online by Cambridge University Press

term "democracy" in our political context is going to have a lot to do with substantive moral and political issues about how we should live and how we should organize our political institutions.[51] Many theorists tend to downplay the fact that arguments that are directly *about words* can have these features. But even granting this fact, it can be difficult to *tell* when a dispute is one or the other. The disputes have similar linguistic features, they each generate the psychological impression of being at conflict, and, as noted, they turn on similar substantive matters.[52]

Note also that—given default cultural assumptions about when something is a substantive issue worth debating—people are often skeptical of the import of debates *about words*. One upshot of this is that in politically advocating for the use of one way of using a term rather than another in a large group setting (as opposed to, say, a philosophy classroom), it can often be *strategically* disadvantageous to cede that one is advocating using a word to express a new concept.

For example, it would be a politically bad strategy to have a national political campaign in the United States that centered on the following idea: "we advocate that Americans stop caring so much about DEMOCRACY, the concept you normally have expressed by "democracy," and instead that Americans care more about such-and-such other related ideal, DEMOCRACY*." In a cultural setting where there is a rich set of positive feelings toward things that are called "democratic"—feelings that are not part of the literal meaning of the word "democracy" but are strongly associated with its deployment in our culture—it is much better to represent your campaign as involving the implementation of *real* democracy, the sort of thing that is responsive to the *true nature* of democracy.[53]

These considerations jointly suggest the following conclusion: people involved in politically weighty activity (e.g., a charged legal case that could have important political ramifications) are likely to resist the idea that the things they take to be canonical moral disputes are in fact metalinguistic negotiations. Even if that is what is really going on.

One might still worry that there is a specifically linguistic problem that we are overlooking. In typical cases, ordinary speakers—if they reflect on the matter—can tell the difference between what they literally say and what they communicate by some pragmatic mechanism. So why would this not be the case when speakers use the pragmatic mechanism of metalinguistic usage? Our response here is to reject the initial assumption. While it can be the case that speakers have some type of introspective access to the difference between semantically and pragmatically communicated content, this is in fact often *not* the case. It is true that with some classic cases of pragmatic communication, it is natural to think that speakers have reasonably reliable intuitions about how they are communicating information. In one classic

51. For further discussion of this second claim, *see* Plunkett & Sundell, *supra* note 9.
52. *See id.*
53. For further discussion of this idea, *see id.*; and Plunkett & Sundell, *supra* note 49.

https://doi.org/10.1017/S1352325213000165 Published online by Cambridge University Press

case, a speaker says that "there's a gas station around the corner" in order to communicate that that there is an *open* gas station around the corner. In such a case, the speaker might well, on reflection, come to know that she did not literally express the stronger claim.

But contrast this with cases of relevance implicature or quantity implicature. When a speaker says, "Betty was able to attend the meeting," it might not be at all clear to the speaker or to her listeners that, for all that was literally expressed, Betty may have had the ability to attend but chosen not to. Perhaps a reflective speaker would come to realize this. Perhaps not. But the matter is not reliably clear to ordinary speakers, and in any case could be settled only by empirical considerations and not with reference to ordinary intuitions.

This latter point is even clearer in the case of scalar implicatures. When a speaker says that "there are three cars in the parking lot," the classic Grician "lower-bound" analysis (discussed briefly in section 4) would have it that only the weaker "*at least* three cars" proposition is literally expressed.[54] But this analysis is now highly controversial, the locus of its own literature in philosophy of language and empirical semantics.[55] If the issue is not clear to the *theorists*, we could hardly expect the distinction between semantic and pragmatic modes of communication to be clear to ordinary speakers. There is simply no good reason to think that in general an analysis on which some part of communicated content is put on the pragmatic side of the ledger can be correct only if speakers have the intuition that they are communicating pragmatically. Especially for more subtle pragmatic mechanisms such as metalinguistic usage of terms—something that is less familiar to many *philosophers* than classic cases of relevance or quantity implicature—why think that ordinary people would be aware that they were making use of this mechanism?

There is much more to say about the sort of error that our metalinguistic account attributes to speakers in general or to legal actors in particular.[56] But the considerations described above represent the core of our response. It should be noted, however, that in the context of a comparison with Dworkin in particular, worries about attribution of error to ordinary speakers are unlikely to be decisive. Dworkin's account of seeming variation cases is equally open to the charge of attributing too much error to legal actors.

Dworkin himself is aware of this worry. He writes, "few people who use the concept of democracy would agree that what a democracy is depends on which political theory provides the best justification of paradigms of the concept. Most would insist that they rely on a criterial or commonsense account of the matter, or none at all."[57] However, he insists that the fact

54. *See Paul H. Grice. Logic and Conversation, in* STUDIES IN THE WAY OF WORDS (1989).
55. *See, e.g.*, Gennaro Chierchia, *Polarity Phenomena and the Syntax/Pragmatics Interface, in* STRUCTURES AND BEYOND (A. Belletti ed., 2004).
56. *See* Plunkett & Sundell, *supra* note 9; and Plunkett & Sundell, *supra* note 49.
57. DWORKIN, HEDGEHOGS, *supra* note 1, at 163.

https://doi.org/10.1017/S1352325213000165 Published online by Cambridge University Press

that speakers do not recognize that they are using interpretative concepts is not fatal to his theory. As he puts it, "we nevertheless need the idea of an interpretive concept to explain their behavior. . . . People are not always or even often aware of the buried theoretical structure needed to justify the rest of what they think."[58] We would put things slightly differently here. Our goal is not to *justify* what people think. It is to explain what people are doing when they engage in certain instances of thought and talk—something, we submit, that does not need to involve justification at all. But Dworkin's key thought here—that positing error does not necessarily count against an explanatory theory—is entirely right. The question is ultimately which of a rival set of explanations is better—explanations that will have a range of different explanatory advantages and problems that need to be carefully weighed against each other. With this in mind, we turn now to a more detailed comparison of Dworkin's interpretivism and our metalinguistic analysis.

VI. THEORETICAL ADVANTAGES OF THE METALINGUISTIC ANALYSIS

Suppose we grant that the metalinguistic analysis of seeming variation cases (including bedrock legal disputes) has some degree of plausibility. How might we then choose between our view and Dworkin's interpretivism? Our view has a number of advantages over Dworkin's, and we submit that they are, jointly, decisive. We describe a number of those advantages here.

First, our account of bedrock legal disputes does not involve positing the existence of any new kind of concept (or, similarly, any new "interpretative phase" that a concept might go through under certain conditions). The thesis that metalinguistic negotiation exists—and even that it is widespread—is compatible with a wide range of views about concepts. This is because the thesis does not itself depend on strong commitments about the nature of the candidate concepts that speakers advocate in metalinguistic negotiations. For instance, the thesis does not involve taking a stand on debates about internalism or externalism about conceptual content or debates between neodescriptivists and inferential role theorists. Of course, as we emphasize at length in other work, our work on metalinguistic negotiations has import for those views—particularly certain forms of contextualism, relativism, and strong internalism—that are commonly thought to be problematic because they cannot account for genuine disagreement.[59] But bolstering the case for a theory is not the same as entailing it, and the view we are offering in this paper does not itself depend on any strong commitments about the nature of the concepts that speakers advocate for.

58. *Id.* at 163.
59. *See* Plunkett & Sundell, *supra* note 9.

https://doi.org/10.1017/S1352325213000165 Published online by Cambridge University Press

Neutrality with respect to difficult questions about the nature of concepts is itself a theoretical virtue. But there are further, related virtues in this area. The thesis that *some* concepts are noninterpretive (at least some of the time) is common ground between Dworkin (even in *Hedgehogs*) and his opponents. Thus, based on a principle of conserving theoretical resources, it makes sense to pursue an account that does not introduce that additional type of concept but rather aims to explain the relevant phenomena using only the resources that are taken to be common ground. Exactly what those resources are will of course retain some elements of controversy—as noted, theorists with differing theories of concepts can make use of metalinguistic negotiation. But it is methodologically a good thing to explain as much as possible with preexisting resources, whatever they are, before introducing additional material like a brand new kind of concept. Other things being equal, conservation of theoretical resources and continuity with our current best accounts of thought and talk are both good things.

The advantage here is not *just* that our view does not require the introduction of a new kind of concept in general. It is also that it avoids introducing Dworkin's idea of *interpretative* concepts in particular. In much of contemporary philosophy, it is taken as common ground that whatever concepts are, they are things that play important explanatory roles in our overall theory of thought and talk. They are the sorts of things (1) that are a constituent component of thoughts (such that an individual's possessing the concept DOG makes it possible for her to think specifically *dog*-thoughts), and (2) that we express when we use words (and hence can be pointed to as the meaning of words in one important sense of "meaning"). Dworkin claims that interpretive concepts are the sorts of things that can also play an important role in explaining disputes of a certain sort. But, as noted above in Section III, we worry that there is not a sufficiently clear account of what interpretive concepts are—and how they differ from noninterpretive concepts—to evaluate their role in an explanatory account of seeming variation cases. We can now add to this the worry that there is not a sufficiently clear account of what interpretive concepts are to evaluate their role in an account of thought and talk more generally.

A final reason to prefer our positive account of legal seeming variation cases—and our account of bedrock legal disputes in particular—concerns the way in which thinking of LAW as an interpretive concept positions the relationship between law and morality. As Dworkin emphasizes, a consequence of taking a given concept to be interpretive is that it makes exchanges involving different views about what falls under that concept into "*value* disagreements rather than disagreements of [nonnormative] fact or disagreements about dictionary or standard meanings."[60] The reason for this is fairly straightforward: insofar as the application conditions for an interpretive concept involve the best justification of a set of practices—and

60. DWORKIN, HEDGEHOGS, *supra* note 1, at 6.

https://doi.org/10.1017/S1352325213000165 Published online by Cambridge University Press

justification is understood in a familiar normative sense—this means that facts about what falls in the extension of an interpretive concept are determined in part by normative and evaluative facts about what justifies a set of practices. The sort of justification needed, Dworkin suggests, is one that will need to involve appeal to normative and evaluative facts that are distinctively *moral* ones.

In the case of law, suppose we take *law* to be the thing—whatever it is—that is picked out by the concept LAW. Interpretivism about LAW then leads directly to a form of *legal antipositivism*—roughly, the view according to which what the law is in a given jurisdiction at a given time is ultimately grounded in moral facts as well as social facts.[61] Many philosophers want to reject legal antipositivism in favor of *legal positivism*—roughly, the view according to which the law in a given jurisdiction at a given time is ultimately grounded in social facts but not moral facts.[62] On this view, what the law *is* in a given jurisdiction at a given time does not depend on any facts of moral merit, whether of the moral merit of a law, a set of practices, or anything else.

Legal positivism—defined in the way we sketch it above—is incompatible with Dworkin's interpretivism about LAW. In contrast, our view about what is going on in bedrock legal disputes is *neutral* with respect to whether positivism or antipositivism is true. This is because the ordinary, *noninterpretive* concept that a speaker uses the word "law" to express—for example the concept C1 or C2—*might or might not* involve specifically moral application conditions as part of that concept's application conditions. One person involved in a metalinguistic dispute might advocate a concept that involves application conditions referring to moral facts, whereas her opponent might not. Nothing in our view commits one either way with respect to the question of whether moral (or otherwise normative) conditions are included among the application conditions. It depends on what a given concept's application conditions actually turn out to be. We take it as a theoretical virtue of our proposed positive view of bedrock legal disputes that, in contrast to Dworkin's interpretivism, it is neutral on a major and divisive issue in the philosophy of law and thus can be drawn on by positivists and antipositivists alike.

61. In this context, one can think of the "moral facts" as a subset of normative and evaluative facts—and, crucially, the sort that Dworkin thinks are needed to justify the total set of social practices that LAW is deployed in. There are, of course, interesting further questions about what exactly makes a normative fact a specifically "moral" one, but they are issues that do not matter for us this paper.

62. This general formulation of the positivism/antipositivism debate draws heavily on Mark Greenberg, *How Facts Make Law*, *in* EXPLORING LAW'S EMPIRE: THE JURISPRUDENCE OF RONALD DWORKIN (S. Hershovitz ed., 2006); Mark Greenberg, *Hartian Positivism and Normative Facts: How Facts Make Law II*, *in* EXPLORING LAW'S EMPIRE: THE JURISPRUDENCE OF RONALD DWORKIN (S. Hershovitz ed., 2006); and SHAPIRO, *supra* note 45. For further discussion of this basic way of framing the debate, *see* David Plunkett, *A Positivist Route for Explaining How Facts Make Law*, 18 LEGAL THEORY 139–207 (2012).

https://doi.org/10.1017/S1352325213000165 Published online by Cambridge University Press

Before moving on, we want to pause here to note an important upshot of what we have just said about legal positivism—an upshot that is not about a theoretical advantage our metalinguistic view has over Dworkin's interpretivism but rather about how to best understand the positivism/antipositivism debate itself. If our general discussion of metalinguistic negotiations in this paper is correct, then when philosophers themselves engage in the dispute over positivism and antipositivism, it is an open question whether they share the same concept. At the very least, the thesis that they do not share the same concept is entirely consistent with the fact that disagreements between positivists and antipositivists are entirely genuine. This has potentially important consequences for how we should think about the debate between positivism and antipositivists. In framing the positivism/antipositivism debate, the traditional understanding suggests that positivism will be vindicated if the concept LAW turns out to involve application conditions that concern only social facts and not moral facts, whereas antipositivism will be vindicated if LAW turns out to involve application conditions that concern moral facts in addition to social facts. This follows from thinking that *law* itself—the thing whose nature is up for debate in the dispute between positivists and antipositivists—is the thing, whatever it is, that is picked out by some shared concept, LAW.

Our work in this paper involves the idea that different people involved in legal practice employ distinct concepts when they use the term "law"—and hence casts into doubt whether they all employ some shared concept, LAW, which anchors the dispute between positivists and antipositivists. Similarly, it is not clear that the philosophers involved in the positivism and antipositivism dispute employ the same concept when they use the term "law," and hence it is not clear that they are all debating the nature of the same thing. Positivism might thus be right about the nature of one thing (the thing that certain people refer to when using the concept they express by "law") but wrong about the nature of another.

One might initially be tempted to think that this result would undercut the thesis that there is a genuine disagreement between positivists and antipositivists. Yet, given what we have argued in this paper, such a dispute—between philosophers who pick out different concepts with their word "law"—might very well reflect a genuine disagreement. It would just be a disagreement that is different in kind from how those participants may initially have thought of it and hence would involve a reconceptualization of the full range of issues that are involved in the debate between positivists and antipositivists. In short, part of that debate might involve a tacit normative disagreement about which of various competing concepts is best suited to play a given functional role in organizing our thought, talk, and practice.

Far from rendering such disagreements nongenuine, or merely terminological, such a reconceptualization can serve to clarify the stakes and the considerations relevant to participants on both sides of the debate. What are the relevant circumstances? What functional role is the concept answering

https://doi.org/10.1017/S1352325213000165 Published online by Cambridge University Press

to the word "law" meant to fill? What, in the circumstances that really matter, serves to make one candidate concept preferable to another? We leave it for future work to explore these broader methodological implications of our argument in this paper for philosophical disputes over positivism and antipositivism, as well as the basic question of whether and to what extent those disputes are metalinguistic.

VII. IS OUR DISAGREEMENT WITH DWORKIN SUBSTANTIVE?

Dworkin suggests that in order to vindicate our intuitions about the presence of genuine disagreement across disputes, theorists need to posit a special kind of concept, namely "interpretive concepts." We disagree. We think that traditional understandings of concepts and word meanings can give the kind of analysis of those disputes that Dworkin demands without the resources he thinks are necessary. This in turn leads us to different accounts of how best to explain seeming variation cases (including, crucially, bedrock legal disputes).

But what about the status of this disagreement between Dworkin and us? Unlike some opponents of Dworkin, we grant that at least *some* bedrock legal disputes exist—which, among other things, involves holding that they involve a type of disagreement that persists even when most or all of the typically adjudicating empirical facts are mutually known by the parties involved—and that they can concern an important topic (and, indeed, an important *normative* topic) that is well worth arguing about. Moreover, at the end of the last section, we suggest that the debate between positivists and antipositivists itself might best be conceived of as (at least partially) a tacit normative dispute—a conclusion that resonates with Dworkin's own treatment of the debate over positivism in *Law's Empire* and beyond. With this much in common between Dworkin and us, perhaps Dworkin's view and ours are themselves terminological variants. If they are, does our dispute express a *merely* terminological disagreement or is it a disagreement where matters of substance hang on our terminological choices?

Dworkin's interpretivist analysis of a given bedrock legal dispute (or another seeming variation case) and a metalinguistic analysis of that dispute could turn out to be terminological variants, one might argue, in virtue of the fact that we have concerned ourselves quite explicitly with the theoretical resources of linguistics and the philosophy of language, while Dworkin's project is only very indirectly an empirical one about language. Perhaps Dworkin has put his finger on a certain class of phenomena and proposes an analysis at a fairly high level of abstraction (one adequate to his own purposes), and one possible implementation of that proposal is the more detailed view on offer here. If this understanding of the dialectic were correct, then we would show not that Dworkin's argument for interpretivist

https://doi.org/10.1017/S1352325213000165 Published online by Cambridge University Press

concepts is unsound but rather that one way for a concept to be interpretive is for it to lend itself to metalinguistic usage and disagreement.

While we are happy to recognize that Dworkin does draw attention to crucial aspects of legal thought and practice, we reject any understanding of the dialectic that sees a metalinguistic view of bedrock legal disputes as a particular way of spelling out Dworkin's interpretivism. Notions such as "concept," "word meaning," and the like are useful only insofar as they play certain explanatory roles in our more general theorizing about language, thought, and behavior. We hope we demonstrate here the possibility of a satisfying analysis of bedrock legal disputes that works nicely with and sits easily against an enormous background of philosophical and empirical work on concepts, conceptual analysis, semantics, and pragmatics. It seems misleading in the extreme to view this type of account as an implementation of a more abstract account positing purportedly fundamental and empirically unmotivated divisions within categories that on our account remain unified.

To appreciate the problem with this proposed way of representing the relationship between our view and Dworkin's, consider the following question: What explanatory role does Dworkin's more abstract and disjunctive notion of "concept" (one that includes both noninterpretative concepts and interpretative concepts) play in our understanding of language, thought, and behavior? What role *could* it play if one holds fixed the thesis—central to this reading of the relationship between our view and Dworkin's—that the distinctions Dworkin draws our attention to can be explained without reference to any actual difference between kinds of concepts but simply with reference to different communicative mechanisms employing a single, unified category of concept? Questions like this lead us to conclude that if our disagreement with Dworkin is terminological, it is not because our own view could be understood as an implementation of his.

But perhaps there are other ways in which our dispute with Dworkin could fail to express a genuine disagreement. Dworkin proposes that bedrock legal disputes are settled with reference to the norms that best justify a set of practices. We propose that at least some bedrock legal disputes are settled by the question of which of two or more competing concepts is most appropriate to the circumstances. But how will we spell out what it is for a legal concept to be most appropriate to its circumstances? If the substantive normative facts that settle the disagreement turn out to be the same on our view and Dworkin's, then perhaps we still manage to draw a distinction without a difference.[63]

We also reject this strategy for arguing that our disagreement with Dworkin is not genuine. There are two reasons for this. The first of these reasons has already been described: Even if our theory and Dworkin's were to yield the same conclusions, ours would still make use of fewer theoretical resources, posit fewer distinctions within the relevant categories, and

63. Thanks to Ronald Dworkin for helpful discussion on this issue.

https://doi.org/10.1017/S1352325213000165 Published online by Cambridge University Press

fit more easily within established theories of language, cognition, and behavior. Even setting these considerations aside, however, we have a second reason for rejecting this strategy, namely, that it is far from clear that the important normative facts will in fact be the same on the two views. The facts that best justify *the use of one concept over another in a context* are hardly *guaranteed* to coincide with the facts that best justify *a whole set of practices*. Thus there is no reason to think that Dworkin's view and our own in fact line up with respect to the substantive normative conclusions to which they lend themselves.[64]

Finally, it is worth remembering an important and broader theoretical difference between our proposed view and Dworkin's. As we emphasize above, the metalinguistic view of bedrock legal disputes that we offer in this paper is neutral with respect to the positivist/antipositivist debate. In contrast, Dworkin's view is not neutral in this way. Now, perhaps both Dworkin and we are wrong that interpretivism about LAW commits one to a form of antipositivism. However, if we are not wrong about this, then this gives a further reason to resist the idea that our dispute with Dworkin lacks substance. Indeed this difference should be particularly salient to Dworkin and other interpretivists who understand their argument for interpretivism to be an important part of a larger argument for legal antipositivism.

Because of these considerations, we take it that we provide here the outlines of a genuine alternative to Dworkin's account of bedrock legal disputes, an account that rejects the idea of interpretive concepts and instead makes use only of standard kinds of concepts (whatever those may be). A crucial component of this account is the claim that some bedrock legal disputes are expressed in metalinguistic negotiations about which of a set of competing concepts is to be used in a given context—an idea that a wide range of philosophers can make use of in combination with a wide range of noninterpretivist views on concepts. Whatever the long-term prospects for such a view, we submit that metalinguistic negotiations exist, that they are a widespread and deeply embedded feature of ordinary communication, that they seriously undermine Dworkin's general disagreement-based argument for interpretivism, and that they provide an important and underappreciated theoretical resource for philosophers of law.

64. This general point is even more important if one takes interpretivism's commitment to "fit with the practices" to be an additional criterion on top of the general justificatory demand to "best justify" those practices. This reading of Dworkin is widespread, and there are many passages in DWORKIN, EMPIRE, *supra* note 1, that suggest such a reading. If such a reading of interpretivism is correct, then, whereas interpretivism demands fit with past legal practices, there is no general requirement that a proponent of a given concept needs to be at all concerned with fitting (let alone justifying) the practices of a community. We are inclined to follow Mark Greenberg's lead in thinking that fit is best understood as an aspect of justification, as in Greenberg, *How Facts*, *supra* note 62; and Greenberg, *How Facts II*, *supra* note 62. Dworkin's comments in response to Greenberg in Ronald Dworkin, *Response*, *in* EXPLORING LAW'S EMPIRE: THE JURISPRUDENCE OF RONALD DWORKIN (S. Hershovitz ed., 2006) imply that he agrees with the suggestion.

https://doi.org/10.1017/S1352325213000165 Published online by Cambridge University Press

WORKS CITED

Barker, Chris. 2002. "The Dynamics of Vagueness." *Linguistics and Philosophy* 25:1–36.

——. 2013. "Negotiating Taste." *Inquiry* 56(2–3):240–257.

Burgess, Alexis, and David Plunkett. 2013a. "Conceptual Ethics I." *Philosophy Compass* 8(12):1091–1101.

——. 2013b. "Conceptual Ethics II." *Philosophy Compass* 8(12):1101–1110.

Chalmers, David J. 2011. "Verbal Disputes." *Philosophical Review* 120(4):515–566.

Chierchia, Gennaro. 2004. "Polarity Phenomena and the Syntax/Pragmatics Interface." In A. Belletti, ed., *Structures and Beyond* (Oxford: Oxford University Press).

Dworkin, Ronald. 1986. *Law's Empire* (Cambridge, MA: Belknap Press).

——. 2006. "Response." In S. Hershovitz, ed., *Exploring Law's Empire: The Jurisprudence of Ronald Dworkin* (Oxford: New York: Oxford University Press).

——. 2011. *Justice for Hedgehogs* (Cambridge, MA: Harvard University Press).

Egan, Andy. 2010. "Disputing about Taste." In R. Feldman and T. Warfield, eds., *Disagreement* (Oxford: Oxford University Press).

Gibbard, Allan. 1990. *Wise Choices, Apt Feelings: A Theory of Normative Judgment* (Cambridge, MA: Harvard University Press).

——. 2003. *Thinking How to Live* (Cambridge, MA: Harvard University Press).

Greenberg, Mark. 2006. "How Facts Make Law." In S. Hershovitz, ed., *Exploring Law's Empire: The Jurisprudence of Ronald Dworkin* (Oxford: New York: Oxford University Press).

——. 2006. "Hartian Positivism and Normative Facts: How Facts Make Law II." In S. Hershovitz, ed., *Exploring Law's Empire: The Jurisprudence of Ronald Dworkin* (Oxford: New York: Oxford University Press).

——. 2011. "The Standard Picture and Its Discontents." In L. Green and B. Leiter, eds., *Oxford Studies in Philosophy of Law*. Volume 1 (Oxford: Oxford University Press).

Grice, H. Paul. 1989. "Logic and Conversation." In *Studies in the Way of Words* (Cambridge, MA: Harvard University Press).

Jackson, Frank. 1998. *From Metaphysics to Ethics: A Defence of Conceptual Analysis* (Oxford: Clarendon Press).

Kaplan, David. 1989. "Demonstratives: An Essay on the Semantics, Logic, Metaphysics, and Epistemology of Demonstratives." In Joseph Almog, John Perry, and Howard Wettstein, eds., *Themes from Kaplan* (Oxford: Oxford University Press).

Kennedy, Chris. 2007. "Vagueness and Grammar: The Semantics of Relative and Absolute Gradable Predicates. *Linguistics and Philosophy* 30(1):1–45.

Leiter, Brian. 2009. "Explaining Theoretical Disagreement." *University of Chicago Law Review* 76(3):1215–1250.

Lewis, David. 1983. "New Work for a Theory of Universals." *Australasian Journal of Philosophy* 61(4):343–377.

Ludlow, Peter. 2008. "Cheap Contextualism." *Philosophical Issues* 18:104–129.

Margolis, Eric, and Stephen Laurence. 1999. *Concepts: Core Readings* (Cambridge, MA: MIT Press).

Plunkett, David. 2012. "A Positivist Route for Explaining How Facts Make Law." *Legal Theory* 18(2):139–207.

https://doi.org/10.1017/S1352325213000165 Published online by Cambridge University Press

Plunkett, David, and Timothy Sundell. 2014. "Antipositivist Arguments from Legal Thought and Talk: The Metalinguistic Response." In G. Hubbs and D. Lind, eds., *Pragmatism, Law, and Language* (New York: Routledge).

——. 2013. "Disagreement and the Semantics of Normative and Evaluative Terms." *Philosophers' Imprint* 13(23):1–37.

Shapiro, Scott. 2011. *Legality* (Cambridge, MA: Harvard University Press).

Sider, Theodore. 2012. *Writing the Book of the World* (Oxford: Oxford University Press).

Smith, Michael. 1994. *The Moral Problem* (Cambridge, UK: Blackwell).

Sundell, Timothy. 2011. "Disagreement, Error, and an Alternative to Reference Magnetism." *Australasian Journal of Philosophy* 90(4):743–759.

——. 2011. "Disagreements about Taste." *Philosophical Studies* 155(2):267–288.

Toh, Kevin. 2005. "Hart's Expressivism and His Benthamite Project." *Legal Theory* 11(2):75–123.

——. 2011. "Legal Judgments as Plural Acceptances of Norms." In L. Green and B. Leiter, eds., *Oxford Studies in Philosophy of Law.* Volume 1 (Oxford: Oxford University Press).

U.S. Department of Justice, Office of the Assistant Attorney General. 2002. Memorandum for Alberto R. Gonzales, Counsel to the President, Re: Standards of Conduct for Interrogation under 18 U.S.C. secs. 2340–2340A.

United Nations, General Assembly. 1984. Convention Against Torture and Other Cruel, Inhuman or Degrading Treatment or Punishment. Opened for signature December 10, 1984, (entered into force June 26, 1987). United Nations Treaty Series 1465.

Wedgwood, Ralph. 2007. *The Nature of Normativity* (Oxford: Oxford University Press).

https://doi.org/10.1017/S1352325213000165 Published online by Cambridge University Press