Ergo AN OPEN ACCESS JOURNAL OF PHILOSOPHY

Deliberative Authority and Representational Determinacy: A Challenge for the Normative Realist

TRISTRAM MCPHERSON
The Ohio State University

This paper aims to illuminate the deep explanatory challenge to normative realism that underlies two familiar objections to that view. According to the first objection, the realist cannot explain the distinctive deliberative significance of our normative thought. According to the second objection, the realist cannot explain how our normative thought has determinate representational content. This paper argues that, properly understood, these seemingly distinct objections are in fact mutually reinforcing sides of a single deep challenge to the realist's ability to explain what we know about our normative thought: the *deliberation/representation challenge*. I spell out this challenge and suggest that it is an excellent candidate for the deepest problem in the philosophy of mind for the normative realist. This paper also shows that addressing this challenge to the normative realist—as opposed to the moral realist—is illuminating. There are natural and promising replies to each side of the challenge available to the moral realist that are not available to the normative realist. But focusing on the challenge to the normative realist helps us to see that these replies fail to address the challenge in its deepest and most general form.

Controversy about the credibility of normative realism is endemic to contemporary metaethics. Some take realism to be "*obviously*, the default position," while others, to put it mildly, do not.[1,2] In the face of such persistent

1. The quote is from Enoch (2018: 30), emphasis his. By way of contrast, Hartry Field describes Enoch's brand of normative realism as "obviously false" on the dust jacket of Enoch (2011).

2. In this paper, single quotation marks (e.g., 'cat') are used strictly to mention linguistic items. Double quotation marks (e.g., "cat") are used for a variety of tasks including quoting others' words, scare quotes, and mixes of use and mention. Terms in small caps (e.g., CAT) pick out concepts. I will not provide or assume a specific theory of concepts, but I will assume that concepts are representational mental entities, which are the constituents of propositional thoughts.

Contact: Tristram McPherson <dr.tristram@gmail.com>

http://dx.doi.org/10.3998/ergo.12405314.0006.045

controversy, it can be valuable to step back from the myriad arguments for and against realism, and seek to understand the challenges that face this view in their deepest and most general form. This paper aims to achieve this deeper understanding with respect to a pair of familiar challenges to normative realism in the philosophy of mind.

The most familiar contemporary incarnation of the first challenge is Christine Korsgaard's objection that the realist cannot answer the "normative question" (1996). It is famously controversial how to understand this objection, and whether it presents the realist with a real challenge at all. I argue that the clearest and most general version of this challenge is for the realist to explain what we know about the distinctive role of normative beliefs in deliberation (§3). If, as the realist believes, such thoughts are simply beliefs about certain types of properties, why do they play such a deliberative role? Call this the *deliberative authority* challenge.

The most familiar contemporary instance of the second challenge is Terence Horgan and Mark Timmons's *Moral Twin Earth* argument (e.g., 1992). Horgan and Timmons frame their objection as concerning the inability of certain sorts of realist to explain certain facts about genuine moral disagreement across possible linguistic communities. I argue that this is not the deepest version of this challenge, which is rather for the realist to explain the possibility of our normative concepts having acceptably determinate representational content (§4). One way of seeing the force of this general challenge is that apparently competent speakers can disagree deeply, broadly, and seemingly intractably about (for example) what we ought to do. This can make it seem difficult for the realist to explain how normative concepts could have determinate representational contents. Call this the *representational determinacy challenge*.

One reason to suspect that these two challenges are connected is that instances of each have been presented as a way of extending G. E. Moore's (1903/1993) "Open Question" argument.[3] However, given the notorious difficulties of interpreting that argument in a way that makes it probative, this is far from compelling. This paper argues that this suspicion can be vindicated: properly understood, these seemingly distinct challenges are in fact two sides of a single deep challenge. The realist needs to explain how a single state could possess both deliberative authority and representational determinacy (§5). The unity of the challenge is illuminating for at least two reasons. First, the two sides of the challenge are mutually reinforcing. For example, I show that it would be much easier to explain the representational determinacy of a state that did *not* have the deliberative significance of our normative thoughts. Second, I show that when

3. For these connections to the open question argument, see for example Darwall, Gibbard, and Railton (1992/1997: 4) and Horgan and Timmons (1992) respectively.

properly understood the two challenges stand and fall together. For example, were the realist able to adequately explain how our normative thoughts had determinate representational content, we would thereby have solved the deliberative authority challenge.

In my view, this unified *deliberation/representation challenge* is the deepest challenge to the normative realist in the philosophy of mind. It also has some claim to be the deepest challenge to the normative realist, period. This is because the familiar metaphysical and epistemological challenges to the realist arguably cannot even get going if we cannot give an adequate realist construal of our normative thoughts.

It is worth emphasizing that this paper lays out an explanatory challenge, rather than levying an objection. To turn the challenge into an objection one would have to show that no plausible combination of theories of representational content and the nature of normative thought available to the realist could meet the explanatory challenge. And that is far too large of a task for a single paper.

Explanatory challenges can only be as compelling as what they claim needs explaining. In light of this, I seek to ground my formulation of the deliberation/representation challenge in some highly plausible data about our normative thought, which I lay out in Section 1. In doing so, I pay unusually close attention to what shifts when we focus our metaethical attention on normativity rather than morality. This shift in focus has been a prominent trend in recent metaethics.[4] But it is not always clear how the shift matters. This paper seeks to demonstrate that—at least in this case—the shift can help us to strip away what is extraneous from a metaethical challenge, to reveal its philosophical core. Thus, I show that there are natural and promising replies to each side of the challenge available to the moral realist that are not available to the normative realist. But focusing on the challenge to the normative realist helps us to see that these replies fail to address the challenge in its deepest and most general form.

1. Normative Thought

This section begins to execute the methodology just mentioned. It introduces the sort of normativity that is my target and explains why it is of distinctive philosophical interest. It then introduces some data concerning what changes when we shift our focus from moral thought to the relevant sort of normative thought.

To begin, notice that there is a perfectly good meaning of the word 'norma-

4. Some enormously influential exemplars of the shift include Gibbard (2003), Street (2006), Schroeder (2007), Wedgwood (2007), Parfit (2011). Of course, the shift also has notable exceptions, such as the considerable recent literature responding to the account of moral concepts offered by Cuneo and Shafer-Landau (2014).

tive,' on which any standard that we can conform to or contravene is normative. Call such norms *merely formally normative*. For example, consider the norm which states that every sentence in this paper must begin with the letter 'A.' It is likely that nobody has ever cared about that norm. Unlike that norm, some norms (such as the rules of etiquette or chess) have *currency* in some way in our psychologies: our deliberation and decisions are often responsive to conclusions framed using these normative concepts.[5]

There is reason to think that mere formality and having currency do not exhaust the interesting sorts of normativity. For example, many philosophers think that there is a *normative* asymmetry between morality, on the one hand, and the rules of etiquette and chess, on the other. One way to state such a contrast is to make claims that link morality to further norms, which are presumed to be part of the *most* normative family of norms. Such "most normative" norms are labelled in the literature with a variety of terminology: reasons in the "standard normative sense," "reasons-implying" norms, the "all-things-considered" ought, etc.[6] In this paper I will follow my (2018) which introduces the technical term 'authoritative' to mark these "most normative" norms. There is plausibly a family of authoritatively normative concepts: AUTHORITATIVE REASON, AUTHORITATIVE VALUE, etc. But here, for simplicity, I will focus on the concept AUTHORITATIVE OUGHT. The concept AUTHORITATIVE OUGHT is used to pick out a standard just insofar as it is the most normative standard with ought-structure. This has a consequence that will be important in Section 4: it is part of this concept that it is *uniquely* authoritative: there could not be a distinct concept OUGHT* that has a conflicting extension but was just as normative as AUTHORITATIVE OUGHT.[7]

It is worth emphasizing that this orientation deploys rather than cashes out the "most normative" metaphor. Attempting to do so would take me too far afield, however.[8] So here I will follow the bulk of the literature and assume that such talk is meaningful and theoretically important. With this assumption in hand, I can clarify that this paper concerns challenges to realism about authoritative normativity.

Authoritatively normative concepts are an especially natural focus for metaethical attention in their own right. They also provide a useful means for formu-

5. I take the 'merely formally normative' terminology from Railton (1999/2003: 323). I take the 'currency' terminology from Copp (1995: 22). See also Woods (2018) for helpful discussion.

6. For these labels, see Scanlon (1998), Parfit (2011), and Wedgwood (2004), respectively. For more such labels, and why they can be misleading, see McPherson (2018).

7. It is compatible with this way of understanding the concept that there might be genuine AUTHORITATIVE OUGHT dilemmas: circumstances in which you both authoritatively ought to do A and authoritatively ought to refrain from doing A.

8. For worries about our ability to cash out the metaphor, see Copp (2005a; 2005b), Tiffany (2007), Baker (2018). Recent replies to these challenges include McPherson (2018) and Wodak (2019).

lating important debates about other normative standards. For example, these concepts provide an extremely natural way of formulating a familiar hypothesis about the normativity of morality:

Moral Rationalism If you morally ought to do A, then you authoritatively ought to do A

Certain modern classics concerning the foundations of morality suggest that an adequate theory of morality will entail something like Moral Rationalism.[9] And we can state the intended contrast with etiquette (e.g.) by denying that one always authoritatively ought to do what etiquette requires one to do.

Authoritatively normative concepts also usefully regiment other recently prominent metaethical questions. For example, philosophers have begun to systematically investigate whether *structural rationality* or *well-being* or *epistemic* or *legal standards* are normative.[10] Such philosophers are not well understood as asking whether these standards are formally normative, or whether they commonly have currency in our psychologies. Of course they are, and do! Rather, such philosophers are best understood as querying the relationship between these standards and *authoritatively normative* conclusions.

Moral Rationalism may be true (and even knowable a priori to be true). But it is not *transparently* true: Moral Rationalism is a substantive and interesting thesis about which philosophers disagree. This shows an interesting contrast between the concepts MORAL OUGHT and AUTHORITATIVE OUGHT. As it is deployed in claims like Moral Rationalism, AUTHORITATIVE OUGHT is *transparently* authoritative, while MORAL OUGHT is (at least) less transparently so. Because of this, the concept AUTHORITATIVE OUGHT also provides us with a clear way to deny that there is a deep normative asymmetry between morality and other normative standards, such as etiquette: one can do this by denying that Moral Rationalism is true.

My aim in the remainder of this section is to identify what changes when our theoretical focus shifts from moral concepts to authoritatively normative concepts. One might begin this project by offering a theory about the nature of moral concepts and authoritatively normative concepts. But because such theories

9. Central examples include Smith (1994) and Korsgaard (1996). For recent discussion see Part One of Jones and Schroeter (2018). The mere entailment suggested by Moral Rationalism is too weak to secure what most moral rationalists want. Most moral rationalists seek to secure underlying explanatory relationships between morality and authoritative normativity—and perhaps also our nature as agents or reasoners—that would ground the relevant entailments. Because these explanatory relationships can take a wide variety of forms, I have used entailment as a simplifying proxy, even though there are possible ways of securing an entailment that would not satisfy most rationalists.

10. See for example Kolodny (2005), Broome (2013), Kieswetter (2017), Lord (2018), Fletcher (2019), and some of the papers in Plunkett and Toh (2019).

are highly controversial, doing so would limit the interest of this paper. Instead, I will spell out what some of the crucial *data* are, concerning the shift from moral to authoritatively normative concepts.

The data I focus on concern certain patterns in thought that are characteristic of persons who possess these concepts in the normal way, and not merely deferentially. To make this last qualification clear, focus on the concept SPIN, which is deployed in parts of contemporary physics concerned with entities like quarks. I take it that, on a natural picture of concept-possession, I possess this concept. But I'm no physicist, and my possession of this concept is plausibly explained in large part by my deferential relations to people who actually understand it.[11]

In the cases I want to focus on, one's grasp of a concept is not overwhelmingly deferential in this way. In these cases, possession of a concept by its ordinary users is correlated with tendencies to deploy the concept in thought in certain characteristic ways. For example, possessors will tend to form beliefs involving the concept on the basis of characteristic types of evidence or grounds. And they will tend to form certain further mental states on the basis of tokenings of the concept. Call these *upstream* and *downstream* tendencies, respectively. One sort of data I consider concerns these sorts of tendencies associated with ordinary possession of a concept.

I also consider data concerning the range of such tendencies that is treated as compatible with possession of a concept, absent some special explanation. For example, imagine a speaker who appears to systematically use the word 'cat' to attempt to talk about cookie tins. Other things equal, we would likely conclude either that there is something badly wrong with this person's grasp of the concept CAT, or that they are using the word 'cat' to express a wholly different concept. Notice that only some sorts of puzzling talk supports such conclusions: a person who sincerely said "I love my cat jumping on my face in the morning, and peeing in my shoes!" is very puzzling, but does not raise doubts about whether they possess CAT.

To be maximally illuminating, and yet keep things manageable, I focus only on the tendencies associated with three structurally similar first-person thoughts. Specifically, I focus on the contrast between a paradigmatically non-authoritative normative thought, a moral thought, and an authoritatively normative thought:[12]

11. On some views of concept-possession, one cannot possess a concept deferentially in this way. Most of what I am saying here can be translated without loss by those who accept such views.

12. My example "thoughts" include the word 'my,' which suggests indexical or *de se* structure. If you theorize thought in a way that avoids such structure, feel welcome to translate my examples as appropriate.

Chess	This is my only option that is a legal chess move
Moral	This is my only option that I morally ought to perform
Authoritative	This is my only option that I authoritatively ought to perform

I will begin by focusing on contrasts between these types of thoughts on the upstream side. Consider the characteristic tendencies associated with the *Chess* thought. This thought is characteristically tokened while playing chess, after seeking to identify the moves that constitute one's legal chess moves, and identifying only one. More generally, although variation is clearly possible, we would expect that only a fairly narrow range of types of upstream states and circumstances—almost all involving thoughts about chess—would tend to elicit thoughts like *Chess*. Further, we would expect a relatively high degree of consensus among users concerning the relevant grounds. Finally, where there is disagreement about the proper grounds for a conclusion like *Chess*, we would expect users to overwhelmingly recognize a shared means of adjudicating that disagreement, such as reference to some statement of the rules of chess.

Next consider the thought I dubbed '*Moral*.' The range of characteristic bases for forming a thought like this are more varied than for *Chess*, but not unconstrained. Philosophers sometimes talk about "moral considerations" that favor an option, and (where they think about such things) deliberators typically judge that they morally ought to perform an act on the basis of such considerations. Clusters of paradigmatic moral considerations include those associated with well-being, respect, and cooperativeness. The penumbra of considerations whose status as moral is more controversial include (for example) a range of self-oriented reasons, and reasons grounded in the preservation of species or of aesthetic value.

The range of considerations that a speaker *could* treat as a basis for *Moral*, without raising questions about her competence with the concept MORAL OUGHT, is broader than this. However, as Philippa Foot showed us, this range appears to have limits. To adapt one of her examples (1958–59: 92), if someone insisted that the mere fact that an option involved clasping one's hands entailed that one morally ought to do it (and not because of some further moral consideration in the background), we would begin to question their grasp of the concept MORAL OUGHT.

A further contrast between MORAL OUGHT and LEGAL CHESS MOVE is that moral disagreement is, famously, much less tractable. At least in the contemporary context, there is no means of adjudicating moral disagreement that is widely accepted in the way that the documented rules of chess are.

Finally, consider the thought I dubbed '*Authoritative*.' The range of bases on which people typically form thoughts of this kind is arguably broader than for *Moral*: along with moral considerations, one might base a thought of this kind on

considerations about one's beloved, considerations connected to one's enlightened self-interest, considerations connected to one's projects, etc.[13] The same is true for the range of *intelligible* bases. Consider again, for example, the fact that an option involves hand-clasping. Some philosophers take treating this as a basis for forming a thought like *Authoritative* to be as puzzling as it would be to treat it as a basis for forming a thought like *Moral*. Others, however, take this to be a perfectly conceptually intelligible (if psychologically puzzling) basis for forming a thought like *Authoritative*.[14]

We can sum up these contrasts as follows: there is much less uniformity in the upstream tendencies associated with MORALLY OUGHT than LEGAL CHESS MOVE. And, although the contrasts are subtler, there is arguably less uniformity still for AUTHORITATIVE OUGHT.

Next consider the downstream tendencies associated with tokening *Chess*. Having tokened this thought, one might seek to recheck whether its content is true, or one might (given certain other salient thoughts) offer a draw or resign. But most characteristically, one would form the intention to play the chess move this thought identifies.

This tendency, however, is also *fragile* in a familiar way: if I notice that my child has hurt himself just as I have tokened *Chess*, I am likely to simply ignore the game, and help him. More generally, unless one is deeply invested in one's game, or worried about being inconsiderate to one's opponent, almost any apparent reason to perform another action can be enough to defeat the tendency to play the move identified in *Chess*.

The tendency to act on the content of thoughts like *Moral* is typically *less fragile* than that for thoughts like *Chess*. This is well-illustrated by Michael Smith's example of concluding that he morally ought to give to a famine relief canvasser. Given this conclusion, Smith claims plausibly that only certain specific sorts of further information would remove your expectation that he then gives to the canvasser: namely, that he thinks there are stronger reasons to do something else, or that he is weak of will (1994: 6–7). At a coarse grain, this is parallel to the case *Chess*. The core contrast is that we do not expect almost any reason to be enough to defeat one's tendency to act morally: typically only a weighty competing consideration will do this explanatory work.

There is another way of dispelling this puzzlement, not mentioned in this

13. Here I assume that authoritatively normative thoughts are ubiquitous and ordinary. An important alternative view is that such thoughts are arcane technical ruminations ordinarily had only by a few philosophers. That alternative view raises important complications which I cannot adequately address here.

14. This division was striking among the audience for an ancestor of this paper at the 2018 Central APA. I am indebted to Josh Gert for discussion that helped me to better understand the conceptual puzzlement reaction.

passage from Smith: we could learn that the unmoved moral judge does not believe that moral conclusions entail authoritative reasons. Many excellent philosophers have defended views that at least suggest this as an implication. Think of Plato's Thrasymachus on justice, or Nietzsche's account of morality, or of various philosophers in the Humean tradition.[15] If a person accepted such a view, their failing to see that there is any deliberative significance to their conclusion that they morally ought to give to the canvasser would be akin to a chess player starting to make an artwork with their pieces in mid-game: psychologically surprising, perhaps, but not an obvious indication of any conceptual confusion or defect.[16]

Finally, consider the downstream tendencies associated with the thought *Authoritative*. Again, there is a presumption that one will act in accord with such a thought. But again, people do not always act in accord with such thoughts. Consider the analogue of Smith's question here: what would dispel our puzzlement about someone failing to act on their salient occurrent thought of this form?

It is clear that weakness of will is again one explanation, as it was in the case of MORAL OUGHT.[17] The other explanations in the moral case do not transfer as easily to the authoritative case. For example, to think that one authoritatively ought to perform an act, but that one has more authoritative reason to perform another, is ordinarily deeply confused. Similarly, while philosophical reflection could lead someone to believe that they authoritatively ought to ignore moral considerations, it is transparently confused to think that one authoritatively ought not to do what one authoritatively ought to do.[18]

We can sum up these contrasts as follows: as we move from thoughts involving LEGAL CHESS MOVE to those involving MORAL OUGHT and then AUTHORITATIVE OUGHT, the downstream significance of the relevant normative thoughts appears to become stronger, and less intelligibly defeasible.

This section introduced the authoritatively normative concepts that are the

15. For example Foot (1972), Williams (1985), and Street (2012).

16. Recall that I am focusing here on a class of *data* that we have to work with, and not on the nature of moral concepts. This data is compatible with the view that Moral Rationalism is a conceptual truth (for this idea, see Smith 1994: Ch. 3).

17. Note that the relevant sort of "weakness" can be systematic. Some people are simply moved little (if at all) by certain classes of their normative beliefs, either emotionally or deliberatively. For example, I have known smokers who believed that they ought never smoke again, but for whom this belief was deliberatively inert. Not only did they continue to smoke, they routinely structured their deliberation around smoking. For example, they would make plans in ways that ensured that they could buy another pack later in the day.

18. I am seeking to establish a contrast here, not seeking to show that relevant doubts are completely unavailable. For example, one could intelligibly have the first sort of doubt if one thought that authoritative ought-dilemmas were possible. And Eklund (2017: Ch. 7) suggests an important way to construct an intelligible form of doubt about the significance of one's AUTHORITATIVE OUGHT thoughts.

target of the discussion in this paper. It then sketched the contrasting tendencies in thought associated with thoughts tokening three different normative concepts. This brief discussion suggests two conclusions. First, with respect to both the upstream and downstream tendencies, MORAL OUGHT and AUTHORITATIVE OUGHT are more similar to each other than they are to LEGAL CHESS MOVE. But at the same time there are subtle contrasts, which we might boil down in this way: AUTHORITATIVE OUGHT is associated with even less upstream determinacy than MORAL OUGHT, and with stronger and less defeasible downstream significance. I will appeal to these data in understanding our two challenges to the normative realist. Before doing that, however, I briefly characterize the target realism.

2. Characterizing the Target Realism

This paper develops a challenge for *realism* about authoritative normativity (for brevity in what follows, I will call this *normative realism)*. However, 'realism' is used in metaethics to mean many different things. So in this section I very briefly clarify the (orthodox) characterization of realism that I treat as the target of the challenge. This characterization consists of five commitments.

The first commitment is *cognitivism*: this is the claim that ordinary thoughts in which AUTHORITATIVE OUGHT features—thoughts like *Alice ought to intervene*—are beliefs, which represent the world as being a certain way. This contrasts with views on which normative thoughts are desire-like states (or are wholly constituted by them).

The second commitment is *invariantism*. This consists of two claims. First, AUTHORITATIVE OUGHT makes (at least nearly) the same contribution to the content of thoughts that token it, across contexts of tokening.[19] Second, any normative thought with a completely determinate intension has the same truth-value across different evaluators (unless that thought has some *other* constituent which requires a relativist interpretation).[20]

The third commitment is *acceptable representational determinacy*. On many accounts, the representational import of our words and concepts can be *imprecise.* For example, prominent ways of theorizing vague predicates such as 'is bald' or 'is a heap' entail that there are certain entities to which these predicates neither

19. Notice that this claim about the concept is compatible with the plausible thesis that the English word 'ought' is context-sensitive. I include the qualification 'at least nearly' because I take normative realism to be compatible with some vagueness, and context sensitivity provides one important way of theorizing vagueness (e.g., Graff 2002). For a helpful introduction to context-sensitivity metaethics, see Silk (2018).

20. For a helpful introduction to the variety of relativist theories in metaethics, see Stojanovic (2018).

determinately apply nor determinately fail to apply. I take realism to be compatible with some representational indeterminacy. For example, I take it to be compatible with realism that there is no precise contribution to an effective charity that marks the lower bound of the extension of REQUIRED, as applied to me now, such that my giving a penny less would fall outside of the extension, and my giving a penny more would fall in the extension of SUPEREROGATORY. However, I take realism to require that this indeterminacy be limited. For example, I will treat it as incompatible with realism that the intension of AUTHORITATIVE OUGHT is indeterminate between the intension suggested by maximizing act-utilitarianism and that suggested by a satisficing act-egoism.

The realist's fourth commitment is to a normative ontology. The core idea of this commitment is that it is not enough for normative realism that our thought be a certain way: the world must "cooperate," furnishing appropriate facts or properties for that thought to be about. This fourth commitment has three consequences, the first of which is more familiar than the others. First, it requires that (at least some of) our authoritatively normative concepts have non-empty extensions.[21] Second, it requires that the properties and relations picked out by these concepts are themselves authoritatively normative. It thus rules out nihilism about authoritative normativity: the view that there are no actually instantiated authoritatively normative properties.[22] Third, recall from Section 1 the idea that AUTHORITATIVE OUGHT is uniquely authoritative: that there could not be another ought-concept, OUGHT* that was equally authoritative as the former concept, but had a different extension. Given realism, this requires that the extension of OUGHT* is not an authoritatively normative relation that is distinct from the one picked out by AUTHORITATIVE OUGHT. The reason for this constraint is that, if there was such a relation, reality would fail to vindicate the presumption of unique authority which is central to our grasp on authoritatively normative concepts.

The realist's fifth commitment is *metaontological*. Influential contemporary views like *quasi-realism* and *quietism* suggest that metaphysical-sounding claims, such as those made in the preceding paragraph, are correctly interpreted in a way that robs them of their genuine metaphysical import (e.g., Blackburn 1993; Gibbard 2003; Scanlon 2014). The realist denies these interpretations.

These conditions entail that this characterization of realism rules out a considerable number of possible views about the normative. However, it is worth

21. This commitment distinguishes realism from the *error theory*. According to the most familiar form of error theory, our normative thought purports to be representational, but there is nothing in the world for such thought to represent. For a helpful introduction to error theory, see Olson (2018).

22. The view that normative concepts might have extensions without there being normative properties or relations is not very familiar; Eklund (2017: Ch. 6) dubs this view 'presentationalism,' and provides an extended discussion of this possibility.

noting that it is at least in some respects *less* demanding than some other familiar characterizations of realism. For example, it does not require that authoritatively normative properties be *non-natural*, and it is (at least pending further argument) compatible with the view that authoritatively normative facts are fully grounded in (e.g.) facts about agents' attitudes, or facts about social practices.

3. The Deliberative Authority Challenge

This section introduces the deliberative authority challenge to the sort of realist view just sketched. One striking thing about this challenge is that, while many excellent philosophers have found it quite gripping,[23] others have struggled to see how it could be a coherent challenge at all.[24] After introducing the difficulties in formulating the challenge, I develop a formulation that draws on the data about normative thought introduced in Section 1.

As noted in the introduction, the most familiar contemporary incarnation of this challenge is due to Christine Korsgaard (e.g., 1996; 1997: 240–241). According to Korsgaard, the realist cannot "answer the normative question" (1996: 30–34) or explain how reasons "get a grip" on the agent (1997: 240–241). However, the interpretive difficulties concerning this challenge arise in an especially straightforward way for the ur-form of the challenge in contemporary metaethics, offered by P. H. Nowell-Smith. Nowell-Smith addressed G. E. Moore's thesis that our moral thoughts are beliefs about certain non-natural properties. Nowell-Smith proposed to grant (for the sake of argument) that there were such properties, and that we could come to know about them. He then asked, "Why should I *do* anything about these newly-revealed objects?" (1954: 41, emphasis original).

The difficulty of interpreting this challenge is immediately apparent. For when confronted with this question, we need to ask how to read 'should' in the quoted sentence.[25] If we read it as *morally should*, then the question seems easy to answer. After all, if we grant the non-naturalist's metaphysical hypothesis that morally should is a certain non-natural relation, then it is trivial that we morally should perform the actions to which we bear that relation.

23. How precisely to classify challenges in this vicinity is controversial, but instances of such challenges arguably include Korsgaard (1996; 1997: 240–241), Gibbard (2002: 180), Bedke (2014), Dreier (2015), and Dasgupta (2017).

24. For example, Parfit (2006) forcefully emphasizes the way in which some prosecuting the challenges in this vicinity seem to simply conflate normative or deliberative significance with motivational force.

25. A similar issue arises for Korsgaard: officially, the normative question concerns what "justifies the claims that morality makes of us" (1996: 9–10). The key question in this case is how to interpret the term 'justifies.'

A natural and less trivial reading arises if we read this 'should' as *authoritatively should*. That permits the following gloss on the challenge: why authoritatively should I perform the action that has the non-natural property of being such that I morally should do it? On this gloss, the challenge is that, if there were non-natural moral properties, they would not bear on what we authoritatively ought to do. This would constitute an objection to non-naturalism for anyone who accepts Moral Rationalism.

So formulated, however, the objection appears both disappointing and easily met. It is disappointing because it *states* that the (non-naturalist) realist can't explain Moral Rationalism, but it does not explain why. One wants to know: what is the problem supposed to be? A natural hypothesis is that the objector is presupposing an anti-realist theory of authoritative normativity. But then the moral realist's response seems straightforward: to accept a realist theory of authoritative normativity, in a way which secures the connection.[26]

The objection initially appears even less promising when levied instead against the normative realist, instead of the moral realist. If we seek to offer a normative reading of Nowell-Smith's question (authoritatively ought I to do what I authoritatively ought to do, if realism is true?) the answer is again a trivial: "yes." There is no obvious gap to appeal to here that is akin to the moral/authoritative gap that made for a less trivial version of the objection in the moral case.

We could instead try framing the challenge in psychological terms. For example, we could read it as suggesting that if people believed that normative realism were true, then they might not be motivated to do what they authoritatively ought to do.[27] But even if this psychological conjecture were true (which is unclear), it is hard to see why that would be evidence of the falsity of normative realism. After all, sometimes the truth doesn't motivate us. As David Lewis famously quipped, ". . . who ever thought that philosophy could replace the hangman?" (1996: 307).

The objector seems to face a dilemma: either the challenge is merely motivational ("Do I care about normative properties . . . ?"), in which case the answer is that one might not care. Or it is normative, in which case taking seriously the fact that the target view is realism *about authoritative normativity* seems to foreclose the objection. It might thus seem that there is no interesting objection to normative realism here.

But this conclusion is hasty. There is a deep challenge in the vicinity, which arises because of the relationship between realism and the nature of normative

26. Smith (1994: 183–184) suggests the most straightforward such account, on which moral reasons are just a substantively characterized *subset* of the authoritative reasons.

27. There are many passages in Korsgaard—for example, the talk of reasons "getting a grip"—which suggests such a psychological reading. Again, see Parfit (2006) for critical discussion.

thought. A recent treatment of the challenge by Jamie Dreier focuses on this relationship, in a way that is both illuminating and instructively flawed.

Dreier's formulation of the challenge focuses on what some philosophers have called the *enkratic principle*: the idea that we are rationally required to be motivated to do what believe we (authoritatively) ought to do (2015: 177). The explanandum here is not the *normativity* of the normative per se, but rather the seemingly obvious structurally rational link between certain authoritatively normative thoughts and our motivational states. Dreier suggests that the enkratic principle cries out for explanation, and that the realist view he targets (a realism which also accepts that *reasons* are the fundamental normative relations) cannot provide that explanation (2015: 179).

The immediate problem with Dreier's version of the challenge is that it is not at all clear why we should grant Dreier's assumption that the reasons-fundamentalist normative realist is unable to explain the enkratic principle. There is a familiar project in ethics of defending certain *inter-normative reductions*: for example, the buck-passer attempts to explain facts about value in terms of facts about reasons. And recently, Benjamin Kiesewetter (2017: esp. §9.5) and Errol Lord (2018: esp. §2.4.5) have argued in detail that we can explain in terms of reasons everything that needs explaining in the vicinity of the (alleged) enkratic principle.[28] This is part of their attempt to give a general reductive account of rationality in terms of reasons. And there appears to be nothing in their accounts which is unfriendly to normative realism. Stepping back, we can see that Dreier's formulation has the same general vice as the formulation of the challenge to the moral realist in terms of Moral Rationalism, considered above. In both cases, the challenge targets an alleged internormative connection (in one case Moral Rationalism, in the other the enkratic principle). But it is hard to see why the *realist* should have special difficulties explaining such internormative connections.

Nonetheless, Dreier's gloss on the challenge makes important progress. Unlike the straightforward readings of the challenge previously canvassed, his focuses directly on the role of normative beliefs in the agent's psychology. And clearly part of the intuitive force of the challenge is an alleged disconnect between normative realism and the agent's psychology. We can improve on Dreier's formulation by replacing his focus on the enkratic principle with a focus on the data discussed in Section 1.

As we saw, AUTHORITATIVE OUGHT thoughts appear to have a distinctive role in deliberation. But on the realist picture, such thoughts are simply beliefs about some bit of reality. The deep challenge here is that it is unclear why a mental state that is simply a belief with a distinctive content could have this distinc-

28. As a referee helpfully points out, Kieswetter and Lord do not purport to explain the enkratic principle itself as a principle of rationality, but rather explain our judgments about, e.g., rational failure that we might have seemed to need the enkratic principle to explain.

tive deliberative role. To make the challenge vivid, imagine there were some normative "halos of goodness" out there. And suppose that there were states in our psychology—HALO beliefs—that functioned to *represent* those halos. It seems like the deliberative role of such HALO beliefs could be wholly contingent on our attitudes towards halos. Such a belief could certainly be deliberatively relevant when combined with certain other intentions or desires. But it might seem quite mysterious how such a belief could be deliberatively relevant *by itself*, in the way that our AUTHORITATIVE OUGHT thoughts appear to be.[29]

Shifting our attention from moral concepts to authoritatively normative concepts has helped to clarify this challenge in two ways. First, we saw that when the challenge is applied to the moral realist, it is natural to read the challenge as concerning Moral Rationalism. But as I emphasized, the challenge is not compelling on this reading. Second, as I argued in Section 1, the deliberative role of AUTHORITATIVE OUGHT is more determinate and less intelligibly defeasible than that of MORALLY OUGHT. This means that the realist has fewer options for addressing the challenge to AUTHORITATIVE OUGHT. For example, some moral realists appear prepared to grant that whatever "authority" in deliberation MORAL OUGHT has depends on various further mental states, like caring about morality (e.g., Railton 1986: 165). It is much harder to see how the deliberative authority of AUTHORITATIVE OUGHT thoughts could be similarly contingent.

It is worth contrasting the formulation of this challenge that I have offered with the motivational or internormative formulations, considered above. As I emphasized, it seems possible to answer any internormative challenge by appealing to claims about normative structure, combined with embracing realism about more sorts of normativity. And it is comparatively easy to shrug off the motivational formulation of the challenge for the simple reason that there is no obvious datum that we are always motivated by our normative judgments. It is far harder to shrug off my formulation of the challenge, for the sorts of data that my challenge calls on the realist to explain are some of our clearest data about what is distinctive about thoughts containing authoritatively normative concepts.

4. The Representational Determinacy Challenge

In the preceding section, I argued to a formulation of the target challenge. In this section I proceed differently. The representational determinacy challenge

29. My talk of halos here is inspired by a memorable passage from Gibbard: "What if . . . you encounter a halo of goodness? We'd need to learn more about halos. If touching a halo of goodness makes you go poof and turn into a halo too, the thing to do, I'd think, is to get away" (2002: 180). One way of reading this is as dramatizing the sort of deliberative contingency suggested in the text.

is straightforward to state, so I begin by briefly stating and motivating it. I then bring out its significance by explaining its relation to the more familiar Moral Twin Earth argument. This will put me in a position to explain (in the next section) why this challenge and the deliberative authority challenge are two sides of a single deep challenge.

As I explained in the Introduction, the representational determinacy challenge questions the ability of the normative realist to explain how our authoritatively normative concepts could have acceptably determinate representational contents. The idea that this is a significant challenge can be initially motivated by the combination of two facts. First, as we saw in Section 1, the upstream role of AUTHORITATIVE OUGHT thoughts appears to be heterogeneous. Second, it is a datum that facts about the upstream role of a concept are typically correlated with its representational content. How exactly such correlations are to be explained by our theory of representational content is a matter of great debate. But the datum itself is highly plausible.

We can illustrate this datum with the concept LEGAL CHESS MOVE. What grounds the fact that moving a pawn forward two squares can fall into the extension of this concept? In one way or another, the answer is surely going to appeal to or explain our collective upstream tendencies in thought using that concept: our tendency to classify such moves as LEGAL, our tendency to defer to rules which permit such moves, etc.

We can reinforce this point by noting how easy it is to think about a distinct *variant* of chess: you and I could sit down, and agree to play a variant (which I will dub 'chessX') where pawns can only ever move forward one square. If we did so, we would plausibly be tokening LEGAL CHESSX MOVE thoughts, whose extension *excludes* moving a pawn forward two squares. It is very plausible that the explanation of the difference in extension between our LEGAL CHESSX MOVE thoughts and LEGAL CHESS MOVE thoughts will either appeal to, or explain, the systematic contrast between the upstream tendencies associated with these concepts.

With this in place, consider the relevant contrast between LEGAL CHESS MOVE and AUTHORITATIVE OUGHT. As I emphasized in Section 1, the upstream tendencies are much less uniform for AUTHORITATIVE OUGHT than for LEGAL CHESS MOVE. We tend to treat people with even fairly striking differences in their upstream tendencies as competent users of AUTHORITATIVE OUGHT. Not so with LEGAL CHESS MOVE. Here, we tend to treat even relatively minor divergences from the rules of chess as conceptual confusions to be cleared up (or, if intractable, as evidence of deployment of a different concept), and not as substantive disagreements.

The challenge is plain: uniformity in upstream tendencies is the very feature that underwrites our confidence that the concept LEGAL CHESS MOVE has reasonably determinate representational content. And that feature is highly attenuated

in the case of AUTHORITATIVE OUGHT. Notice that this challenge is at least slightly more difficult for the realist about AUTHORITATIVE OUGHT than it is for the moral realist. This is because—as we saw in Section 1—the characteristic upstream tendencies of AUTHORITATIVE OUGHT appear even more varied and less constrained than those of MORALLY OUGHT.

It might seem that the representational determinacy challenge is easily met. After all, we know that in some cases, the intensions of our concepts appear to be stable, even across significant changes in the sort of upstream tendencies we have been discussing. Even if this does not seem to be true of LEGAL CHESS MOVE, there is no reason to think it could not be true of other normative concepts.

The most prominent relevant literature here is framed in terms of linguistic as opposed to mental content. So in introducing these ideas, I will briefly shift focus to (meta-)semantics. The most famous candidates for referential stability over variation in use are natural kind terms. For example, many philosophers take it that the intension of 'atom' did not change, even after we learned that atoms in fact had internal structure, and that this structure could be altered in nuclear reactions. This despite the fact that lacking such structure was previously taken to be *definitional* of 'atom.'

In the 1980s, some philosophers sought to develop a theory of reference for moral terms that treated them as natural kind terms. The most famous example was Richard Boyd's (1988). Simplifying greatly, Boyd's proposal was that the reference-determining relation is *causal regulation*. The sorts of upstream tendencies I have been discussing will often play a role in explaining how a kind causally regulates use of a term. But it is clear that different such tendencies can causally link a language user to a single kind. For example, use of the word 'water' by a chemist and a small child can both be causally regulated by H_2O, even if there are very different upstream tendencies associated with their respective uses. This might suggest that a causal regulation meta-semantics is well-suited to explain the possibility of shared content across the sort of heterogeneous upstream tendencies that I just used to motivate the representational determinacy challenge.

Boyd's attempt to achieve referential determinacy by appeal to causal regulation, however, faces a powerful objection, initially formulated by Terence Horgan and Mark Timmons (e.g., 1991; 1992). Horgan and Timmons ask us to imagine a scenario in which the difference between Earth and what they call "Moral Twin Earth" is that a slightly different property causally regulates the use of the word 'good' on Earth than does so on Moral Twin Earth. Crucially, besides these differences, the roles of the words 'good' on the two planets (and especially the sorts of downstream tendencies discussed in §1) are stipulated to be nearly identical.

Because the case has been tailored to have two different properties causally regulating use of 'good' on the two planets, Boyd's theory entails that 'good'

refers to different properties in the mouths of Earthlings and Moral Twin Earthlings. Horgan and Timmons claim this is an implausible result. They claim that, were I to say that breaking a certain promise is "good," while a Moral Twin Earthling says that breaking that very promise is not "good" (using their twin-English term), we would be disagreeing. And they claim that the moral realist who accepts Boyd's theory cannot explain this disagreement.

This paper is focused on the content of thought, not on the semantics of words. But an account like Boyd's can be straightforwardly adapted to address the realist's representational determinacy challenge. So adapted, the view claims that the representational content of a concept is explained by what causally regulates its tokening in thought. Suppose then that there is a single relation that causally regulates our use of MORAL OUGHT. On the adapted view, this supposition could explain how this concept has determinate representational content, despite the broad range of upstream tendencies associated with its use. Unfortunately for this strategy, it is also easy to see how a Moral Twin Earth-style challenge applies to the adapted theory. If the Twins' very similar concept is causally regulated by a different relation, then it will have a different representational content than MORAL OUGHT, with similarly (allegedly) implausible results for disagreement (in this case, disagreement in thought).

As with the deliberative authority challenge, the Moral Twin Earth argument illustrates how carefully attending to the shift from moral concepts to authoritatively normative concepts can be illuminating. I will explain two payoffs. First, a promising line of reply to Moral Twin Earth is much less promising in the authoritative case. Second, it becomes possible to see that the deepest version of the challenge does not concern disagreement at all. Consider these points in turn.

First, one of the most prominent lines of reply to Moral Twin Earth involves arguing that the proponent of a theory like Boyd's can explain disagreement with one's Twin, by denying that such disagreement must be *moral* disagreement. For example, David Merli suggests that the disagreement in Moral Twin Earth cases may be non-moral but normative disagreement about what to do (2002: 232–233, cf. also Copp 2000: 3), and David Plunkett and Tim Sundell (2013) argue that it can be understood as normative *metalinguistic* disagreement about which moral-ish concepts to deploy in a certain way in a conversational or deliberative context.

This style of reply is much less promising when addressed to a Twin Earth scenario for authoritatively normative concepts.[30] For, as I have emphasized, AUTHORITATIVE OUGHT is a concept characterized by a distinctively authoritative role in deliberation. It is certainly possible to accept that someone has made a true

30. See Dunaway and McPherson (2016) for such a scenario. Plunkett and Sundell (2013: §6.2) make the difficulties of applying their strategy to (something like) authoritatively normative concepts plain, in more detail than I do here.

authoritatively normative claim, but "disagree" with that claim in one of the ways suggested above. For example, the weak-willed smoker could be understood as disagreeing about what to do with the person who tells them that they authoritatively shouldn't buy another pack. But surely this is not a charitable way of thinking about a Normative Twin Earth case: neither we nor our Twins are charitably understood as thinking of themselves as weak of will in the way the clear-eyed weak-willed smoker does. Shifting our focus to authoritatively normative talk, then, reduces the realist's promising options for replying to the Moral Twin Earth style of challenge.

The second way that the shift in focus to authoritatively normative concepts matters is that it allows us to see that Horgan and Timmons's framing of the challenge in terms of disagreement is not the deepest form of the challenge. In the authoritative case, we can understand the challenge as simply to explain how normativity can "get into thought," given certain realist commitments, as I will now illustrate.

Again, treat a causal regulation theory of content-determination as our foil. Now consider a Normative Twin Earth, in which our twins have a concept OUGHT* which is associated with just the same upstream and downstream tendencies as our AUTHORITATIVE OUGHT, but is causally regulated by a different property. (One reason to think this is possible is just that the upstream tendencies associated with AUTHORITATIVE OUGHT are so heterogeneous.) By the causal regulation theory, these two concepts represent these two different properties.

Now, recall one of the realist commitments introduced in Section 2: that if a concept is authoritatively normative, then the property it picks out is authoritatively normative too. With this in mind, what is the realist to say about the normativity of the pair of properties that are the intensions of AUTHORITATIVE OUGHT and OUGHT*? There appear to be only three salient hypotheses:

H1 Both properties are equally authoritatively normative for anyone
H2 Only one of these properties is authoritatively normative
H3 Our property is authoritative for us, their property is authoritative for them

H1 is ruled out by one of the core assumptions about authoritative normativity introduced in Section 1: that the authoritative ought in *uniquely authoritative;* that is, that there cannot be distinct conflicting authoritative concepts. But AUTHORITATIVE OUGHT and OUGHT* are conflicting concepts, which are imagined to have extremely similar downstream roles in deliberation. So if each picked out equally authoritatively normative relations, then it would be deeply puzzling what resources a realist could have for insisting that one of these concepts was authoritative while the other was not.

Next consider H2. In general, there is nothing especially puzzling about a concept with a particular role in thought failing to pick out a given property, relation, or kind. For example, consider a person who has a concept that plays a lot of the roles in their thought that WATER tends to play in our thought. Most of us find it highly plausible that whether this person is able to think about water with their concept can depend on their environment. (For example, whether this environment contains water, or instead a distinct, doppelganger substance.) But intuitively, this is not so for authoritatively normative concepts. To see this, suppose for simplicity that AUTHORITATIVE OUGHT latches onto an authoritatively normative property, but OUGHT* does not. This means that our OUGHT*-deploying Twins are failing to have authoritatively normative thoughts. And this suggests in turn that our very ability to have authoritatively normative thoughts depends on a kind of environmental luck that is prima facie absurd. Imagine trying to resolve an apparent dramatic conflict between what morality and prudence demand of you. And stipulate that there is an authoritatively normative ought-relation that holds among your options. You try to think about this relation by deploying a concept with the tendencies in thought of AUTHORITATIVE OUGHT. But—because of some fact about your environment—you simply cannot think about any authoritatively normative relation. That is very hard to believe.[31]

Finally consider H3. This is also unacceptably bizarre. For it implies that which relations are authoritatively normative for me depend upon facts about what causally regulates certain of my concepts. According to H3, in a world where a different property causally regulates certain of my concepts, I can ipso facto be subject to different authoritative norms. But it is hard to think of a more normatively irrelevant feature than which properties happen to causally regulate my thought.

Given that H1–H3 are unacceptable, the realist appears forced to reject the causal regulation account of reference-determination, for reasons that have nothing to do with data about disagreement.[32]

This section has argued that the representational determinacy challenge is not easily met. The contrast with LEGAL CHESS MOVE creates a presumptive explanatory burden. And the discussion of causal regulation has shown that the burden cannot by discharged simply by appealing to some "external" content

31. If it were possible, H2 would also raise a significant skeptical challenge (compare Eklund 2017: 14–15): how could we know whether our concept in fact latches onto an authoritative property? (I do not mean to endorse this skeptical challenge, only to note that it would be worth attention if we accepted H2).

32. Eklund (2017: Ch. 1) argues that the causal theory cannot satisfy what he calls the "ardent" realist, for reasons that have nothing to do with intuitions about genuine disagreement. My arguments are distinct from Eklund's, focusing not on the ardent realist's distinctive metaphysical hypothesis, but on what I take to be plausible realist hypotheses about the content of normative thought.

determination relation, like causal regulation. For the very features that appear to permit causal regulation to secure representational determinacy appear also to generate inconsistency with realist commitments about normative thought. Further, I have shown that locating the Moral Twin Earth argument within this broader challenge helps us to see that a prominent focus of the literature on that argument—on the conditions for genuine disagreement—do not cut to the heart of the deepest relevant challenge to realism.

5. Unifying the Challenges

We are now in a position to see why the deliberative authority challenge and the representational determinacy challenge are best understood as two faces of a single deep challenge to the normative realist—what I have dubbed the deliberation/ representation challenge.

The general form of this challenge is to square the tendencies in thought of the concept AUTHORITATIVE OUGHT with the core commitment of realism about this concept: that it represents a determinate relation. Each face of the challenge focuses our attention on one part of the data concerning AUTHORITATIVE OUGHT introduced in Section 1. The deliberative authority side of the challenge grants (for reductio) that normative thoughts of the form *I authoritatively ought to . . .* are ordinary beliefs, and challenges the realist to explain how such beliefs could play the role in deliberation that is part of our data about AUTHORITATIVE OUGHT. The representational determinacy challenge instead focuses on our data about the upstream tendencies in thought of this concept, and asks: how could a concept associated with such tendencies represent an acceptably determinate relation?

One reason it is valuable to see that the two challenges are unified is that it is crucial to the strongest form of each of these sides of the challenge that it tacitly assumes the desideratum at the heart of the other side of the challenge. We can see this by noting that if we abandon one of these desiderata, it becomes much easier to explain how we might be able to satisfy the other. Consider two examples.

First, if the realist were not committed to vindicating the deliberative authority data, a Boydian solution to the representational determinacy challenge might be entirely satisfactory. Boyd's theory arguably secures adequate representational determinacy. And, as I briefly mentioned in the preceding section, there is a family of promising replies to Horgan and Timmons's disagreement-based version of the objection, which appeal to the idea that Moral Twin Earth cases exhibit genuine but *non-moral* disagreement. As I explained there, this reply appears much less promising for the realist about authoritatively normative concepts. But that worry relies crucially on assumptions about the deliberative

authority of such concepts. The point was that the assumed deliberative authority of these concepts was precisely what made it uncharitable to interpret parties to such a dispute as engaging in non-canonical disagreements.

Second, consider someone uncommitted to representational determinacy. They could simply embrace the data about the role of AUTHORITATIVE OUGHT in thought, and face no obvious deliberative authority challenge. After all, as we saw in Section 3, the most promising way of understanding that challenge is precisely to explain how realism can be squared with this data about the role of this concept in our thought. But if we reject representational determinacy, and (for example) imagine a form of "realism" on which the intension of AUTHORITATIVE OUGHT is *extraordinarily* vague, it is no longer clear that there is anything puzzling about how the role in thought of this concept could explain this sort of representational content. (There are of course *many* other things that are puzzling about the imagined theory.)

Further, once we treat each side of the deliberation/representation challenge as appropriately intensified by the other, we can see that adequately addressing one side of the challenge would also address the other. For example, suppose that the realist had a compelling reply to the representational determinacy side: she could explain why a concept with the characteristic role in thought of AUTHORITATIVE OUGHT represents a determinate property. Then the deliberative authority side would be answered too. For the objector asked: how could a belief about some bit of reality have the sort of deliberative authority that AUTHORITATIVE OUGHT thoughts do? The reply would be: we have explained how the latter thoughts themselves can represents a determinate bit of reality; what more could you want?

We can illustrate this by considering an important recent instance of the deliberative authority challenge, due to Shamik Dasgupta. A key element of Dasgupta's challenge (in this case, to the non-naturalist) is that they must "play fair in naming" their alleged non-natural properties, and not simply stipulate or assume that they are normative (2017: 300ff.). Dasgupta does not consider a response that focuses on a theory of content determination. But suppose that Dasgupta's foil offered such a theory, and it solved the representational determinacy challenge. By doing so, they would, in effect, have shown that they were playing fair: for they would have shown that the non-natural property they have targeted is the one that we are talking about when we ask the relevant normative questions.

The general form of the unified challenge is to square the tendencies in thought of the concept AUTHORITATIVE OUGHT with the core commitment of realism about this concept: that it represents a determinate relation. The unified challenge is potent because on the one hand, it is natural to think that we need something more than the tendencies we find in normative thought in order to

explain the representational determinacy of that thought. But it is unclear why any such "something more" is required for a mental state to have the deliberative authority characteristic of AUTHORITATIVE OUGHT.

One way to illustrate the generality of the deliberation/representation challenge is that an important dilemma for constitutivism can be understood as an instance of this challenge. The dilemma is posed this way by David Enoch:[33]

> the more you pack into whatever it is you claim is constitutive of agency, the less plausible is the claim that it is so constitutive . . . . the less you pack into whatever it is you claim is constitutive of agency, the less by way of norms of practical reason you can extract from it. (2010: 213)

From the perspective of this paper, the constitutivist strategy Enoch addresses can be understood as one interesting strategy for meeting the unified challenge I have posed. That our normative concepts are appropriately related to what is constitutive of agency is hoped to explain the downstream tendencies of such concepts. And packing content into what is constitutive is supposed to achieve representational determinacy. Enoch's point is that these attempts pull against each other. And this is an instance of the general point that I have been emphasizing: that the two sides of the challenge seem quite generally to pull against each other.

Conclusion: The Significance of the Unified Challenge

I conclude this paper by explaining how the formulation of the deliberation/representation challenge that I have offered is illuminating. In short, it is illuminating in its unity, depth, and generality. I briefly illustrate each of these slogans in turn.

First, I have just emphasized the *unity* of the two sides of this challenge. To see the significance of this, note that leading proponents of the deliberative authority side of the challenge such as Korsgaard, Dreier, and Dasgupta do not mention content determination either as a constituent of the challenge as they understand it, or as a potential resource with which to address it. To put it flippantly, this means that they are prosecuting a challenge concerning the relation between mind and world without attending to the very sort of theory that would most plausibly allow us to sensibly link them.

Second, consider the *depth* of the unified challenge. I argued in Sections 3–4

33. The generalization is a bit pointed, *ad hominem*, since Enoch (2011) does not appear to be especially worried by either the deliberative authority challenge or the representational determinacy challenge.

that there were apparently promising replies to prominent instances of both constituent challenges, when they are addressed to moral concepts. These were the "more realism" reply to a formulation of the deliberative authority challenge that focused on vindicating Moral Rationalism, and the non-moral disagreement reply to Moral Twin Earth. But once we have formulated the general challenge, it becomes clear that these replies do not touch the deepest form of the relevant challenges.

The final way that the deliberation/representation challenge I have offered is illuminating is in its *generality*. There are at least two dimensions to this generality.

First, there is topical generality. Once we see the general form of the challenge I have developed for the normative realist, we can see that the moral realist faces a version of that general challenge too. As we saw in Section 1, although the concept MORAL OUGHT is in some ways a middle case between LEGAL CHESS MOVE and AUTHORITATIVE OUGHT, it appears *much* closer to the latter than the former. So we can frame a version of the general challenge that I have formulated that targets the moral realist. Suppose (just for example) that the moral realist appeals to causal regulation to solve the representational determinacy challenge. But with that solved, the deliberative authority challenge is exacerbated: if MORAL OUGHT thoughts are ordinary beliefs about whatever happens to causally regulate those beliefs, what explains their distinctive deliberative tendencies?

Recall from Section 1 the recent debates about whether structural rationality, epistemology, well-being, and the law are normative. One might suspect that these debates get going in part because the tendencies in thought of rationality, epistemic, well-being, and legal concepts are (like moral concepts) *more similar* to that of AUTHORITATIVE OUGHT than they are to LEGAL CHESS MOVE. If this is so, realism about such concepts may also face an important version of the unified challenge sketched here.

The second way the challenge is general concerns the range of metanormative views that the challenge targets. As we have seen, important instances of the two constituent challenges are often addressed to relatively narrow targets. For example, as I noted in Section 3, Dreier's version of the deliberative authority challenge (2015) targets the *reasons-fundamentalist*, which combines realism with a specific hypothesis about normative structure. But the general challenge as I have posed it applies to the realist whatever her account of normative structure. It *may* be that adopting some competitor to reasons-fundamentalism (such as value-, fittingness-, or rationality-fundamentalism, or a no priority view) would somehow help to address the challenge. But this is a matter of whether certain views possess helpful resources to address a quite general challenge, not a matter of which views face that challenge.

The same generality holds for the representational determinacy side of the challenge. Thus, the Moral Twin Earth argument was initially presented as a

challenge to Boyd in particular, and later extended to apply to all naturalistic moral realists (Horgan & Timmons 2000). But the deep challenge here applies more broadly yet, to anyone who takes our normative thoughts to have reasonably determinate representational content. For example, the problem is not made easier in any obvious way by a non-naturalistic hypothesis in the metaphysics of ethics. For it is not initially clear how positing non-natural properties could help to explain why the tendencies in thought associated in thought with AUTHORITATIVE OUGHT enables that concept to represent some such property. Indeed, a variant of the challenge holds for other cognitivist views in ethics, such as familiar forms of error theory and fictionalism: on these views, the thought that *I ought to do such-and-such* is assumed to have reasonably determinate representational content, even if reality itself may be empty of normativity. But this means that proponents of these views face the task of explaining how such representational determinacy is possible.[34]

In my view, this challenge, in its full generality, is the fairest formulation of the hardest problem for the normative realist (and indeed the cognitivist) in the philosophy of mind. Contrast the most familiar such challenge: that of accounting for so-called "judgment internalism." Judgment internalism is the thesis that there is some sort of necessary connection between making a moral (or normative) judgment and being motivated to act in accord with it. The problem with formulating a challenge via such a thesis is that there is no way of making judgment internalism usefully precise that is not intensely controversial.[35] It can thus appear open to the realist to shrug off the challenge.[36] It is much harder to deny or shrug off the sorts of data briefly sketched in Section 1 of this paper. And it is very hard to deny that the realist needs some explanation of this data.

My hope is that formulating the challenge in the general way I have here will enable progress in assessing the realist research program in metaethics. I see two important tasks ahead. First, clearly formulating a challenge can help to see what is required to *evade* it. To offer one stark example, one can imagine a version of cognitivism that evaded the challenge by simply denying representational determinacy. My own ambitions are to address, rather than evade this challenge. From this perspective, the challenge calls for a clear and constructive reply. Understanding whether this can be done, and what commitments the realist must take on to do so will put us in a much better position to evaluate the realist's overall research program.

34. A related point regarding moral fictionalism is the central theme of Hussain (2004). The challenge to these views is a *variant* of the one offered here, because they need not accept some of the realist commitments that I have shown to generate the full challenge.

35. For a recent survey, see Faraci and McPherson (2018).

36. For a model execution of this sort of strategy, see Enoch (2011: 247–259).

Acknowledgments

I am grateful for feedback on ideas that went into this paper to participants in my Spring 2018 Graduate Seminar, and to audiences at the Central APA, the Abu Dhabi Normativity and Reasoning Workshop, and the University of Missouri. I am also indebted to Christian Coons, David Faraci, Jamie Fritz, Josh Gert, David Plunkett, and two referees for *Ergo* for helpful discussion or comments on drafts of this paper. A draft of this paper was awarded honorable mention in the 2019 Sanders Prize in Metaethics essay competition.

References

Baker, Derek (2018). Skepticism about Ought Simpliciter. In Russ Shafer-Landau (Ed.), *Oxford Studies in Metaethics* (Vol. 13, 230–252). Oxford University Press. https://doi.org/10.1093/oso/9780198823841.003.0011

Bedke, Matthew S. (2014). A Menagerie of Duties? Normative Judgments Are Not Beliefs about Non-Natural Properties. *American Philosophical Quarterly*, *51*(189), 189–201.

Blackburn, Simon (1993). *Essays in Quasi-Realism*. Oxford University Press.

Boyd, Richard N. (1988). How to Be a Moral Realist. In Geoffrey Sayre-McCord (Ed.), *Essays on Moral Realism* (181–228). Cornell University Press.

Broome, John (2013). *Rationality through Reasoning*. Wiley-Blackwell.

Copp, David (1995). *Morality, Normativity, and Society*. Oxford University Press.

Copp, David (2000). Milk, Honey, and the Good Life on Moral Twin Earth. *Synthese*, *124*(1–2), 113–137.

Copp, David (2005a). Moral Naturalism and Three Grades of Normativity. In Peter Schaber (Ed.), *Normativity and Naturalism* (7–45). Ontos-Verlag. https://doi.org/10.1515/9783110327694.7

Copp, David (2005b). The Normativity of Self-Grounded Reason. *Social Philosophy and Policy*, *22*, 165–203. https://doi.org/10.1017/s0265052505052076

Cuneo, Terence, and Russ Shafer-Landau (2014). The Moral Fixed Points. *Philosophical Studies*, *171*(3), 399–443. https://doi.org/10.1007/s11098-013-0277-5

Darwall, Stephen, Allan Gibbard, and Peter Railton (1997). Toward Fin de siècle Ethics: Some Trends. In Stephen Darwall, Allan Gibbard, and Peter Railton (Eds.), *Moral Discourse and Practice* (3–47) Oxford University Press. Reprinted from *The Philosophical Review*, 101(1), 115–189. 1992. https://doi.org/10.2307/2185045

Dasgupta, Shamik (2017). Normative Non-Naturalism and the Problem of Authority. *Proceedings of the Aristotelian Society*, *117*(3), 297–319. https://doi.org/10.1093/arisoc/aox016

Dreier, James (2015). Can Reasons Fundamentalism Answer the Normative Question? In Gunnar Björnsson, Caj Strandberg, Ragnar Francén Olinder, John Eriksson, and Fredrik Björklund (Eds.), *Motivational Internalism* (167–182). Oxford University Press. https://doi.org/10.1093/acprof:oso/9780199367955.003.0009

Dunaway, William, and Tristram McPherson (2016). Reference Magnetism as a Solution to the Moral Twin Earth Problem. *Ergo*, *3*(25), 639–679. https://doi.org/10.3998/ergo.12405314.0003.025

Eklund, Matti (2017). *Choosing Normative Concepts*. Oxford University Press.
Enoch, David (2010). Schmagency Revisited. In Michael Brady (Ed.), *New Waves in Metaethics* (208–233). Palgrave.
Enoch, David (2011). *Taking Morality Seriously*. Oxford University Press.
Enoch, David (2018). Non-Naturalistic Realism in Metaethics. In Tristram McPherson and David Plunkett (Eds.) *Routledge Handbook of Metaethics.* (29–42). Routledge. https://doi.org/10.4324/9781315213217-2
Faraci, David, and Tristram McPherson (2018). Ethical Judgment and Motivation. In Tristram McPherson and David Plunkett (Eds.) *Routledge Handbook of Metaethics* (308–323). Routledge. https://doi.org/10.4324/9781315213217-20
Fletcher, Guy (2019). Taking Prudence Seriously. In Russ Shafer-Landau (Ed.), *Oxford Studies in Metaethics* (Vol. 14, 70–94). Oxford University Press. https://doi.org/10.1093/oso/9780198841449.003.0004
Foot, Philippa (1958–59). Moral Beliefs. *Proceedings of the Aristotelian Society, New Series*, *59*, 83–104.
Foot, Philippa (1972). Morality as a System of Hypothetical Imperatives. *Philosophical Review*, *81*(3), 305–316. https://doi.org/10.2307/2184328
Gibbard, Allan (2002). Reply to Hawthorne. *Philosophy and Phenomenological Research*, *64*(1), 179–183.
Gibbard, Allan (2003). *Thinking How to Live*. Harvard University Press.
Graff, Delia (2002). Shifting Sands: An Interest-Relative Theory of Vagueness. *Philosophical Topics*, *28*(1), 45–82. https://doi.org/10.5840/philtopics20002816
Horgan, Terence and Mark Timmons (1991). New Wave Moral Realism Meets Moral Twin Earth. *Journal of Philosophical Research*, *16*, 447–465. https://doi.org/10.5840/jpr_1991_19
Horgan, Terence and Mark Timmons (1992). Troubles for New Wave Moral Semantics: The 'Open Question Argument' Revived. *Philosophical Papers*, *21*(3), 153–175. https://doi.org/10.1080/05568649209506380
Horgan, Terence and Mark Timmons (2000). Copping Out on Moral Twin Earth. *Synthese*, *124*(1–2), 139–152.
Hussain, Nadeem (2004). The Return of Moral Fictionalism. *Philosophical Perspectives*, *18: Ethics*, 149–187.
Jones, Karen and Francois Schroeter (Eds.) (2018). *The Many Moral Rationalisms*. Oxford University Press. https://doi.org/10.1093/oso/9780198797074.001.0001
Kieswetter, Benjamin (2017). *The Normativity of Rationality.* Oxford University Press.
Kolodny, Niko (2005). Why Be Rational? *Mind*, *114*(455), 509–563. https://doi.org/10.1093/mind/fzi509
Korsgaard, Christine (1996). *Sources of Normativity.* Cambridge University Press.
Korsgaard, Christine (1997). The Normativity of Instrumental Reason. In Garrett Cullity and Berys Gaut (Eds.), *Ethics and Practical Reason* (215–254). Clarendon.
Lewis, David (1996). Desire as Belief II. *Mind*, *105*(418), 303–313. https://doi.org/10.1093/mind/105.418.303
Lord, Errol (2018). *The Importance of Being Rational.* Oxford University Press.
McPherson, Tristram. (2018). Authoritatively Normative Concepts. In Russ Shafer-Landau (Ed.), *Oxford Studies in Metaethics* (Vol. 13, 253–277). Oxford University Press. https://doi.org/10.1093/oso/9780198823841.003.0012
Merli, David (2002). Return to Moral Twin Earth. *Canadian Journal of Philosophy*, 32(2), 207–240. https://doi.org/10.1080/00455091.2002.10716518

Moore, G. E. (1993). *Principia Ethica* (rev. ed.). Thomas Baldwin (Ed.). Cambridge University Press. Original work published 1903.
Nowell-Smith, P. H. (1954). *Ethics*. Penguin.
Olson, Jonas (2018). Error Theory in Metaethics. In Tristram McPherson and David Plunkett (Eds.) *Routledge Handbook of Metaethics* (58–71). Routledge. https://doi.org/10.4324/9781315213217-4
Parfit, Derek (2006). Normativity. In Russ Shafer-Landau (Ed.), *Oxford Studies in Metaethics* (Vol. 1, 325–380). Oxford University Press.
Parfit, Derek (2011). *On What Matters* (Vol. 2). Oxford University Press.
Plunkett, David and Timothy Sundell (2013). Disagreement and the Semantics of Normative and Evaluative Terms. *Philosophers' Imprint*, 13(23), 1–37.
Plunkett, David and Kevin Toh (Eds.) (2019). *Dimensions of Normativity*. Oxford University Press.
Railton, Peter (1986). Moral Realism. *Philosophical Review*, *95*(2), 163–207.
Railton, Peter (2003). Normative Force and Normative Freedom. In Peter Railton (Ed.), *Facts Values and Norms* (322–352). Cambridge University Press. Reprinted from *Ratio*, *12*(4), 320–353. 1999. https://doi.org/10.1111/1467-9329.00098
Scanlon, Thomas (1998). *What We Owe to Each Other*. Belknap.
Scanlon, Thomas (2014). *Being Realistic about Reasons*. Oxford University Press.
Schroeder, Mark (2007). *Slaves of the Passions*. Clarendon.
Silk, Alex (2018). Metaethical Contextualism. In Tristram McPherson and David Plunkett (Eds.), *Routledge Handbook of Metaethics* (102–118). Routledge. https://doi.org/10.4324/9781315213217-7
Smith, Michael (1994). *The Moral Problem*. Blackwell.
Stojanovic, Isidora (2018). Metaethical Relativism. In Tristram McPherson and David Plunkett (Eds.), *Routledge Handbook of Metaethics* (119–134). Routledge. https://doi.org/10.4324/9781315213217-8
Street, Sharon (2006). A Darwinian Dilemma for Realist Theories of Value. *Philosophical Studies*, *127*(1), 109–166. https://doi.org/10.1007/s11098-005-1726-6
Street, Sharon (2012). Coming to Terms with Contingency. In Jimmy Lenman and Yonatan Shemmer (Eds.), *Contructivism in Practical Philosophy* (40–59). Oxford University Press.
Tiffany, Evan (2007). Deflationary Normative Pluralism. *Canadian Journal of Philosophy*, *33*(Suppl.), 231–262. https://doi.org/10.1353/cjp.0.0076
Wedgwood, Ralph (2004). The Metaethicist's Mistake. *Philosophical Perspectives 18: Ethics*, 405–426.
Wedgwood, Ralph (2007). *The Nature of Normativity*. Oxford University Press.
Williams, Bernard (1985). Morality: The Peculiar Institution. In Bernard Williams (Ed.), *Ethics and the Limits of Philosophy* (174–197). Routledge.
Wodak, Daniel (2019). Mere Formalities. *Canadian Journal of Philosophy*, *49*(6), 828–850. https://doi.org/10.1080/00455091.2018.1433795
Woods, Jack (2018). The Authority of Formality. In Russ Shafer-Landau (Ed.), *Oxford Studies in Metaethics* (Vol. 13, 217–229). Oxford University Press.