1

The Fragmentation of Authoritative Normativity

Tristram McPherson and David Plunkett

Introduction

A prominent theme in recent work in metaethics has been an increasingly explicit focus on *normativity*. Some of this recent work involves discussion of normativity in a wide sense that includes, for example, the sort of "generic" or "merely formal" normativity (purportedly) possessed by the rules of board games and social clubs. The core focus, however, has tended to be on a (purportedly) "distinctively normative" or "fully normative" sort of normativity: one that (purportedly) is "authoritative" or carries "normative force".[1] Put roughly, were this kind of normativity to exist, it would settle what agents "really and truly" or "genuinely" should do (or think, feel, etc.).[2] For many philosophers, part of what makes morality distinctively worth investigating is that it appears to have some kind of intimate connection to this "authoritative" kind of normativity.[3]

Of course, the existence of authoritative normativity is controversial. Some philosophers argue that there are no authoritatively normative facts or properties.[4] Others argue that we cannot even make sense of the idea of

[1] For central uses of these different locutions, see, respectively, (Dorsey 2019), (Scanlon 1998, 19), (McPherson 2018), and (Enoch 2011, 2).

[2] Our way of making this broad cut between "generic" (or "merely formal") normativity and "authoritative" (or "robust") normativity, draws from (McPherson 2018) and (McPherson and Plunkett 2017), which in turn draw on (McPherson 2011).

[3] For example, many philosophers accept versions of "moral rationalism" that purport to link moral reasons (or requirements, etc.) to authoritatively normative ones. For some recent discussion of different versions of such theses, see the papers collected in (Jones and Schroeter 2018).

[4] For example, it is natural to impute this view to (Olson 2014) and (Streumer 2017).

Tristram McPherson and David Plunkett, *The Fragmentation of Authoritative Normativity* In: *Oxford Studies in Metaethics, Volume 19*. Edited by: Russ Shafer-Landau, Oxford University Press. DOI: 10.1093/oso/9780198911876.003.0001

"authoritative normativity".[5] In this chapter, we largely set aside these relatively familiar worries to discuss what we take to be an important but neglected possibility. This is that our thought and talk about authoritative normativity is *discordant* in the sense captured by the following thesis:

Discord:

(i) There are multiple, equally optimally apt candidate notions of "authoritativeness".

(ii) The norms picked out by different optimally apt candidates for "authoritativeness" can conflict with each other.

We do not endorse Discord in this chapter. Rather, we think that this thesis is important, and deserves sustained attention. To encourage such attention, we offer a prima facie argument for Discord, and explore its significance.

We proceed as follows. In §1.1, we clarify Discord and explain our core strategy for arguing for it. In §1.2, we implement this strategy by identifying two sorts of data that can be used to argue for a plurality of equally apt candidate notions of "authoritative normativity". In §1.3, we explain how taking Discord seriously matters for metanormative inquiry. In §1.4, we discuss the prospects of pushing back against Discord by denying one or the other of its two conjuncts. The first way is to defend a kind of *uniqueness* about authoritative normativity, according to which there is a single most apt candidate authoritativeness concept. The second way is to defend a kind of *harmony* among the apt candidate authoritativeness concepts, according to which (put roughly) the norms picked out by these concepts do not deeply conflict.

1.1 Authoritatively Normative Concepts and Discord

In this section, we introduce the assumptions about authoritatively normative thought and talk that we use to generate our prima facie case for Discord. We start by discussing some of the difficulties of offering an illuminating characterization of authoritatively normative concepts. We suggest that the most promising way forward is to offer a *functional* characterization of such

[5] The clearest example here is (Baker 2018). It is also possible to read (Tiffany 2007) as making this sort of claim. Both Baker and Tiffany draw on ideas from (Foot 1972) and (Copp 1997).

concepts. We illustrate this strategy by introducing the most developed extant functional characterization strategy, which one of us (McPherson) has recently put forward. This illustration in turn enables us to clarify Discord. It also paves the way for the following section, where we use McPherson's account as a foil in developing our prima facie case for Discord.

To begin, note that the term 'authoritative normativity' is tricky to theorize about because it is neither a paradigmatic folk term nor a paradigmatic technical term. We take it to be uncontroversial that the term itself is a piece of technical terminology introduced by one of the co-authors of this chapter.[6] However, that introduction was an effort to make clearer and more precise an idea that appears to be built into, or presupposed by, much existing normative thought and talk. For example, suppose that Jacob writes down a series of rules that apply to anyone who enters his treehouse. Many people feel like they intuitively "get" the contrast between what one should do *relative to Jacob's treehouse rules* and what one "really and truly" should do.[7] Talk of 'authoritative normativity' is an attempt to regiment and clarify the normativity involved in the second side of this contrast.

Part of the motive to clarify the notion of authoritativeness is provided by the fact that, despite the widespread feeling just mentioned, several careful participants in discussion about this topic doubt that there is any clear idea of "authoritativeness" to be had. For example, Philippa Foot once claimed that the relevant use of 'ought' "makes no sense at all".[8] And Derek Baker suggests that most attempts at relevant "clarification" of this use amount either to the use of metaphors that need interpretation, or to mere table-thumping.[9]

To begin to see the challenge involved in explaining the idea of authoritative normativity, consider the simple proposal that authoritative normativity is simply the normativity characteristic of what agents *ought* to do (or believe, or feel, etc.). On its own, this gloss won't do. This is because the word 'ought' that is used in this explication is context-sensitive.[10] And in many contexts, sentences of the form "S ought to do A" express

[6] (McPherson 2018).

[7] Note that this intuitive contrast is complicated by the apparent normative significance of a social norm's being "in force" (Woods 2018). Note also that many people feel a contrast between the *relative* normative importance of different norms (e.g. moral norms vs. etiquette norms vs. Jacob's treehouse rules), which suggests there is also an important question concerning how *relatively* authoritative norms are. We put the issue of relative authoritativeness to the side in this chapter, and focus on maximal authoritativeness.

[8] (Foot 1978/2003, 169).

[9] (Baker 2018, 234).

[10] As on the dominant view of the semantics of 'ought' developed by (Kratzer 2012). For discussion, see (Chrisman 2016, section 2).

non-authoritative normative content. For example, suppose one of the rules of Jacob's treehouse prescribes speaking in rhyme. And suppose one says, "Naftali ought to have spoken in rhyme" in a context where those rules are salient in conversation. It is natural to interpret this statement as expressing a non-authoritative normative thought: roughly, a thought about what Naftali ought to have done *relative to the treehouse rules*.

One can orient to "authoritative normativity" by talking about what one ought to do "simpliciter", or of what one "really and truly" ought to do—as we ourselves did in the Introduction. We could turn this orientation into an explanation of authoritativeness if 'ought' was semantically like 'flat'. But it is not. In the latter case, there are degrees of flatness, and talk of what is "really and truly" flat plausibly creates a contextual standard that is at least relatively close to perfect flatness. But there are not, in our ordinary ways of thinking or talking, underlying "degrees of oughtness" where the term 'authoritative ought' might pick out the perfect extreme. So while these locutions may help orient us to authoritativeness, they do not seem to help to illuminate its nature.

Derek Parfit arguably seeks to orient us to "authoritativeness" by talking about norms "in the reasons-implying sense" as opposed to the "rule-implying sense".[11] But, again, this orientation does not illuminate the nature of authoritativeness. This is because, like the term 'ought', the term 'reason' can plausibly be naturally used to express a non-authoritative concept. To see this, start with the standard gloss on "normative reasons" (the kind that matter here). This is that reasons "count in favor of" (or count against) certain actions, beliefs, attitudes, etc. (And, perhaps, some reasons interact in more complicated ways, such as by undercutting or changing the force of other reasons). The standard thought (put roughly) is that facts about reasons, so understood, combine to deliver normative verdicts, such as ones about what one ought to do, think, or feel. For our purposes here, the key issue is that, just as "ought" claims can be context-sensitive, so too it seems that claims about "reasons" can be too. For example, suppose that in offering legal advice, you "offer reasons" for a particular course of action. Your claims can be true in light of what we might call the relevant "legal reasons" (those that determine what you *legally* ought to do) without imputing anything about "authoritative reasons" (those that determine what you *authoritatively* ought to do). Plausibly, in introducing his terminology, Parfit is thinking of "reasons" in

[11] (Parfit 2011, §88).

what T.M. Scanlon infectiously dubbed "the standard normative sense".[12] But just as plausibly, "the standard normative sense" is another way of saying "the authoritative sense". This makes it clear that Parfit's gloss is assuming, rather than explaining, the idea of authoritative normativity.[13]

The general lesson of this discussion is the following. One way of thinking about the core notion of authoritativeness is as the notion of a distinctive *normative* asymmetry among norms. It is tempting to attempt to capture this asymmetry using normative words like 'ought', 'reason', 'good', etc. But as we have seen, this strategy faces a general problem: these normative words can be used smoothly to talk about structural relations within whatever sorts of normative standards are made salient in the context of use. Because of this, even if these words might help provide initial orientation to authoritatively normative thought, they are unpromising tools for fully illuminating what distinguishes the authoritatively normative concepts from "merely generic" normative concepts like GOOD CHESS MOVE.[14]

In the face of these difficulties, we find it most promising to seek to characterize authoritative normativity in terms of a distinctive functional role. It is easiest to illustrate this strategy by stepping away from authoritativeness per se, and focusing on a particular authoritatively normative concept, AUTHORITATIVELY NORMATIVE PRACTICAL OUGHT, and its use in the context of first-person deliberation. It is natural to think that the thought "I authoritatively ought to do A" *settles* deliberation, and its authoritativeness consists in this settling role.[15] If this thought is correct, the challenge of explaining the authoritativeness of this concept reduces to the challenge of providing an illuminating functional characterization of the "settling role" of this sort of thought.[16]

[12] (Scanlon 1998, 19). [13] For more on this point, see (McPherson 2018).

[14] Christine Korsgaard emphasizes that, by using certain "normatively loaded" words, philosophers can in effect confuse themselves into mistakenly thinking that they have answered what she calls the "normative question" (see Korsgaard 1996, 42–3). Put roughly, for Korsgaard, the "normative question" arises when agents confront situations where doing what morality requires is hard and ask why they should follow the demands of morality. We are making a structurally similar point to Korsgaard's, about how using certain language can provide the illusion of having illuminated what we mean by "authoritatively normative".

[15] Compare the functional role Gibbard proposes to associate with the locution of "the thing to do" in (Gibbard 2003). Cf. also (McPherson 2018, 254).

[16] It is attractive to think (as (McPherson 2018) discusses) that there is a family of authoritatively normative concepts: AUTHORITATIVE OUGHT, AUTHORITATIVE REASON, AUTHORITATIVE VALUE, etc. But the latter concepts, by their nature, do not play a settling role in thought. A nontrivial commitment of the broad strategy that we focus on here is that the settling-based account of AUTHORITATIVE OUGHT can form the heart of a unified account of all authoritatively normative concepts.

In "Authoritatively Normative Concepts", one of us (McPherson) proposes that we can clarify the relevant settling role by focusing on contexts of deliberation where inter-normative conflict is salient. McPherson's leading example is this:

Sticky Situation: You find yourself in a sticky situation. You conclude that morality requires you to stay and help, while prudence dictates that you take the money and run. Torn, you ask yourself: *given all of this, what ought I to do?*[17]

The idea is that, in posing the italicized question in this context, an agent is seeking to deploy an "ought"-concept with a distinctive functional role: the role of being distinctively suitable for settling this deliberative conflict. We can call this the *deliberative settling role*. McPherson takes the deliberative settling role to give us the best grasp on the notion of an authoritatively normative concept. Of course, this gloss leaves an obvious question: what does *distinctive suitability* amount to here? McPherson offers a constitutive account, according to which, put roughly, "distinctive suitability" amounts to being the constitutive success norm for the sort of deliberative activity the agent is engaged in, in cases like Sticky Situation.

Much could be said about McPherson's account. For the purposes of our argument here, we propose to simply accept that McPherson has a promising strategy for illuminating authoritative normativity. We will now clarify our central thesis—Discord—in light of this strategy. Recall this thesis:

Discord:

(i) There are multiple, equally optimally apt candidate notions of "authoritativeness".

(ii) The norms picked out by different optimally apt candidates for "authoritativeness" can conflict with each other.

The first conjunct of Discord says that there are multiple, *equally optimally apt* candidate accounts of "authoritativeness" in thought and talk. What is it for an account to be *apt*? We understand *aptness* here to be a matter of capturing the underlying pretheoretical ideas that lead people to think that they "get" the contrast between "authoritative" and "merely generic" norms.

[17] (McPherson 2018, 254). Similar cases have a long history in metaethics. One striking early example is Williams' discussion of the "deliberative ought" in (Williams 1965/1973).

We think of McPherson as attempting to provide (part of) an account of 'authoritativeness' that is apt in this sense. The first conjunct of Discord says, in effect: there are multiple such candidate accounts that have equal claim to best capture the pretheoretical ideas. The second conjunct of Discord then says that the norms picked out by these apt candidates can issue conflicting verdicts about the same objects (e.g. an option or state of affairs).

In what follows, we will frequently be talking about equally optimally apt candidate accounts of authoritativeness, and the norms that form the intensions of these candidates. For brevity, we will use:

(i) 'authoritativeness candidate' as a shorthand for an optimally apt candidate account of "authoritativeness" in thought and talk,

and

(ii) 'candidate authoritative norm' as a shorthand for a norm picked out by an authoritativeness candidate.

Our clarification of Discord enables us to contrast this thesis with more familiar ideas. For example, take Parfit's reading of Henry Sidgwick's "dualism of practical reason".[18] On the view that Parfit attributes to Sidgwick, there are two "incomparable" normative standards: the standard given by what is *impartially best*; and the standard given by what would be *best for the actor*. Parfit understands this incomparability in terms of a single underlying authoritatively normative currency: for Parfit's Sidgwick, the dualism consists in the fact that we always have *sufficient normative reason* to act in accordance with either the impartial or prudential standard, where normative reasons are understood as the unified currency of authoritative normativity.[19] This view domesticates the conflict: *ultimately*, it is permissible either to do what is impartially best, or to do what is best for oneself.

You can imagine a more radical variant of Sidgwick which suggests that there is a single authoritativeness candidate, but it applies equally well to the prudential and impartial norms. On this view, we can be required by one maximally authoritative norm to take a certain option, while another maximally authoritative norm forbids doing so. Unlike on Parfit's version of Sidgwick's view, this kind of view entails that there would be no "more ultimate" authoritative standard, according to which, taking the relevant option would be simply permissible.[20]

[18] (Parfit 2011), discussing (Sidgwick 1874/1981). [19] (Parfit 2011, Vol. 1, 144).

[20] It is plausible to read David Copp as endorsing a version of this view, in work such as (Copp 1997).

In contrast to both of these views, consider a version of Sidgwick's dualism that would exemplify Discord. On this version, there are two authoritativeness candidates. One of these candidates has as its intension what is impartially best, while the other has as its intension what would be best for the actor. Assuming the impartial norm can conflict with self-interest, this would be an example of Discord. The contrast with the previous possibility is that, here, there is conflict between authoritativeness candidates, rather than conflict within a single authoritativeness candidate.

1.2 A Case for Discord

In this section, we use McPherson's account as a foil, to develop our prima facie case for Discord. We begin by developing two strands of what we call the *discord data*. The first strand involves pointing to the apparent possibility of multiple authoritativeness candidates deriving from competing specifications of what we above called the "deliberative settling role". The second strand involves identifying alternative functional roles that are different from the deliberative settling role, which also seem to underwrite authoritativeness candidates. We then go on to use this discord data to offer our prima facie argument for Discord.

Our case for Discord can be initially motivated in the following way. As we mentioned above, we understand *aptness* as a matter of capturing the underlying pretheoretical ideas that lead people to think that they "get" the contrast between "authoritative" and "merely generic" norms. But, in general, our sense of what would provide the relevant contrast will be somewhat imprecise. So we might predict that if any account can capture the relevant contrast, there will be other, conflicting accounts that can do so equally well. Our strategy in this section is to make good on this motivation by illustrating how it can be pursued using McPherson's account as a foil. Some readers may be unsympathetic to McPherson's account. Those readers can treat the main arguments to come as exemplifying our general strategy for defending Discord.

1.2.1 Discord Data: Competing Specifications

In this subsection, we explore the idea that we might get multiple authoritativeness candidates by identifying competing specifications of what it is to

realize the "deliberative settling role", as McPherson understands it. We then briefly explain why we think a similar strategy will likely be fruitful when applied to some competitors to McPherson's account.

To apply our strategy to McPherson's account, we need to describe more of that account. McPherson offers a substantive proposal concerning what realizes the deliberative settling role: the constitutive success norms for the activity of *non-arbitrary selection* (where, e.g. selection by coin-flipping is a paradigm of arbitrary selection). McPherson homes in on the relevant notion of "arbitrariness" by offering a series of examples. And they suggest that these examples present a pattern that readers will be able to extend to novel cases. But note that, even if this is true, we are unlikely to be able to decide the "arbitrariness" of every novel case. And if we are not, this means that there are multiple candidate intensions for 'arbitrary', each of which is fully compatible with the pattern of judgments that McPherson uses to orient readers to the concept.

Now consider a schematic version of McPherson's central principle:

Constitutive-S: S practically ought$_i$ to do A in context of action C $=_{def}$ the constitutive success conditions for S's activity of non-arbitrary$_i$ selection concerning C require doing A.[21]

The key point is that for each of the candidate intensions (*i*) for "arbitrary" mentioned in the previous paragraph, we can adapt the rest of McPherson's account to generate an activity of *non-arbitrary$_i$* selection. We can then treat that activity as an input to the Constitutive-S schema, to generate an account of what an agent practically *ought$_i$* to do. Now, suppose that the features mentioned in McPherson's discussion are sufficient to entail that the concept PRACTICAL OUGHT is a candidate authoritative norm (i.e. a norm picked out by an authoritativeness candidate). Given that each concept PRACTICALLY OUGHT$_i$ shares these features, it is plausible that these concepts will be candidate authoritative norms as well. That is, we seem to potentially have a dense cloud of closely related but distinct candidate authoritative norms.

A different way of generating multiple authoritativeness candidates focuses on different ways that a norm might be related to the activity of deliberative settling. Constitutive-S focuses on the relation of being the

[21] Cf. (McPherson 2018, 270). Note, as in (McPherson 2018), we use '$=_{def}$' to mean that the material on the right-hand side of the operator is the definition of the material on the left-hand side of the operator.

constitutive success conditions for the activity. But other constitutivists have taken different relations to an activity to be the crucial ones for grounding normativity. For example, Christine Korsgaard emphasizes a link to a certain kind of "inescapability", while Paul Katsafanas provides an account that bottoms out in constitutive "aims" of an activity.[22] We can generate different candidate authoritative norms by amending Constitutive-S to replace mention of "constitutive success conditions" with one of these competing characterizations of constitutivity.

As we have noted, McPherson provides a substantive gloss on the deliberative settling role in terms of non-arbitrary selection. A final way to generate additional authoritativeness candidates appeals to alternatives to McPherson's general account of the deliberative settling role.[23] First, *à la* Ralph Wedgwood, we might think of *settling on doing A* as a state linked to intention to do A by an "enkratic" norm of structural rationality.[24] The precise character of "enkratic" structural norms are controversial. In light of this, we can imagine a range of slightly different authoritativeness candidates individuated in part by inferential links to intention via different "rational" principles. Second, *à la* Allan Gibbard, we might think of the state of settling on doing A as fully constituted by *planning* to do A.[25] There are competing ways of understanding what "planning states" are. Each such competitor can be used to formulate an authoritativeness candidate.

In each of the examples we have discussed, we have the outlines of an argument that multiple possible normative concepts are equally apt candidates to count as "authoritative". The general strategy illustrated by these examples is to focus attention on *what it is* in a plausible account that is supposed to ground the "authoritativeness" of certain normative concepts, and then note that the relative plausibility appears likely to be shared by a number of other (subtly different) concepts.

1.2.2 Discord Data: Contrasting Functional Roles

As we emphasized at the outset, our leading foil (McPherson 2018) explains authoritative normativity in terms of a functional role: namely, providing

[22] See (Korsgaard 2009, 1), (Katsafanas 2013, 39), and (McPherson 2018, 260–1).

[23] The general thought that there might be multiple plausible specifications of the relevant inferential role is closely related to what (Eklund 2017, §3.5) calls the "embarrassment of riches" problem.

[24] Compare (Wedgwood 2007).

[25] Compare (Gibbard 2003).

first-personal deliberative guidance. Building on this idea, McPherson then proposes that the paradigmatic "authoritative" norm is the constitutive norm for an activity appropriately tied to this functional role. We just showed how we can generate a plurality of apparent authoritativeness candidates based on this functional role. In this section, we consider ways of generating authoritativeness candidates by highlighting alternatives to that functional role. The key motivation for this subsection is that providing first-person deliberative guidance is only one among multiple functional roles that intuitively "authoritatively" normative concepts play in our lives.[26] To see this, consider the following examples.

First, consider what we might call the "evaluation role". Just as we can make "ought" claims relative to any number of standards, we can evaluate an act or state of affairs in any number of ways—as best *for me*, as *most polite*, as *morally best*, as best *according to the aristocratic ideal*, etc. But it is natural to ask, when given an act that is morally admirable, but impolite and imprudent, if it is good *overall or simpliciter*. In parallel with McPherson's account, we can think of the functional role of the concept AUTHORITATIVELY GOOD as being to settle one's *evaluative perspective*.

We can make two points about a concept that fills this "evaluation role" (if there is one). First, suppose that we draw an initial, intuitive cut between the "authoritatively" normative and the non-authoritative, "merely generic" norms. Jacob's treehouse rules (as such) seem merely generic, while PRACTICAL OUGHT seems to pick out normative facts that are authoritative. The norms which tell us what is "good simpliciter" intuitively land on the authoritative side as well.

The second point is that settling evaluation is different from settling deliberation. One way to illustrate this point is to notice that a consequentialist and one sort of non-consequentialist could seemingly agree in their ethical evaluations. They might agree, for example, that in Judith Jarvis Thomson's "transplant case", the world would be authoritatively better if an unscrupulous surgeon were to carve up one healthy person in order to save five by distributing their victim's organs.[27] But they might nonetheless disagree about what the doctor authoritatively ought to do.

Next, consider what we might call the "*fittingness* role" for a norm. Consider mental states such as being angry at Alice, or believing that water is wet. It is a now-familiar thought that we can ask whether such

[26] For connected thoughts, see (Plunkett 2020) and (Southwood 2016).
[27] *À la* (Thomson 1976), developing an example from (Foot 1967, 10).

states are "fitting", "merited", or "appropriate", where the answer to this sort of question is not, on its face, a matter of whether one "authoritatively ought" to be in that state.[28] For example, suppose that your employer is bigoted, vindictive, and perceptive. Given that they are perceptive, they will notice if you are indignant at their bigotry, and given their vindictiveness, they will make your life miserable if they notice. Given this situation, it may well be that you ought to prevent yourself from feeling indignation. But, despite this, such indignation remains a fitting response to their bigotry. One way of thinking about fittingness conditions is that they again have a distinctive functional role: they are the constitutive correctness conditions for the mental states that they apply to.

We will briefly mention three other examples of potentially distinctive normative functional roles. First, consider a sort of "first-personal endorsement" role, which tracks what follows from an agent's "normative perspective".[29] There are also two prominent sorts of "second-personal" functional roles for norms: governing the appropriateness of advice, and of second-personal objection.[30] And, famously, there are norms of structural rationality, which govern our psychologies.[31] In each case, it seems possible that we can identify a distinctive functional role that the relevant normative concepts play in our psychology, which might seem authoritative in a way that contrasts markedly with the functional role "allowing me to track what will get me into trouble in Jacob's treehouse club."[32]

Notice two things about the examples just sketched. First, in all of these cases, the relevant sorts of normative concepts are frequently put forward as normatively weighty or significant, in a way that contrasts with, for example, the requirement to put on a baseball hat when entering Jacob's treehouse. That is, prima facie, these are all authoritativeness candidates. Second, the

[28] For discussion, see the essays collected in (Howard and Rowland 2023).

[29] (Southwood 2013).

[30] For emphasis on the (purported) "second-personal" nature of many normative concepts (including, purportedly, our moral ones), see (Darwall 2006).

[31] For a recent discussion that defends the distinctiveness of structural rationality, see (Worsnip 2021).

[32] There are also further (more controversial) candidate roles one might discuss here. For example, Southwood (2016) argues for the importance of a distinct "hypological" role—roughly the role of providing the basis for apt *criticism*. Or, consider the question of how practical judgments about "what to do" relate to normative judgments about what one should do. Some suggest that they amount to the same thing (see (Gibbard 2003)) while others argue that they are importantly distinct (see (Hieronymi 2009)). If they are distinct, it might well be that there is a further candidate role tied to "authoritativeness" connected to the former, distinctively "practical" role of "what to do" judgments.

norms picked out by these concepts can seemingly conflict. For example, for someone with bad normative views, structural rationality can seemingly require them to do things that (in fact) they authoritatively ought not to do. The bigoted boss case above illustrates that similar conflicts can seemingly arise between fittingness norms and the "authoritative ought".

1.2.3 A Prima Facie Case for Discord

We have just put forward a set of prima facie "authoritativeness"-constituting functional roles for normative or evaluative concepts. In §1.2.1, we sketched a set of prima facie ways for there to be significant competing *specifications* of the deliberative guidance role. A similar case for competing specifications could be sketched for each sort of normative functional role just discussed in §1.2.2. At first blush, these examples are each authoritativeness candidates that are relevant to defending Discord. This is because these examples suggest that there are a range of distinct and seemingly equally apt candidates for what it is that makes a normative concept *authoritative*.

This "discord data" doesn't by itself establish that Discord is true. It just provides the basis for an initial case for it. To understand this case, contrast the following three theses about the authoritativeness candidates:

Discord:

(i) There are multiple authoritativeness candidates.

(ii) The norms picked out by different authoritativeness candidates can conflict with each other.

Harmony:

(i) There are multiple authoritativeness candidates.

(ii) The norms picked out by different authoritativeness candidates cannot conflict with each other.

Uniqueness:

There is a unique maximally apt candidate notion of authoritativeness.

With these three options made salient, we can state the prima facie case for Discord by considering the implications of the discord data for the plausibility of Uniqueness and Harmony, respectively.

First, consider Uniqueness. The discord data presents a straightforward reason to doubt Uniqueness. Begin with the competing specifications data. Suppose, as part of that data suggests, that there is a cloud of candidate specifications of the deliberative guidance role. What could make it the case that *just one* of them is the "real" deliberative guidance role, suitable to characterize the nature of authoritatively normative concepts? Next, consider the contrasting functional roles discussed in §1.2.2. As we emphasized at the outset, it is attractive to follow McPherson in hoping that identification of a distinctive functional role might help to illuminate what we are even trying to communicate in talking of the "authoritatively" normative. But we might think of the different functional roles sketched above as providing different *flavors* of authoritativeness. Insofar as these are equally apt candidates for authoritativeness, this suggests strong reasons to reject Uniqueness.

To see the force of this point, consider a possible objection. Various contemporary philosophers defend reasons-, evaluation-, and fittingness-*first* accounts of authoritative normativity.[33] This fact might seem to suggest that the discord data simply reflects conflicting *substantive* views about what it is to be an authoritatively normative concept.

We think there is clearly something to this line of thought. However, we are skeptical that, as stated, it does much to undermine the prima facie case for Discord. Arguably the "x-first" debates just mentioned simply presuppose Uniqueness. But, as we have seen, there are plausibly distinctive functional roles intimately connected to fittingness, evaluation, and reasons/ought. Given their connection to distinctive functional roles, it is unclear what, if anything, would warrant our taking them to be competing candidates for the fundamental basis of a single authoritatively normative structure.

Next, consider a prima facie case from the discord data against Harmony. For ease of exposition, we will focus on contrasting functional roles to make the case, though we think a parallel case could be made by focusing on competing specifications of a given functional role instead. If one grants that there really are distinctive normative functional roles, then Harmony is seemingly hard to maintain. For example, it is extremely plausible that the attitude it is fitting to have can come apart from the attitude it would be best overall to have (indeed, we used such a possibility to illustrate the idea of a fittingness norm, in the case of the bigoted boss). And this seems like a real

[33] For examples of each approach, see, respectively, (Schroeder 2021), (Maguire 2016), and (Chappell 2012). For doubts about widely shared assumptions in this debate, see (Wodak 2020).

conflict: one might be torn about whether to have the fitting attitude, or the one that it is best to have.

Similar points can be made with respect to the evaluative and practical perspectives. Several non-consequentialist philosophers have followed Philippa Foot in thinking that consequentialists have a "compelling idea", which remains compelling even after one has rejected the theory.[34] The functional pluralism we have proposed suggests an attractive version of this idea: whatever is true about 'ought' or 'permissible', part of what is compelling is the apparent authoritativeness of the evaluative functional role itself.

That's hardly a conclusive case against either Uniqueness or Harmony. As we discuss briefly in §1.4, much more can be said on behalf of each of these views. Our aim in this chapter, however, isn't to fully argue against these views and in favor of Discord. In our view, Discord is almost invisible in recent discussions of authoritative normativity. Our aim here is to make the case that it is worth taking seriously as a theoretical option, given the discord data. For now, we take ourselves to have done this. With that in mind, we now explore some of the broader philosophical consequences of this initial case.

1.3 Implications

In this section, we consider some implications of our discussion. We begin by exploring the implications of the discord data for the debate between proponents and "skeptics" of the idea of authoritative normativity. We then consider four sorts of implications that the truth of Discord itself would have for contemporary metanormative inquiry.

1.3.1 The Discord Data and "Skepticism" about Authoritative Normativity

In our view, contemporary debates over whether the very idea of authoritative normativity makes sense can be significantly improved by attention to the discord data. Consider two examples.

[34] (Foot 1985, 198).

First, defenders of authoritative normativity sometimes appear to presuppose something like Uniqueness.[35] However, it is striking that their arguments do not typically do anything to establish Uniqueness over alternatives like Discord or Harmony. To illustrate, consider again McPherson's argument in "Authoritatively Normative Concepts". That argument emphasizes the distinctive normative significance of a certain functional role tied to first-personal deliberation. But nothing in that paper even considers the sorts of alternative functional roles sketched above, or the potential case for Discord that they motivate. Or consider Daniel Wodak's recent attempt to defend the idea of authoritative normativity, by arguing for fictionalism about thought and talk about "merely formal" norms.[36] On this view, to say that one is required to do something, relative to the rules of Jacob's treehouse, is like saying that Sherlock Holmes is a detective. However, note that it would be a much harder stretch to go fictionalist about, for example, our thought and talk about fittingness or authoritatively evaluative normative concepts. The arguments of these papers thus simply do not address the discord data, and its potential significance for the very possibility of authoritative normativity.[37]

Now consider those skeptical about authoritative normativity (such as Derek Baker). Such skeptics typically argue for the position that (drawing on Evan Tiffany) we can call "deflationary normative pluralism", according to which there are a plurality of different kinds of normative facts, but none that are authoritatively normative.[38]

Such skeptics often argue in part by denying that we can even *understand* what it would be for one sort of norm to be "normatively supreme" or to "trump" another.[39] A central burden facing such arguments is that it seems very plausible that, for example, the moral norms in some relevant sense "trump" the rules written in a kid's treehouse. So these arguments have the familiar vices of philosophical arguments for the unintelligibility of what appear to be plausible thoughts.

It is notable that, like their opponents, many skeptics about authoritative normativity have seemed to simply assume that authoritativeness and

[35] For example: "Our concept of an authoritatively normative property seems to require a kind of uniqueness not required by our concept of a generically normative [property]." (McPherson 2020, 499).

[36] (Wodak 2018).

[37] Similar points also apply to the primitivism about (authoritative) normativity suggested in (Dorsey 2019).

[38] (Tiffany 2007).

[39] See (Baker 2018), drawing on ideas from (Copp 1997).

Uniqueness go together. For example, this assumption seems built into much of the language that Baker uses in putting forward his skeptical case. Making this assumption explicit suggests a potentially more promising strategy for arguing for deflationary normative pluralism. The strategy can be glossed as follows. First, argue that Uniqueness is a presupposition of "authoritative normativity" talk. Second, use the discord data to argue that it is impossible to satisfy this presupposition. A key virtue of this strategy is that it makes the alleged lack of an asymmetry in normative authority between morality and the treehouse rules a theoretical *consequence* of a case for "deflationary normative pluralism", not a putative intuitive *motivation* for it.

1.3.2 Metanormative Implications of Discord

In this subsection we briefly sketch four ways that taking seriously the possibility that Discord is true illuminates underexplored possibilities in metanormative inquiry.

First, Discord may suggest a new way of theorizing profound "normative dilemmas". Some philosophers are skeptical about the possibility of genuine normative dilemmas, according to which (put roughly), agents ought to do two things, but can't do both at once. One reason for this skepticism is the thought that there is something in the nature of morality or practical reason that is inconsistent with the very idea of normative dilemmas.[40] However, even if there cannot be *moral* dilemmas, or dilemmas within the sort of practical reason constitutively tied to the nature of first-person deliberation, there might nonetheless be *authoritativeness* dilemmas. For example, perhaps an agent could be "authoritatively" pulled in different directions qua deliberator, evaluator, and possessor of a mental state.

Second, Justin Clarke-Doane has recently argued for a "new open question argument".[41] Very roughly, Clarke-Doane claims that, given the plurality of possible normative standards, no conclusion about what follows from such standards can settle what to do. A natural objection to Clarke-Doane would be to argue that this is not true of the *authoritative* normative standard, by the very nature of this sort of standard. However, if Discord were true, a more modest version of Clarke-Doane's conclusion might hold:

[40] See (Sayre-McCord Manuscript) for connected discussion.
[41] (Clarke-Doane 2020, §6.5).

it might be an open question which authoritativeness candidate to follow, in cases of conflict.

Notice that Discord is compatible with the possibility that, in many circumstances we might confront, there will be no conflict among the candidate authoritative norms. After all, Discord is an existential, not universal claim: it states that there are *some* deep conflicts between the relevant norms, not that such norms conflict in every case. In the cases where there are no conflicts among authoritative norms, it seems we should be able to use "authoritative dominance" reasoning. That is, in these cases, we may be able to conclude that one option is superior to (or at least not worse than) another, according to *all* of the authoritative norms. Similar dominance reasoning is sometimes invoked in the context of epistemic uncertainty about normative matters.[42] If Discord were true, dominance reasoning could play a key role in practical deliberation even in cases without normative uncertainty. This is one important way in which Discord is different from "deflationary normative pluralism". According to the latter view, any given option will always be required by many normative standards, and forbidden by many others, so the sort of dominance reasoning sketched above will never be available.

Finally, one strand of our (preliminary) case for Discord appealed to *functional* pluralism. Functionally-based Discord provides an interesting motivation for the possibility of a kind of *metanormative* fragmentation. On the view that we favor, the overarching metanormative project can be glossed as follows: to explain how actual normative thought and talk—and what (if anything) such thought and talk is distinctively about—fits into reality.[43] Now focus on the thought and talk—and what (if anything) that such thought and talk is distinctively about—that is tied to the different sorts of functional roles we have discussed. Perhaps different metanormative theories (e.g. expressivism, non-naturalism, naturalistic realism, error theory, etc.) might be true of some of these types of normative thought and talk (and the parts of reality, if any, that it is about) but not others.

Consider one example to motivate this idea. Influential arguments for a surprising range of metaethical views appeal to the distinctive character of first-person deliberation. Think of David Enoch's deliberative indispensability argument for non-naturalism in *Taking Morality Seriously*, Christine Korsgaard's case for her "constructivism" in *The Sources of Normativity*, or

[42] As in (Ross 2006).

[43] See (McPherson and Plunkett 2017).

Allan Gibbard's central use of the deliberatively settling "thing to do" locution in defending expressivism in *Thinking How to Live*.[44] Suppose that one of these forms of argument is sound, or at least on the right track. If Discord is true, this might nonetheless fail to establish the relevant view as the correct view about normative thought and talk tied to evaluation or fittingness.

The possibilities we have sketched in this section underscore ways that Discord, and the discord data that motivates it, complicate and potentially significantly transform fundamental debates in contemporary metaethics. Attention to the discord data opens up new terrain for theorizing about normativity (and our thought and talk about it). Given that issues about authoritative normativity show up not only in debates about ethics, but also in other areas (such as epistemology, political philosophy, philosophy of law, and aesthetics), exploration of this terrain matters for metanormative inquiry quite generally.

1.4 Resisting Discord

In this chapter thus far, we have made a preliminary case for Discord. We then explored some implications of taking seriously either Discord or the data that we have used to motivate that thesis. Our aim in the chapter is not to establish the truth of Discord. Rather, it is to put this thesis on the table for serious discussion. With that in mind, we use this section to briefly survey several important strategies for resisting Discord without abandoning belief in authoritative normativity. We will begin by sketching strategies for defending Uniqueness, and then turn to strategies for defending Harmony.

1.4.1 Strategies for Defending Uniqueness

Here, we briefly introduce five strategies for defending Uniqueness.

We begin with what we take to be an extremely natural strategy for pushing back against the functional pluralism data: to insist that there just really is *one* functional role that provides a maximally apt candidate for characterizing the authoritatively normative concepts. Unsurprisingly, we

[44] See (Enoch 2011), (Korsgaard 1996), and (Gibbard 2003).

think that the most promising candidate here is the deliberative settling role. Consider how someone with this perspective might view the fittingness role. Such a person might grant that fittingness facts are important, insofar as they tell us about the well-functioning of some of our *parts* (e.g. certain of our attitudes). But the *real* normative question, such a person might insist, arises in the context of first-person deliberation. And within that context, she might reason, sometimes it makes sense to embrace or bring about the ill-functioning of some of one's parts, in order to act as one should.

One limitation of this strategy for defending Uniqueness is that it does nothing to address the specification discord data. And indeed, by itself, this strategy seems helpless to do so. This is because of a point that we made in §1.2.1: it is implausible that our grasp on the nature of authoritative thought is fine-grained enough to put us in a position to rule out competing specifications.

A second strategy is possible, if our class of discordant norms has some extensional overlap. In this case, it might seem possible to "construct" the concept of a uniquely maximally authoritative norm. Here is the idea. Suppose (for simplicity) that there are two authoritativeness candidates, A and B, and they each pick out a norm as authoritative (ought_A and ought_B respectively). Now suppose you say: the uniquely optimal candidate is the one that picks out a norm $\textit{ought}_S$ such that one ought_S to perform an act just in case you both ought_A and ought_B to do so.[45]

There are three natural kinds of worry about this strategy. First, it only works if there is a non-empty intersection of the intensions of all of the authoritativeness candidates. Second, if the intersection is very limited, picking out only a few cases, one might think that it doesn't do much to alleviate worries about the conflict between ought_A and ought_B, given that they would still conflict in almost all cases. More deeply, one might wonder whether one had really identified a "most authoritative" standard, or simply done something else. For example, perhaps one has just found a convenient label for the intersection points where the conflict among the authoritativeness candidates is not practically relevant.

A third strategy is especially salient for responding to the specification data. Here, one might be optimistic about appealing to "worldly help" to defend Uniqueness. It is a familiar thought that, in many domains where we embrace some kind of "realism", the determinacy and stability of the

[45] Thanks to Shyam Nair for illuminating discussion of possibilities closely related to the one we describe here.

content of our thought and talk are explained partly by the world itself. For example, on this sort of view, the word 'water' can stably pick out H_2O, despite the imprecision and inconstancy of our uses of this word. Analogously, even if we can't ourselves be more determinate about what is involved in "authoritative" normative thought, maybe one possible specification of "non-arbitrariness" in fact "carves nature at its normative joints", and this specification plays a role in explaining what the uniquely authoritative practical ought facts are. We think there is much to be explored in this line of thought.[46] But it is far from obvious that this move can work. To appreciate why, we can note that some philosophers—including ones seemingly otherwise inclined toward normative realism—seem moved by the possibility that the joints of nature could simply cross-cut the authoritative norms.[47] The idea, in short, is that "joint-carvingness" is one thing, and normative "importance" another. If that is right, it poses a deep challenge to this strategy for defending Uniqueness.

A fourth strategy is to defend Uniqueness as a kind of brute primitive fact. Imagine that we think of authoritativeness as a non-natural higher-order property of norms. Then it might simply brutely adhere to a norm that can be thought about via a precise specification of a single functional role.

A fifth strategy is to defend a mysterian view here that is analogous to mysterian views in the philosophy of mind.[48] That is, one might argue that perhaps there is an explanation of why a certain sort of normative fact is uniquely authoritative, but our limited human minds simply cannot grasp what that explanation is.

An important challenge for proponents of the third, fourth, and fifth strategies is this: even if Uniqueness *could* be explained in one of these ways, what basis do they provide for thinking that Uniqueness is actually true, as opposed to a mere possibility?

1.4.2 Strategies for Defending Harmony

We now introduce interesting strategies for defending Harmony.

The first strategy argues that, although there are multiple authoritativeness candidates, they each have non-overlapping domains of application

[46] For a detailed discussion of using this sort of strategy to explain the semantic stability of normative terms, see (Dunaway and McPherson 2016).

[47] (Eklund 2017, §2.3). [48] For example, see (Stoljar 2006).

and thus never come into conflict. To see how this might go, consider a simplified example. Suppose (just for illustration) that there are two sorts of candidate authoritative norms: epistemic norms and practical norms. Suppose further that one can have epistemic requirements only to believe, disbelieve, or suspend judgment. And suppose finally that one can never have a practical requirement *to believe* that P (or to disbelieve or suspend that P), but only one to *bring it about that* one forms the belief P (etc.).[49]

One might respond to this strategy by pointing out that there is still a potentially distressing sort of tension involved in finding oneself (epistemically) required to believe that P, while also being (practically) required to bring it about that you do not. We could address this sort of worry by appealing to a second sort of strategy.

This strategy begins by arguing that we can secure Harmony if we can show that the sort of tension that can hold *among* candidate authoritative norms can also arise *within* a highly unified normative system. To illustrate the latter sort of tension, consider the thesis of "global consequentialism".[50] Suppose, for simplicity, that the correct evaluative standard is *happiness maximization*. According to the global consequentialist, we can directly evaluate each of an agent's acts by this standard. But we can also use this standard to directly evaluate their *dispositions*, *deliberative tendencies*, and everything else that we might want to evaluate. Given this variety of direct evaluative targets, global consequentialism can entail that an agent *ought to perform* a certain action (e.g. to lie to their friend), but at the same time, *ought to be robustly disposed not to perform* actions of that sort (e.g. because being such that one would be prepared to lie to one's friend in such a context will tend to have bad consequences in other contexts).

Now return to the alleged practical/epistemic tension introduced just above. The proponent of the second strategy might argue that this "tension" is just the same sort of tension that global consequentialism can entail regarding acts and dispositions. If such a tension can be produced by a unified normative theory like global consequentialism, it might seem like it does not show a theoretically *deep* disharmony among the authoritative norms.

A third way that one might seek to argue for Harmony is to argue for explanatory relations among the candidate authoritative norms. For example,

[49] For a related claim, compare (Shah 2003, 463), who states that "we cannot recognize the relevance of [non-evidential] factors in our deliberative reasoning about what to believe."

[50] (Pettit and Smith 2000).

suppose (again just for simplicity) that our two candidate authoritative norms are the epistemic and practical norms. And suppose that the facts about the epistemic norms are grounded in facts about the practical norms. This might seem to suggest that, even if there can be conflicts between the epistemic and practical norms, these conflicts are not relevantly deep.

Notice that all of these strategies appear, at best, helpful only for addressing the functional pluralism data. In the case of the first two strategies, this is because it is very plausible that any case of specification pluralism will involve candidate authoritative norms with overlapping domains of direct application. In the case of the last one, this is because it is implausible that, for example, in the case of two candidate specifications of 'arbitrariness', the facts picked out by one are grounded in the facts picked out by the other.

We think there is much more to be said about each of these strategies. It might well be that one (or more) of them can be made to work. But we hope to have shown why it is far from obvious that any of them will. Thus, we think there is a good prima facie case to be made for Discord, and that this thesis is worth taking seriously as a live theoretical possibility about the structure of authoritatively normative thought, which, perhaps, reflects something about the structure of normativity itself.

The discussion in this subsection arguably reveals a striking lack of clarity in what counts as "deep" conflict or harmony among a set of norms. That's because it isn't entirely clear which of the possibilities sketched in that section would involve "harmony" or "discord" among the relevant norms. We conclude this section by sketching three further dimensions to this unclarity.

The first concerns the correct individuation of "norms" as potential objects of conflict. To illustrate, consider moral norms. It is natural to think of moral norms as encompassing (at least) relations of moral obligation and permission, moral value, moral virtue, moral reasons, etc. For the purposes of thinking about "harmony" and "discord", does it make sense to think of these as a single unit (the "moral norms"), or as a collection of norms that might turn out to be in intramural discord? Depending on how one resolves this sort of question, certain views might get classified either as instances of Harmony or of Uniqueness, or even potentially as a form of Discord (e.g. perhaps in cases of normative dilemmas of the relevant kind).

A second dimension of unclarity concerns the significance for Harmony and Discord of different dimensions of *how much* norms can conflict. Consider three examples. First, we can imagine candidate authoritative norms conflicting subtly or robustly in a given case (e.g. there is slightly

more reason_a to do A than there is reason_b, vs. A being ideal_a but apocalyptically bad_b). Second, we can imagine candidate authoritative norms conflicting in a wider or narrower range of cases. Third, we can imagine conflict arising in common vs. rare vs. merely metaphysically possible cases.

Finally, consider a dimension of unclarity tied to the fact that we defined Harmony and Discord in terms of there being multiple, *equally optimally apt* candidate notions of "authoritativeness". Now consider the case where this condition just *barely* fails to be satisfied. That is, a case where we think there is only one optimal candidate notion, but at least one *ever so slightly less apt* candidate.

How should we think of classifying this multidimensional range of cases with respect to Harmony vs. Discord? We are unsure. We think it is natural to find the hypothesis of Discord distressing. One reason is broadly practical: it might seem to threaten informed, normatively sensitive agents with paralyzing conflicts in their normative thinking. Another reason is broadly theoretical: in short, one might think that (i) if reality contains authoritative normativity, it would need to be the kind of normativity that is uniquely authoritative or (ii) that any normative concept worth thinking of as "authoritative" would be clearly separated from "non-authoritative" normative concepts (and not, say, just be a *slightly better* candidate for "authoritativeness"). We are not sure how cleanly these different concerns map onto the different theses we have been exploring. One possibility is that there is not one distinctively interesting cut between Harmony and Discord. Rather, it might be that once we accept that there is a plurality of apt candidates (or near-candidates) for "authoritativeness", there are a variety of possible sorts of tension that different people might find more or less gripping or worrisome.

1.5 Conclusion

A common assumption in much of the literature on "authoritative normativity" is that there is a univocal kind of concept that we are thinking and talking about with the sort of language that gestures toward authoritativeness. Arguably, this assumption is also common in much of the thought and talk that this literature aims to understand. The guiding idea of this chapter is that thought and talk connected to "authoritative normativity" might be more fragmented than this assumption suggests. In this chapter, we have

offered a preliminary argument that there is a plurality of discordant candidate authoritative norms, each picked out by equally optimally apt candidate accounts of "authoritativeness" in thought and talk. With this case on the table, we have sketched some implications of the possibility of discordance for metanormative inquiry. We have also sketched what we take to be a range of promising ways of resisting our preliminary argument.

At the end of §1.4.2, we suggested that some of these ways of pushing back reveal that it is unclear how to most fruitfully characterize the notions of Discord and Harmony that we used to frame our discussion. We take it that the idea that the authoritative norms might be "fragmented" or "discordant" is significant in large part because it will strike many of us as unsettling. The unclarities we have discussed suggest that it is not easy to puts one's finger on exactly why this is so. We suspect that getting clearer on why the idea of the fragmentation of authoritative normativity is unsettling can help us better understand our aspirations in thinking about authoritative normativity. We thus hope that this chapter spurs further work investigating both the credibility of Discord, and its philosophical significance.[51]

Works Cited

Baker, Derek. 2018. Skepticism about Ought Simpliciter. In *Oxford Studies in Metaethics*, vol. 13, edited by R. Shafer-Landau. Oxford: Oxford University Press, 230–52.

Chappell, Richard Yetter. 2012. Fittingness: The Sole Normative Primitive. *Philosophical Quarterly* 62 (249): 684–704.

Chrisman, Matthew. 2016. *The Meaning of 'Ought': Beyond Descriptivism and Expressivism in Metaethics*. Oxford: Oxford University Press.

Clarke-Doane, Justin. 2020. *Morality and Mathematics*: Oxford: Oxford University Press.

Copp, David. 1997. The Ring of Gyges: Overridingness and the Unity of Reason. *Social Philosophy and Policy* 14 (1): 86–106.

Darwall, Stephen L. 2006. *The Second-Person Standpoint*. Cambridge, MA: Harvard University Press.

[51] Thanks to Mark Budolfson, David Faraci, Joshua Gert, Sean Kim, Devontae Lacasse, Zachary Lang, Dustin Locke, Shyam Nair, Hille Paakkunainen, Adrian Russian, Karl Schafer, Thomas Schmidt, Nicholas Southwood, Zoe Thierfelder, two anonymous referees, and participants in discussions of earlier versions of this chapter at the University of Vermont, the 2022 Madison Metaethics Workshop, and Humboldt University.

Dorsey, Dale. 2019. On Distinctively Normative Norms. *Journal of Moral Philosophy* 16 (4): 414–36.

Dunaway, Billy, and Tristram McPherson. 2016. Reference Magnetism as a Solution to the Moral Twin Earth Problem. *Ergo: An Open Access Journal of Philosophy* 3.

Eklund, Matti. 2017. *Choosing Normative Concepts*. Oxford: Oxford University Press.

Enoch, David. 2011. *Taking Morality Seriously: A Defense of Robust Realism*. Oxford: Oxford University Press.

Foot, Philippa. 1967. The Problem of Abortion and the Doctrine of the Double Effect. *Oxford Review* 5: 5–15.

Foot, Philippa. 1972. Morality as a System of Hypothetical Imperatives. *Philosophical Review* 81 (3): 305–16.

Foot, Philippa. 1978/2003. Postscript to "Morality as a System of Hypothetical Imperatives". In *Virtues and Vices*. Berkeley, CA: University of California Press, 169–73.

Foot, Philippa. 1985. Utilitarianism and the Virtues. *Mind* 94 (374): 196–209.

Gibbard, Allan. 2003. *Thinking How to Live*. Cambridge, MA: Harvard University Press.

Hieronymi, Pamela. 2009. The Will as Reason. *Philosophical Perspectives* 23 (1): 201–20.

Howard, Christopher, and R. A. Rowland. 2023. *Fittingness*. Oxford: Oxford University Press.

Jones, Karen, and François Schroeter. 2018. *The Many Moral Rationalisms*. Oxford: Oxford University Press.

Katsafanas, Paul. 2013. *Agency and the Foundations of Ethics: Nietzschean Constitutivism*: Oxford: Oxford University Press.

Korsgaard, Christine M. 1996. *The Sources of Normativity*. New York: Cambridge University Press.

Korsgaard, Christine M. 2009. *Self-Constitution: Agency, Identity, and Integrity*. Oxford: Oxford University Press.

Kratzer, Angelika. 2012. *Modals and Conditionals: New and Revised Perspectives*. Oxford: Oxford University Press.

Maguire, Barry. 2016. The Value-Based Theory of Reasons. *Ergo: An Open Access Journal of Philosophy* 3.

McPherson, Tristram. 2011. Against Quietist Normative Realism. *Philosophical Studies* 154 (2): 223–40.

McPherson, Tristram. 2018. Authoritatively Normative Concepts. In *Oxford Studies in Metaethics*, vol. 13, edited by R. Shafer-Landau. Oxford: Oxford University Press, 253–77.

McPherson, Tristram. 2020. Ardent Realism without Referential Normativity. *Inquiry* 63 (5): 489–508.

McPherson, Tristram, and David Plunkett. 2017. The Nature and Explanatory Ambitions of Metaethics. In *The Routledge Handbook of Metaethics*, edited by T. McPherson and D. Plunkett. New York: Routledge, 1–25.

Olson, Jonas. 2014. *Moral Error Theory: History, Critique, Defence*. Oxford: Oxford University Press.

Parfit, Derek. 2011. *On What Matters*. Oxford: Oxford University Press.

Pettit, Philip, and Michael Smith. 2000. Global Consequentialism. In *Morality, Rules and Consequences: A Critical Reader*, edited by B. Hooker, E. Mason, and D. Miller. Edinburgh: Edinburgh University Press, 121–33.

Plunkett, David. 2020. Normative Roles, Conceptual Variance, and Ardent Realism about Normativity. *Inquiry* 63 (5): 509–34.

Ross, Jacob. 2006. Rejecting Ethical Deflationism. *Ethics* 116 (4): 742–68.

Sayre-McCord, Geoff. On a Theory of a Better Morality. Manuscript.

Scanlon, T. M. 1998. *What We Owe To Each Other*. Cambridge, MA: Harvard University Press.

Schroeder, Mark. 2021. *Reasons First*. Oxford: Oxford University Press.

Shah, Nishi. 2003. How Truth Governs Belief. *Philosophical Review* 112 (4): 447–82.

Sidgwick, Henry. 1874/1981. *The Methods of Ethics*. Indianapolis, IN: Hackett Publishing Co.

Southwood, Nicholas. 2013. "The Thing To Do" Implies "Can". *Noûs* 47 (3).

Southwood, Nicholas. 2016. Does "Ought" Imply "Feasible"? *Philosophy and Public Affairs* 44 (1): 7–45.

Stoljar, Daniel. 2006. *Ignorance and Imagination: The Epistemic Origin of the Problem of Consciousness*: New York: Oxford University Press USA.

Streumer, Bart. 2017. *Unbelievable Errors: An Error Theory about All Normative Judgments*: Oxford: Oxford University Press.

Thomson, Judith Jarvis. 1976. Killing, Letting Die, and the Trolley Problem. *The Monist* 59 (2): 204–17.

Tiffany, Evan. 2007. Deflationary Normative Pluralism. *Canadian Journal of Philosophy* 37 (5): 231–62.

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

Williams, Bernard Arthur Owen. 1965/1973. Ethical Consistency. In *Problems of the Self*. Cambridge: Cambridge University Press, 166–86.

Wodak, Daniel. 2018. Fictional Normativity and Normative Authority. *Canadian Journal of Philosophy* 49 (6): 828–50.

Wodak, Daniel. 2020. Who's On First? In *Oxford Studies in Metaethics*, vol. 15, edited by R. Shafer-Landau. Oxford: Oxford University Press, 49–71.

Woods, Jack. 2018. The Authority of Formality. In *Oxford Studies in Metaethics*, vol. 13, edited by R. Shafer-Landau. Oxford: Oxford University Press, 207–29.

Worsnip, Alex. 2021. *Fitting Things Together: Coherence and the Demands of Structural Rationality*. New York: Oxford University Press.