

Authoritative Normativity

Tristram McPherson (Ohio State) mcpherson.164@osu.edu
and
David Plunkett (Dartmouth) david.plunkett@dartmouth.edu

Version of June 4, 2024
Forthcoming in Oxford Handbook of Meta-Ethics
Eds. David Copp and Connie Rosati

**Please do not cite, circulate, or quote without permission from an author.
Please quote, cite, and circulate final published version.**

Introduction

Human life is permeated by normative and evaluative thought and talk. For example, people regularly invoke moral, legal, and aesthetic norms, and the rules of sports or games. It is notable that some norms seem to *matter normatively* more than others. For example, many think that moral, political, epistemic, and aesthetic norms matter (much) more than the rules of a kid's treehouse club. This sort of apparent asymmetry in "normative importance" is arguably part of what motivates many philosophers to work in moral and political philosophy, epistemology, and aesthetics, rather than studying the norms of children's clubs.

In this paper, we provide a critical overview of recent discussions about one framework for thinking about this apparent asymmetry in "normative importance". This is a framework that we've advocated for and explored in a number of our recent papers.¹ This framework appeals to a (purported) distinction between two kinds of normativity. On the one hand, all of the above norms are normative in a *generic* sense: put roughly, they all are standards that something (e.g. an action) can conform to or not.² On the other hand, at least some of these norms fail to be *authoritative*: that is, (roughly) they fail to concern what we "really and truly" should do, think, or feel.

One might seek to motivate the distinction between "generic" and "authoritative" normativity in a variety of ways. One approach begins by thinking about deliberation in contexts where (at least a perceived) conflict among norms is salient to the agent. Consider an example from one of us (McPherson):

¹ See, (McPherson 2011), (McPherson 2018), (McPherson and Plunkett 2017), (Plunkett and Shapiro 2017), and (McPherson and Plunkett 2024).

² Compare to (Railton 1999/2003, 323), (McPherson 2018), (Plunkett and Shapiro 2017), and (Copp 2004).

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?³

Intuitively, when you ask this question in Sticky Situation, you don't seem to be trying to think about a norm that is just one among many. Rather – at least in certain ways of spelling out the example – you seem to be reaching for a norm that has *transparently* distinctively normative importance, in a way that (at least plausibly) not even morality and prudence do. We might think of this sort of example as giving us a useful intuitive grasp of the relevant notion of authoritativeness.⁴

As this example suggests, we can potentially use the idea of “authoritative normativity” to theorize about various types of asymmetries in “normative importance”, of the kind that we gestured to above. For example, the final question in Sticky Situation is intended to be directly about an authoritatively normative relation. By contrast, if you think that you always authoritatively ought to do what you are morally required to do, this belief suggests that moral facts – even if not authoritative in themselves – might play an especially important role in fixing what the authoritatively normative facts are.

In what follows, we put forward some of our own views about the topics we address. However, our main aim in this paper is not to convince readers of these views. Rather, it is to provide a critical introduction to the recent literature on “authoritative” normativity (which, in other work, we’ve also called “robust” normativity), and how it contrasts with “generic” normativity (which in other work, we’ve also called “formal” normativity).⁵ In addition, we aim to introduce neglected issues on this topic for further philosophical exploration.

We divide our work into four sections. In the first section, we briefly introduce how we think about “generic” normativity. In the second section, we discuss a range of common ways of glossing the idea of “authoritative normativity”. We argue that these common glosses can be helpful in orienting to the basic idea of authoritative normativity, but can't offer much illumination beyond that. The third section introduces theories that aim to be more informative than such common glosses, including both ones that aim to undermine the idea of the distinction between

³ (McPherson 2018, 254). See also (Wedgwood 2004, 406), (Cullity and Gaut 1997), and (Williams 1965/1973).

⁴ Some people here might not like the idea of thinking of morality and prudence as separate normative standards (as imagined in this example), as opposed to (for example) just considerations of certain kinds within the authoritatively normative. If so, we could then just swap in other standards to make the same point: e.g., focusing on a perceived conflict between legal norms and the rules of a social club.

⁵ See, for example, (McPherson 2011) and (Plunkett 2019). Note that, as will become clear in our discussion, the relevant notion of “authoritative” normativity that we explore is distinct from the kinds of issues about “practical authority” (and its justification) that is an important other strand of ethical theory, as in (Raz 1979/2002). Put roughly, that kind of authority concerns the authority that some agents have in directing the actions of other agents.

authoritative and generic normativity described above, and ones that aim to vindicate it. In the fourth section, we discuss a range of further important questions about the apparent contrast between generic and authoritative normativity that we think have been underexplored. These include the following: Is the “authoritativeness asymmetry” discrete (“on/off”) or gradable? Is there just one “authoritativeness” asymmetry, or a plurality of such asymmetries? Is the idea of a distinction between authoritative normativity and generic normativity part of our ordinary ethical thought and talk, or not?

In the fourth section, we also explore the relationship between “authoritativeness asymmetries” at the “object-level” of normative reality (which includes, e.g., such things as normative facts, properties, and relations, insofar as they exist), and at the “representational-level” of normative thought and talk (which includes, e.g., such things as our normative concepts and words). This contrast matters, because we might give different theories of the two levels. For example, one sort of error theory would take there to be authoritatively normative thought, but no authoritatively normative reality for it to be about. In some cases below, we are explicit about whether an idea (or theory, etc.) about authoritativeness is best understood or developed at the “representation-level” or at the “object-level”, or both. However, in other cases (as in the introduction, above), we elide this distinction for the sake of brevity.

1. Generic Normativity

Our main aim in this article is to explore contemporary discussion of authoritative normativity. To get a handle on much of this discussion, it’s useful to start with a bit more on the kind of normativity we’ve claimed it can be usefully contrasted with: namely, merely “generic” normativity.

In the introduction, we claimed that generic norms are standards that something (e.g. an action) can conform to or fail to conform to. We are here using the terms ‘norm’ and ‘normativity’ in *very* wide senses of these terms, where (for example) ‘norms’ here covers everything ranging from the instructions in cookbooks to the rules of a board game to moral norms.

Take any given object: a baseball; a pattern of linguistic use; or the content of a belief. We can then imagine any number of different norms related to this object. These might be norms that prescribe as follows: make new baseballs that conform to the shape of this baseball, or engage in speech that conforms to this pattern of linguistic use. Or the norms might prescribe the opposite of such things, or any number of more complicated things, such as: conform to the first part of this linguistic pattern, but then not this part, and then change your pattern of use in such-and-such ways starting next week. We can think of these norms as abstract objects. On the view we favor, there are *many* such abstract objects. For example, there is a norm that says “wink every thirty seconds for the next week” and one that says “always chew gum when reading Derek Parfit”.

This brief sketch of generic normativity obviously leaves much open. For example: it leaves open questions about how exactly to understand what it means for norms to be “abstract objects”. Further, we want to be open to alternative ways of understanding generic normativity that avoid commitment to abstracta. So much more could be said about generic norms as such, even if one grants the sketch we’ve given so far. Moreover, there are, of course, many things one might want to contest even in our brief sketch. However, we think this sketch is enough to get discussion of authoritative normativity off the ground.

2. Orientations

In this section, we explore three common strategies that philosophers have used to orient to the idea of authoritative normativity (or of our thought and talk about it): appeal to certain vivid metaphors, to key structural notions within the normative, and to familiar normative subcategories, respectively.⁶ We grant that these strategies can all provide useful orientation, at least in certain contexts. However, we will argue that beyond that, they fall short. The metaphors don’t provide much illumination, and the structural notions and subcategories cross-cut the authoritative/generic distinction.⁷

2.1 Metaphors

Some philosophers talk about normative or practical “oomph”.⁸ The idea is that authoritative norms (insofar as they exist) are somehow normatively *potent* in a way that merely generic norms are not. Other philosophers talk about authoritative norms as “trumping” other norms, drawing on the idea that one card can trump another.⁹ These are metaphors. Talk of “oomph” and “trumping” suggest some sort of asymmetry. But what does that asymmetry consist in? These metaphors do not tell us. Presumably it is an asymmetry in “authoritativeness”. But that asymmetry is just what we are seeking to understand. So these metaphors do not seem to help to illuminate it.

2.2 Structural Notions

⁶ Here we assume, for brevity, that these strategies are all intended to illuminate a single target “authoritativeness” phenomenon. In some cases, another possibility is that they are intended to illuminate different phenomena, which are nonetheless related to “authoritativeness” in some relevant way. This latter option is especially important if it turns out that, as we discuss later in this paper, that there isn’t a single “authoritativeness” phenomenon, but rather a range of related phenomena that are all equally as good candidates for being “authoritative” in a relevant sense.

⁷ In making this argument, we draw on previous arguments from (McPherson 2018). For connected discussion, see also (Baker 2017).

⁸ See especially (Joyce 2006, 62ff). Joyce himself claims that his talk of “practical oomph” is “usefully vague”. For a different use of “practical oomph,” which concerns motivational efficacy, see (Bedke 2019).

⁹ See especially (McLeod 2001). McLeod notes that there are “difficult questions” about what the relevant trumping power “consists in”.

A second approach appeals to structural relations that can apply to or between norms. This approach initially appears more promising. However, extant structural proposals tend to appeal to relations that cross-cut the contrast between authoritative and merely generic normativity. Here, we consider two prominent examples.

First, one might draw inspiration from Immanuel Kant, and think that we can shed light on authoritativeness by appealing to the distinction between categorical and hypothetical imperatives.¹⁰ Simplifying greatly, a hypothetical imperative says: *if you want this, then do that!* By contrast, a categorical imperative simply says: *do that!* Hypothetical imperatives are *escapable* in an important sense: if you do not have the relevant desire, then the imperative simply fails to apply to you. So one might hope that authoritativeness can be understood in terms of categoricity.

Philippa Foot showed that this idea cannot be right.¹¹ Consider our paradigm non-authoritative norm: the rules that 7-year-old Lyra has set for her treehouse club. Suppose that one of these rules states that one must wear a hat at all times in the treehouse. And suppose that you pop your head into the treehouse to announce that dinner is ready, not wearing a hat. Do you care at all about conforming to the rule? You do not. Have you violated it? You have. This is because, as stated, the treehouse club rule is categorical in form. Yet, it seems not to be authoritative in the least. There might, of course, be ways of modifying what one understands “categoricity” to involve, to help it do more work here. But at least on straightforward (and relatively standard) ways of thinking about it, it fails to isolate something distinctive about authoritative normativity, or our thought and talk about it.¹²

Second, consider the idea that we can understand the authoritative ought as the “all things considered” ought. This is an initially appealing idea. If I tell you that if you consider *x* and *y*, you ought to stay and help, but if you consider *all the facts*, you ought to take the money and run, the latter claim seems more authoritative. After all, it looks like the significance of the first “ought” claim depends on simply ignoring relevant facts, whereas the latter one does not. However, as with categoricity, the problem is that this distinction cross-cuts the authoritative vs. merely generic distinction. To illustrate, consider again the paradigmatically non-authoritative treehouse club rules. Facts about what is “treehouse-justified” in light of *all* the facts (including, for example, facts about what morality requires) are not authoritative in the least.

¹⁰ (Kant 1785/2012). (Mackie 1977, 27–30), who can be read as denying the existence of authoritative normativity, seems to identify the sort of normativity he finds implausible with categoricity.

¹¹ (Foot 1972, 308–309).

¹² Note that, in addition to categoricity not being a sufficient condition for making a norm “authoritative”, it is also unclear whether it is a necessary one. There seems nothing incoherent about a broadly “Humean” approach to authoritative normativity, according to which the fundamental authoritative norms are all hypothetical imperatives. Consider here the kinds of Humean ethical theories offered by (Schroeder 2007) and (Street 2012), which can plausibly be read along these lines.

2.3 Appeal to specific normative properties or concepts

A third natural strategy is to seek to illuminate authoritative normativity (or our thought and talk about it) by appealing to some specific normative property or concept. Here, we consider three examples: appeal to “ought”, to “rational”, and to “reasons” (either at the level of properties, or concepts). Here again, the initial problem is that the would-be explanantia appear to cross-cut the authoritative/generic contrast. And attempts to precisify the explanantia to avoid this problem appear to presuppose the very thing they are intended to explain.

First, consider “ought”. In Sticky Situation, introduced above, our protagonist is imagined to ask the question “what *ought* I to do?”. This phrasing might suggest that authoritative normativity is the sort of normativity associated with “ought”. The problem with this suggestion is that people regularly use “ought” claims in a wide range of different contexts, in relation to a wide range of different norms. For example, we make claims both about what one *morally* ought to do, as well as claims about what you ought to do, according to the etiquette standards around here.

This observation is reflected in the most influential semantics for the word ‘ought’, due to Angelika Kratzer. On Kratzer’s account, ‘ought’ is highly context-sensitive, allowing it to express an *extremely wide* range of functions from propositions to truth values.¹³ Put roughly, on Kratzer’s view, in a given context of use, the intension of ‘ought’ is fixed by two parameters: a *modal base* (roughly: the set of possible worlds against which the “ought” claim will be evaluated) and an *ordering source* (very roughly, a way to rank those worlds, to produce truth-conditions for sentences). For example, in Sticky Situation, we might imagine that a *prudential* ordering source ranks the relevant “take the money and run” worlds higher than the “stay and help” worlds. Given these assumptions, the sentence “In Sticky Situation one ought to take the money and run” will be true in a context where prudence is the standard that fixes the ordering source.

If we accept Kratzer’s account of ‘ought’, we cannot understand authoritative normativity simply as the sort of normativity expressed by “ought” claims. To see this, consider the sentence “One ought always wear a hat in the treehouse”, uttered in a context where Lyra’s treehouse norms are salient in a way that fixes the ordering source and modal base for ‘ought’. Then this sentence will be true, despite it’s being false that one always *authoritatively* ought to wear a hat in the treehouse. If Kratzer’s semantics is correct, then, talking about you “ought” to do is not always just a way of talking about a subset of what you authoritatively ought to do.¹⁴

Kratzer’s theory is compatible with the idea that there is an ordering source for “ought” that picks out what one *authoritatively* ought to do. But this means that we need some further theory to understand the nature of the authoritative ought.

¹³ See (Kratzer 2012). For connected discussion, see (Chrisman 2016).

¹⁴ For connected discussion, see (Wodak 2018b).

One might insist that we have the ability to linguistically induce the relevant ordering source using certain words. For example, we might seek to get to the authoritative ordering source by talking about what one “really and truly” ought to do, or what one ought to do *simpliciter*. While such glosses might be helpful in providing orientation (as we think they often are), we do not think they are more deeply illuminating about the nature of authoritative normativity. To see this, contrast a context-sensitive term like ‘flat’: here, there is arguably a privileged “perfectly flat” standard that we can attempt to make salient with the adverbs “really and truly...” or “...simpliciter”. But it is not plausible that the semantics of ‘ought’ can be modeled this way.

We see a similar pattern of orientation without deeper illumination in appeals to other specific normative properties or concepts. Consider two further illustrative examples: rationality and reasons.

Start with *rationality*.¹⁵ One might think that: you can (at least often) rationally ignore Lyra’s treehouse rules, but you can never rationally ignore (e.g.) moral requirements. One might further take this contrast to suggest that authoritativeness consists in something like rational significance. However, there is a two-part challenge to the usefulness of this proposal. First, the term ‘rational’ is used in a dizzying variety of ways in contemporary philosophy, ranging from ways that tie it narrowly to some sort of psychological coherence, to ways that simply use it as a proxy for authoritativeness.¹⁶ It is implausible that all of these uses pick out an authoritatively normative relation.¹⁷ Second, the uses that simply treat it as a shorthand for authoritativeness are potentially misleading, given the varied other connotations of the term. So, in order to provide a substantive explanation of why (a certain way of understanding) rationality is authoritative, one seemingly needs some independent grip on the idea of authoritativeness.

Next, consider (normative) *reasons*. Derek Parfit famously contrasts normativity in the “reasons-implying” sense with normativity in the “rule-implying” sense.¹⁸ Similarly, a prominent debate in contemporary ethical theory concerns whether a sort of “moral rationalism” is true, where that rationalism is often glossed as the thesis that one has reasons to do what morality requires.¹⁹ It is very natural to interpret Parfit and the moral rationalism debate to be using ‘reasons’ as a way of talking about authoritative normativity. But this proposal faces the same core challenge as the appeal to ‘ought’ or ‘rational’.

¹⁵ For an example where appeal to “rationality” seems to be used to try to explain (or capture) “authoritative normativity”, consider Allan Gibbard’s use of ‘rational’ in (Gibbard 1990), where he writes things such as “what it is rational to do settles what to do”. (Gibbard 1990, 49).

¹⁶ For discussion (including of some of the dangers for normative inquiry that arise from this fact, and parallel ones about other core pieces of normative terminology), see (McPherson and Plunkett 2020).

¹⁷ These points are relevant to debates about the “normativity” of rationality, as in (Kolodny 2005).

¹⁸ (Parfit 2011, §88).

¹⁹ For discussion, see (Jones and Schroeter 2018).

The standard contemporary way of understanding normative reasons is as *contributory* entities: ought-facts are explained, it is often said, by the balance of reasons. The basic idea is that certain reasons “count in favor” of certain actions, beliefs, attitudes, etc., while others count against them, and (perhaps) others interact with each other in more complicated ways (e.g., by undercutting or changing the force of a given reason). As we have seen, the term ‘ought’ is context-sensitive. On one plausible way of regimenting terminology, it is natural to think that the *moral reasons* are the considerations that ground the facts about what one morally ought to do. And it is hard to see what is to stop us from talking in similar ways about (e.g.) “mafia” reasons, where these are the considerations that bear on what one ought to do according to the mafia standard.²⁰ So it doesn’t seem like talk of “reasons” by itself isolates something distinctively associated with authoritative normativity.

There is a natural next move here. T.M. Scanlon famously talks of reasons “in the standard normative sense”.²¹ By this phrase, he plausibly means to focus on something like the idea of *authoritative* normative reasons. If by “reasons-implying” Parfit means that *authoritative* reasons are implied (and not just any reasons, such as “mafia reasons” etc.), then his view may be extensionally adequate. But, in that case, Parfit is then just implicitly presupposing, rather than illuminating, the notion of authoritative normativity.²²

The upshot of what we’ve just argued is that appeals to “ought”, “rationality”, and “reasons” can’t significantly illuminate authoritative normativity. We can generalize from these examples to other specific normative concepts or properties that might initially seem like natural things to reach for when trying to explain authoritative normativity. For example, if one thought that VALUE or FITTINGNESS were the fundamental normative concepts, one might instead talk of normativity in the “value-implying” or “fittingness-implying” senses. But, again, we can ask whether we are talking about value (e.g.) in an authoritative sense, or a merely generic sense.

The examples in this section illustrate a general point. Attempts to illuminate authoritative normativity (or our thought and talk about it) by appeal to specific normative properties or concepts (or at least well-known, standard ones of the kind

²⁰ Note that even if people don’t regularly talk about those things using the terminology of “mafia reasons”, we certainly could: these things (understood in the way we just glossed) seem to exist, and our terminology of “reasons” is flexible enough to allow us to refer to them in this way.

²¹ (Scanlon 1998, 18).

²² At this stage, another move one might make is to appeal to the distinction between reasons of the “right kind” vs. ones of the “wrong kind”. Put roughly, this cut relies on the idea of “fittingness” or “warrant”: “right kind” reasons are ones that are “fitting” to have vs. “wrong kind” ones are ones that are not. (For example: many think that one has a “right kind” reason to believe something because it is true, but a “wrong kind” reason to believe something because doing so would make one wealthy). We think that this move won’t help, at least on standard ways of thinking about “fittingness” and the distinction between “right kind” and “wrong kind” reasons. (For more on these topics, see (Howard and Rowland 2023)). In short, this is because the relevant cut here is, once again, one that cross-cuts the distinction between authoritative and merely generic normativity. For connected discussion, see (Lord and Sylvan 2019).

we have discussed here) are unhelpful. This is because the authoritative/merely generic contrast appears to simply cross-cut the relevant properties or concepts. We can sensibly ask what is authoritatively valuable, what one's authoritative reasons are, or what one authoritatively ought to do. But we can also be focused on the (merely generic) norms of the card game bridge, and felicitously ask whether I have been dealt a good hand, what reasons my partner's bid gives me, and whether I ought to lead trump at this point in the game.

If one assumes that (one's use of) a familiar normative concept is authoritative, then one can be vulnerable to the illusion that it provides some grip on what authoritativeness is. And this can in turn obscure one's ability to see that there are deep questions about how to understand authoritativeness that one has left untouched.²³ Part of why we think it is valuable to use the technical terminology of "authoritative" normativity is to help avoid falling prey to this illusion.

3. Theories

We now turn to accounts of the idea of authoritative normativity (or our thought and talk about it) that plausibly survive the ground-clearing we did the preceding section. Put roughly, these theories aim to do something more than just orient us to the basic topic of authoritative normativity, and how it contrasts with merely generic normativity. Instead, they aim to do something more informative, such as provide a better understanding of the concept of AUTHORITATIVE NORMATIVITY, provide a real definition of the property of being authoritatively normative, or explain why the basic distinction between authoritative normativity and merely generic normativity is mistaken. And, importantly, they (at least initially) seem to offer more promise on accomplishing one (or more) of these things than those ideas we discussed in the second section. We focus our discussion on recent work that concerns the (purported) cut between authoritative and merely generic normativity (which may be either explicitly discussed using our terminology, or described using closely-related terminology), rather than trying to survey all the discussions throughout the history of philosophy that can be plausibly read as concerning the issues at hand. We begin with theories that aim to undermine the distinction, and then move to ones that aim to vindicate it.

3.1 Is the "distinction" nonsense?

One might take the lesson of the previous section to be that we should be suspicious of the apparent distinction between "authoritative" and "merely generic" norms. After all, as we have seen, many prominent discussions of this contrast appeal to metaphors, or to notions that simply cross-cut the distinction (if there is a distinction there at all). Philippa Foot famously dubbed the idea that there is a deep

²³ Compare Christine Korsgaard's point that different people treat different words as "normatively loaded", in a way that can make them think they have answered deep questions about normativity when they in fact have not. (Korsgaard 1996, 42).

cut here (which might underwrite the idea that morality is distinctively normatively important, for example) a “fugitive” thought, that can’t be vindicated.²⁴ In what follows, we consider some recent ways of developing this kind of suspicion.

Before proceeding, we want to emphasize the contrast between thinking that the distinction is *nonsense*, and thinking that *there is no such thing as authoritative normativity*. If the distinction is nonsense, then “thinking” “moral norms are authoritatively normative” is much like “thinking” “‘twas brillig and the slithy toves did gyre and gimble in the wabe”.²⁵ One can certainly recite those words in inner speech, but arguably, ‘toves’ and ‘wabe’ fail to pick out even any specific imaginary thing. By contrast, one might think that ‘authoritative normativity’ is meaningful, but that its intension is empty, much like the term ‘largest prime’.

We here focus on the first, *nonsense* hypothesis. This hypothesis might well be coupled with defense of a form of *nihilism* about authoritative normativity, which claims that authoritative normativity doesn’t exist. For example, one might think that the topic of authoritative normativity is tightly tethered to nonsensical thought, that such thought fails to successfully refer to anything, and thus that we should hold that there is no such thing as authoritative normativity. However, the nihilist hypothesis itself is not our main focus here.

One way to argue for the nonsense hypothesis involves an idea that, following Evan Tiffany, we can call *deflationary normative pluralism*.²⁶ This is the claim that there are many sorts of norms (hence the “pluralism”), but no way of making sense of the idea of a “distinctive normative asymmetry” among them (hence the “deflationism”).

Tiffany motivates deflationary normative pluralism by appeal to two central claims. First, there is what he calls *contributory pluralism*. This is the claim (similar to an idea that we discussed in the previous section) that there are different sorts of contributory reasons that might bear on decision. Second, there is what he calls *deliberative pluralism*: there are a variety of standards that we might use to deliberate about how to act in light of a given set of considerations. The “deflationary pluralist” then simply denies that there is a standard which *normatively* stands out among these options in a distinctively “authoritative” way.²⁷

Derek Baker argues that there is a deeper problem, which we can think of as bolstering Foot’s initial claim that no good sense can be made of the “fugitive thought” that there is a kind of normative asymmetry between (e.g.) morality and etiquette. Baker argues that to say that there is an “ought simpliciter” or an “authoritative ought” doesn’t help because, put roughly, this is *just another ought*. What is needed is some way of cashing out the claim that this “ought” is authoritative. But, Baker argues, attempts to do so tend to bottom out in metaphors

²⁴ (Foot 1972, 311).

²⁵ (Carroll 1871/2019).

²⁶ (Tiffany 2007).

²⁷ (Tiffany 2007).

(recall the “oomph” and “trumping” examples from the preceding section). And he argues that the only obvious ways of discharging the metaphors just return us to some ought claim (e.g. “this is the ‘ought’ we ought to use”). This circularity is vicious, given that the proponent of an authoritative ought is trying to articulate a distinction among ought claims. Note that Baker is not just complaining about an alleged primitivism about authoritativeness here. Rather, he is claiming that we don’t even understand what it is that is supposed to be primitive.²⁸

Nathan Howard and N.G. Laskowski offer a different argument for rejecting the distinction between authoritative and merely generic normativity. They agree that in Sticky Situation (or cases like it), there is a question beyond what we morally ought to do, and what we prudentially ought to do. This is the question of what we ought to do *all things considered*. The ‘things’ we consider are reasons, and Howard and Laskowski suggest the following thought: some of our reasons are “moral” and some are “prudential”. The question being asked in Sticky Situation is just: what is the upshot of properly adding up *all of the considerations*? To the extent that we have “clubhouse reasons” or “mafia reasons”, they are just systematically less weighty than moral reasons.

In effect, Howard and Laskowski deny what Tiffany calls “deliberative pluralism”: they suggest that there is just one way to correctly add up considerations. This denial is challenged by an apparent datum we mentioned in §2.2 above: that we can seemingly contemplate what is demanded by (e.g.) the treehouse norms *all things considered*. This suggests that the treehouse norms (or more pernicious ones) are completely capable of applying to *all* the considerations.

Howard and Laskowski respond to a version of this worry as follows: the relevant notion of a reason (and presumably also its weight) is to be understood in terms of its functional role. They write: “this is partly because we’re sympathetic to the view that the concept of a reason is constituted in part by what it does or how it functions and that how it functions determines in part what it is about or to what it refers.”²⁹ And this functional role is supposed to be inconsistent with deliberative pluralism.

We think it is helpful to distinguish two thoughts in the quoted passage. The first is that there is a distinctively “normative” functional role (which we would call an “authoritative” one) played by the relevant concept of a reason. The second is that all “reasons” thought and talk that involves something like the “counting in favor of relation” involves this functional role (this idea is suggested by talk of “*the* concept of a reason”).

We find the second idea quite implausible. Even if one put asides uses of ‘reasons’ that are about such things as motivating reasons (which, put roughly, explain rather than justify action), the term ‘reasons’ (as we have emphasized) looks like it can (in certain contexts) be used to smoothly communicate about considerations that

²⁸ (Baker 2018). For connected discussion, see also (Baker 2017).

²⁹ (Howard and Laskowski 2024).

matter for settling what you *prudentially* ought to do, what you *legally* ought to do, etc., and not just what you *authoritatively* ought to do. Now focus instead on the first idea. Shorn of the second idea, optimism about the ability to spell out a distinctively normative functional role for (some uses of) ‘reason’ suggests a different conclusion from the one that Howard and Laskowski draw. For such a functional role, if illuminating, would potentially provide resources for spelling out the concept **AUTHORITATIVE REASON**, rather than showing that the concept is one that doesn’t make sense, and which we shouldn’t use for philosophical theorizing. The question, of course, is whether that can be spelled out.

3.2 Vindicatory Theories

In this section, we consider four different “vindicatory” theories. These are accounts that take the idea of “authoritative” normativity (or distinctively authoritative normative thought and talk) to make sense. In all but the first case, they seek to illuminate the contrast between generic and authoritative normativity. The attempts that we discuss are not exhaustive, but do helpfully illustrate a range of different “vindicatory” strategies that matter in this context.

The simplest approach is a kind of primitivism. The basic idea here is this: perhaps there is no non-circular way of characterizing authoritative normativity, but we nonetheless understand the idea that there is a putative normative contrast between (for example) morality and Lyra’s treehouse rules. We can think of Baker’s argument in the previous section as challenging this position. Baker’s claim is that we in fact don’t understand this contrast. Rather, we feel pressure to explicate the alleged distinction, but when we attempt to do so, all we can provide is metaphors that have no clear explanatory power.

One thing that makes Baker’s challenge seemingly powerful is that (at least *prima facie*) there seems to be no similar pressure toward primitivism about merely generic norms. What is it for Lyra’s treehouse rules to be generically normative? Put roughly, we can say that they have a certain *form*. They state the sort of thing to which we can aptly say: an act (or an attitude, etc.) *conforms* to this, or *fails to conform*. In turn, for some generic norms that are being used by people in a given context, we can understand their violation as an intelligible (even if misguided) basis for criticism, etc. by those who are invested in those norms. All of this may seem relatively clear and uncontroversial – at least relative to the issue of unpacking the idea of “authoritative” normativity.³⁰ What seems potentially hard to even state is what we would need to add to this status to get authoritative normativity. To posit as an addition something primitive, which we can only grope at with unhelpful metaphors, seems extremely theoretically unappealing.

³⁰ It’s important to stress that this relative claim does *not* entail that there aren’t important philosophical issues here about generic normativity (and our thought and talk about it). We think there very much are, as, for example, is brought out by much of literature on rules and rule-following following the wake of (Kripke 1982)’s discussion of (Wittgenstein 1953/1991).

Daniel Wodak can be understood as seeking to undercut the presupposition that makes this challenge potent: that we should understand authoritative normativity as generic normativity *plus something else*.³¹ Wodak proposes instead what he calls the *fake ducks model*. To be a real duck is not to be a fake duck plus something else. Rather, a fake duck is no duck at all. It is simply something made to look like a duck. Similarly, according to Wodak, a “merely generic reason” (for example) is something made to look like a reason. More precisely, Wodak proposes that to think that (e.g.) Lyra’s treehouse norms are “merely generically normative” is to take a fictionalist attitude towards them. That is, one has roughly the attitude towards the “merely generic” normative claims (e.g., those about “treehouse club reasons” etc.) that one has towards claims about the contents of fictions (such as “Sherlock Holmes is a great detective who lives in London”).³²

If Wodak’s argument is successful, it would achieve two things. First, it would provide a clear candidate explanation of the contrast between thought about a norm as merely generic, and thought about it as authoritative. On his theory: the former sort of thought involves a kind of fictionalist attitude, while the latter does not. Second, if his account is correct, our grasp of the idea of a merely generic norm is parasitic on our grasp of the idea of an authoritative norm. So the sort of worry about primitivism canvassed above – that we seem to have a clear grasp on the idea of generic normativity, and thus the primitivism about authoritativeness seems undermotivated – is undercut. This is because, on Wodak’s account, the idea of authoritativeness is crucial to understanding the idea of the merely generic. Thus, there is no clear grasp on normativity, on this account, that does not involve a grasp on authoritative normativity.

Consider two worries about Wodak’s strategy. First, we might worry about whether Wodak’s strategy fits with our best empirical accounts of relevant parts of thought and talk in cognitive science, psychology, and linguistics. For example: suppose Kratzer’s work on the semantics of ‘ought’ is correct. We then have a uniform semantics for ‘ought’ claims that isn’t fictionalist. It seems that we then need the (purported) fictionalism that Wodak posits to show up in our “ought” thoughts, but not in the semantics of ‘ought’. One might worry that this is an awkward fit.

To bring this worry out, imagine that Sarasdat is playing chess, but the tournament hall where she is playing has just caught on fire. Plausibly, Sarasdat authoritatively ought to leave calmly while she still can. Watching the game on the internet, however, you comment “Sarasdat ought to sacrifice her bishop; if she does, she can force checkmate.” Suppose we grant both that the word ‘ought’ is a context-sensitive expression (e.g., along the lines of Kratzer’s theory which we glossed above) and that your conversational context is one where the norms of chess are salient in the

³¹ (Wodak 2018a).

³² Wodak’s strategy can be seen as one way of developing David Enoch’s claim that “legal reasons” aren’t *real* reasons. As Enoch puts it: “the word ‘legal’ in ‘legal reason’ functions like the word ‘imaginary’ in ‘imaginary friend’” (Enoch 2011, 17, fn. 28). It can also be seen as one way of developing Judith Jarvis’s Thomson’s idea that there is, at the most basic explanatory level, just one real kind of “ought” that we think and talk about. See (Thomson 2001).

relevant way. This suggests that (if the pieces on the board are appropriately arranged) what you say counts as literally true. Why, in order to understand your thoughts, would we need to also interpret you as pretending that the norms of chess are authoritative? The basic challenge here isn't about whether Wodak's account can be made compatible with the details of Kratzer's semantics, but whether it fits with the underlying motivations for that semantics, which includes sensitivity to the context-sensitive way in which we use 'ought'.³³

A second worry about Wodak's account is how it generalizes. A natural way of developing Wodak's story suggests that everything other than thought about authoritative normativity involves a kind of fictionalism. If we take morality and prudence to be conceptually separate from authoritative normativity (as the setup in *Sticky Situation* suggests), Wodak's account implies that we should be fictionalists about the normativity involved in morality and/or prudence. That might be correct. But it's a much more controversial thesis than the parallel kind of fictionalism about the normativity of Lyra's clubhouse rules.³⁴

Wodak's discussion doesn't tell us anything about what it is to think that a norm is authoritative. But many philosophers find the very idea of authoritative normativity more perplexing than the idea of merely generic normativity.³⁵ Even if Wodak's account potentially reduces the theoretical costs of primitivism, we might still hope to find some non-primitivist story about authoritatively normative thought.³⁶

One of us (McPherson) attempts to provide the beginning of such an account in "Authoritatively Normative Concepts".³⁷ They focus on what has seemed to many the most promising starting point for authoritative normative concepts generally: the concept AUTHORITATIVELY NORMATIVE PRACTICAL OUGHT, in the context of first-person deliberation. It is natural to initially suppose that the thought "I

³³ One interesting option for Wodak is to claim that his account can be neutral on the best semantic theory, and instead be a part of the relevant metasemantic story about how 'ought' claims get their meaning, or one about how to best interpret the semantics of 'ought' claims. Such ways of reading (or at least developing) his account might be promising, but it's not clear that this move alone can deal with the worry we are raising here. Note also that this worry becomes arguably deeper if one also brings on board Kratzer's more recent work with Jonathan Phillips, which involves integrating her semantic account of modal claims with an account of modal thought (one which, importantly, doesn't involve the sort of fictionalism posited in Wodak's account). See (Phillips and Kratzer Manuscript).

³⁴ One important moving part in Wodak's view (which might well matter for the challenges we are raising for it here) is whether the "fictionalist" aspect of the relevant normative judgments is best captured by (purported) facts about the content of the judgments (e.g., they involve an "in the fiction" operator) or, instead, by (purported) facts about the kind of attitude that the judgment involves (e.g., they aren't full-blown beliefs strictly speaking, but rather a weaker kind of attitude of the kind involved in playing "make believe" games). For further discussion of this basic distinction between kinds of fictionalist theories, and for a general overview of moral fictionalism, see (Joyce 2017).

³⁵ We raised this idea earlier in this paper when discussing (Baker 2018). See also (Berman 2019). For dissent, see (Herskovitz 2015).

³⁶ Our point here connects to one of the concerns about recent forms of moral fictionalism developed in (Hussain 2004).

³⁷ (McPherson 2018).

authoritatively ought to do A" *settles* what to do.³⁸ Moreover, we might think that its authoritativeness in thought *consists in* this settling role. However, it is not easy to spell out the relevant notion of "settling" in an informative way. On the one hand, if we take it to mean "settling what we authoritatively ought to do", this account is unhelpfully circular. But other glosses on "settling" threaten to change the subject. For example, suppose we understand settling as simply *forming an intention*.³⁹ We then risk losing the normativity of the notion completely. After all, many people report thinking that they ought to be vegan, but continue to form (akratic) intentions to eat meat.

McPherson's strategy is to explicate the idea of "settling" by reference to a distinctive functional role: that of providing *first-personal deliberative guidance*. The relevant functional role is especially obvious in contexts of deliberation where (perceived) inter-normative conflict is salient, such as Sticky Situation, mentioned above. When, in Sticky Situation, the agent asks "what ought I to do?", she 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 our firmest grasp on the notion of an authoritatively normative concept. Of course, this gloss leaves open an obvious question: what does *distinctive suitability* amount to here? 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.

Here are some worries about McPherson's account of authoritatively normative concepts. First, it's not clear that there are any non-trivial success norms for the activity of settling deliberative conflicts. Second, perhaps the constitutive success norms for resolving deliberative conflict aren't appropriately authoritative. It is reasonably clear (for example) what the success norms for playing chess are, and why (and when) we might care about that activity. But given how opaque the success norms for "settling" deliberative conflict are, one might wonder whether we might simply be reasonably indifferent to whether we deliberate successfully.

McPherson's account is highly schematic. This has important benefits. For example, at least *prima facie*, it seems that different philosophers have different views (which are seemingly all fully intelligible and coherent) about which norms are authoritatively normative, and why. Other things being equal, it would be good to be able to explain to someone what common topic these philosophers were all talking about, in a way that didn't impute serious conceptual error to any of them. But the fact that this account is highly schematic is also arguably what gets the first two worries off the ground. So, in response, one might propose a less schematic theory of the idea of authoritative normativity, which takes on more substantive, if also more controversial, claims.

³⁸ Compare this to Allan Gibbard's locution of "the thing to do" in (Gibbard 2003). Cf. also (McPherson 2018, 254).

³⁹ See (Gibbard 2003).

To illustrate how this might go, consider the theory that Linda Radzik calls “Reflective Endorsement Coherentism”.⁴⁰ Radzik’s broadly “subjectivist” view is similar to a range of other recent views in ethical theory (including, notably, Sharon Street’s “Humean Constructivism”).⁴¹ We focus on Radzik’s view, which she explicitly puts forward as a view about “normative authority”, for the sake of illustration.

On Radzik’s view, we start with the norms that are *reflectively endorsed* by the agent. In our example above of the akratic agent thinking about veganism, we might say that the norm *I ought to be vegan* is reflectively endorsed by the agent (even if it is not motivationally efficacious). The agent is then imagined to consider whether they ought to continue to accept that norm. In asking that question, they will deploy their other reflectively endorsed norms (in a way that accords with the norms they endorse for reflecting on their own norms). If the vegan norm would survive this process, then it would count as an authoritative norm for the agent. If it would be rejected at the end of this process, it would not.

In addition to (arguably) avoiding some of the worries about McPherson’s schematic account, one appealing thing about Radzik’s view is that it explains the force of the criticism: *you wouldn’t think that if you really thought it through*. For on Radzik’s view, if this criticism of one of your normative commitments were correct, that normative commitment would ipso facto fail to be authoritative.

Consider four worries about Radzik’s view.

First, a common worry about many broadly “subjectivist” views is that they yield the intuitively wrong results in many cases. To press this point, consider that, on Radzik’s view (as on similar views, such as Street’s), it seems possible that an ideally coherent Caligula authoritatively ought to continue to torture and kill others, because that is what he (coherently) reflectively endorses doing. Many people think that’s false.

Second, it is unclear whether her view adequately accommodates our own doubts about our normative commitments. Consider that it seems possible to think thoughts such as the following: “even if I am ideally coherent by my own lights, I might be getting things *wrong* about what to do”. In other words, an agent might take seriously the possibility that she is in the grip of *fundamental normative error*.⁴² It’s not clear that Radzik’s view can adequately capture such thoughts.

Third, there is arguably a tension between the coherentism and the radical subjectivism within Radzik’s view. To see this tension, imagine that Richard thinks that updating his normative views on the basis of coherentist reflection is a highly

⁴⁰ (Radzik 2002).

⁴¹ (Street 2012).

⁴² Compare (Egan 2007).

unreliable means of doing so: in other words, he reflectively endorses not updating his views based on coherentist reflection.⁴³ But suppose that his views are also riddled with inconsistencies. Radzik's view seems to suggest that his *highly incoherent non-updated* normative views will be authoritative for him. But it might be hard to swallow the thought that incoherent normative views are authoritative for any agent.

Finally, one might worry how informative Radzik's view is about what the basic topic of "authoritative" normativity is. This is because her view seems to take a substantive stand on controversial issues within ethics (as evidenced by, for example, the first worry above). If one uses her view to define the basic topic of "authoritative normativity", then it seems that one is doing so using views that many will (at least at first blush) take to be normatively mistaken – and often deeply so – including for the reasons given above. In light of this, even if Radzik's view (or something similar to it, such as Street's) turns out to be correct as a substantive account of the authoritatively normative facts, it seems harder to accept that it is a correct account of what the very idea of "authoritative normativity" amounts to. For doing so would (at least *prima facie*) seem to impute a deep kind of conceptual error to many philosophers working within ethical theory, in a way that we might well think we have general methodological reasons to avoid.

Stepping back, Radzik's and McPherson's views each seek (in different ways) to provide accounts of authoritativeness that tie it to *intelligibility*. More precisely, the idea is that, if one accepted one of their accounts, there would be a kind of intelligible rationale for acting on one's judgments about what is authoritative. For McPherson, this is because those conclusions are understood as providing guidance that is distinctively deliberatively relevant. For Radzik, this is because authoritative conclusions are understood as ones that (an agent thinks) she would accept if she really thought things through. By contrast, if primitivism is true, there is no distinctive intelligible rationale for why one should act on one's authoritative conclusions, beyond *because they are authoritative*. A general question here is whether there is actually a deep tie between authoritativeness and intelligibility. If there is not, primitivism may be more credible than many people find it. Even if there is, McPherson and Radzik have surely not exhausted the plausible ways of understanding that connection, so finding alternative ways should be rich terrain for those looking to develop competitors to their views.

4. Questions

In this section, we explore five important questions about authoritative normativity, and the authoritative/merely generic contrast.

4.1 Priority between thought and reality

⁴³ The first idea connects to a point made by (Holton 1996), in discussing (Smith 1994).

Suppose that there are authoritatively normative *thoughts*. For example, suppose that Hakim can think “I authoritatively ought to give to charity”. And suppose further that there are authoritatively normative facts. For example, suppose that it is a fact that Hakim authoritatively ought to give to charity. An important question is: what is the relationship between the authoritativeness of thoughts and that of facts? For example: are certain thoughts authoritative in virtue of being about certain (authoritative) facts? Or are certain facts authoritative in virtue of being picked out by certain (authoritative) thoughts? Or is there some other, more complicated relationship here beyond these options?

Before we dive further into this question, it is worth emphasizing that the two suppositions just mentioned are both disputable. As we saw in §2.1, some have suggested that talk of “authoritative normativity” is literally nonsense. One could also dispute the second supposition, by thinking that there is authoritatively normative thought but no authoritatively normative facts. This could involve endorsing a version of one of the following two familiar metanormative theories. First, error theory might be true of authoritatively normative thought. Such thought might purport to be about authoritatively normative facts, but there might simply be no such facts.⁴⁴ By analogy, many physicists once thought about “luminiferous ether” (which was postulated as the medium through which light waves moved), but there turned out to be no such thing. Second, suppose that a (non-quasi-realist) form of non-cognitivism turned out to be true of authoritatively normative thought.⁴⁵ Put roughly, on such a view, our authoritatively normative thought might not play a representational role at all (not even in the minimalist sense of the kind of endorsed, for example, by “quasi-realist” expressivists). Consider a familiar analogy: if Robert thinks “Hooray for the Bluejackets”, this thought does not purport to represent a fact. Similarly, Hakim’s thought that he authoritatively ought to give to charity might consist in a certain non-representational attitude (e.g., a desire) towards giving to charity, such that there aren’t relevant “facts” here for these attitudes to be about.

Suppose we set these possibilities aside (as well as the parallel possibilities concerning generic normativity, and generically normative thought and talk). What then should we make of our guiding question here: which of thought or reality is the primary locus of authoritativeness, and which is “derivatively” authoritative?

This question presupposes that there couldn’t simply be independent accounts of the authoritativeness of normative thought, and of “normative reality”. We find this presupposition attractive. For it is very plausible that we use authoritative thoughts to think about authoritative reality. Indeed, on our preferred gloss of “normative reality” (which we have argued for elsewhere), it is defined as the part of reality, if

⁴⁴ For an overview of metaethical error theory, important versions of which involve a broader metanormative error theory of this kind, see (Olson 2017).

⁴⁵ For an overview of expressivism, see (Camp 2017), and for an overview of “quasi-realism”, see (Cuneo 2017).

any, that normative thought and talk is distinctively about.⁴⁶ If the authoritativeness in the two cases was completely independent, there would be a puzzle about why they happen to align in the way they do.⁴⁷

It's tempting to think that "realists" about authoritative normativity should be drawn to a reality-first view of the priority question. For example, Stephen Finlay says that for the normative realist, "a concept or word is normative in virtue of being about a normative part or feature of the world."⁴⁸ Consider here also Jamie Dreier's argument that a realist view is one according to which part of the account of *what it is* to have a normative thought involves reference to normative properties.⁴⁹

But the reality-first idea faces pressure. For example, Matti Eklund asks us to imagine a linguistic community that introduces a term – 'thgir' – that picks out what our word 'right' picks out. But this community does not use 'thgir' in any way in their deliberation or evaluation.⁵⁰ In this case, 'thgir' does not even appear to be a *generically* normative term, despite picking out the property of rightness.

Suppose that we accept that 'thgir' is not a normative term. At minimum, what this shows is that being about or referring to something normative is not sufficient for a term to be normative. This conclusion leaves open the possibility that for a term to be normative is for it to bear some more complex relationship to normative reality.⁵¹ However, it might seem like the simplest hypothesis is simply that what it is for a thought to be authoritatively normative can be spelled out independently of anything about normative reality. For example, perhaps it can be spelled out in terms of the sort of "deliberative settling role" that McPherson discusses.

4.2 Is 'authoritative' a technical term?

Reading this article, one might easily conclude that 'authoritative normativity' is a technical term. After all, it is a novel bit of terminology that we introduced by gesturing at it in a variety of ways. On the other hand, think about some of the reasons why we introduced this piece of terminology. On our view, it helps us talk about a kind of normativity that many philosophers and many ordinary people alike (at least implicitly) appeal to or presuppose in their thought and talk. Thus, one might think that the core idea of authoritativeness is a familiar folk notion. That is, unlike with "rest mass" or "bipolar disorder", one might take the idea to be familiar in pretheoretical thought, even if we previously lacked a word to express it.

This familiarity matters for thinking about what we are trying to do in theorizing about "authoritativeness". Consider four contrasting things we might seek to do:

⁴⁶ Our use of the term 'normative reality' in this way draws from (McPherson and Plunkett 2017) and (Plunkett and Shapiro 2017).

⁴⁷ For a different take on these issues, see (Leary 2020).

⁴⁸ (Finlay 2010, 334).

⁴⁹ (Dreier 2004).

⁵⁰ (Eklund 2017, 75).

⁵¹ Compare the discussion in (McPherson 2020).

- a. Make explicit the meaning of a piece of folk terminology.
- b. Illuminate the nature of a natural kind.
- c. Refine the meaning of an existing term.
- d. Introduce a novel piece of technical terminology, given its (purported) explanatory or epistemic benefits.

Plausibly, different methodologies are more suitable for some of these projects than they are for others. For example, widely-shared linguistic intuitions are plausibly directly relevant to (a), and somewhat relevant to (c), while they are of unclear significance for (b) and (d). Which of these projects we are focused on may also alter the plausibility of some of the options we have considered. For example, it might be hard to argue for primitivism, given the goal involved in (d), but perhaps there is more promise if we think ‘authoritative’ is a natural kind term (which matters for (b)). Similarly, how we evaluate the alleged link between authoritativeness and intelligibility mentioned in §2 looks different on projects (c) and (d) than it does on projects (a) or (b). (In the case of project (b), it might simply be a totally open question whether there is any tie to intelligibility.) These projects also interact with the priority between thought and reality question introduced above. Project (b), for example, might be a more comfortable fit if one is a “reality-first” rather than a “thought-first”.

4.3 Is authoritative normativity unified, or fragmented?

In the introduction to this paper, we distinguished two intuitive motivations for taking seriously the authoritative/generic distinction. The first was an intuitive contrast between the “normative importance” of different norms, such as between moral norms and the rules of a kid’s treehouse club. The second was the idea that we can ask questions using a normative concept that purports to *settle* conflicts between any and all norms. A *unified* theory of authoritative normativity purports to satisfy both motivations: it suggests that there is a single authoritative standard, which authoritatively settles such conflicts, and can potentially explain such contrasts. But it is also possible to embrace the first, but not the second motivation: to think that some norms are more authoritative than others, but that there is not a unified authoritative standard. Briefly consider two ways of motivating this idea.

First, David Copp suggests that the authoritative/generic distinction can be made by distinguishing standards which address genuine “problems of normative governance”, such as living well together with others, living in accordance with our own values, and regulating our beliefs in ways that enable us to promote our other aims.⁵² Because there are distinct (and sometimes conflicting) standards that best address each of these problems, authoritative normativity is not unified. As Copp puts it, there is no “supreme” normative standard or perspective.

⁵² (Copp 2009, 27-28).

Second, in some of our recent work, we've argued that disunity can be motivated by internal critique of McPherson's own account (discussed in §2 above).⁵³ As we have discussed, McPherson's account attempts to understand authoritatively normative concepts in terms of their relationship to a "deliberative settling role". We suggest that this is only *one* of several intuitively authoritative roles. For example, just as we can distinguish authoritative from non-authoritative "oughts", so too can we *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, given that an act 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 ultimate *evaluative perspective*. If there can be conflicts between, e.g. what is authoritatively best, and what one authoritatively ought to do, then this suggests that authoritative normativity is disunified.

The possibility of disunity of one of the kinds we've just glossed raises some important questions for work on authoritative normativity. First, can we give a unified account of what authoritativeness consists in that is consistent across different parts of normativity (e.g., the deontic vs. the evaluative vs. the aretaic)? Or are there different accounts that work for different parts of it? Second, if there are distinct facts about each of the different kinds of authoritative normativity, how do they interact? For example, if you authoritatively ought to do X, but it would be authoritatively bad if you did so, how should we think about the interaction between those different facts? Does it amount to a kind of discord that we should be worried about?⁵⁴

4.4 Relative authority

Within the class of things that are normative in at least the "generic" sense, we don't seem to just distinguish between those that are "merely" generically normative and those that are fully authoritative. We also (at least *prima facie*) seem to countenance claims about which norms are *more* authoritative than others. For example, even if you don't think that prudential or moral norms are fully authoritative, you might well think they are *more* authoritative than the rules of Lyra's treehouse club.⁵⁵

More strikingly, one might think that the core "authoritativeness phenomenon" we need to explain centers on the sort of *relative* authoritativeness claims, like the ones just made salient. If this idea is correct, it might suggest a worry for the theories of authoritativeness offered by McPherson, Wodak, and Radzik. This is because these theories seem tailored to explain *full* authoritativeness, which, one might worry, is not the core thing in need of explanation.

⁵³ (McPherson and Plunkett 2024).

⁵⁴ (McPherson and Plunkett 2024).

⁵⁵ Another possibility is that these norms have some property other than authoritativeness that marks a practically significant asymmetry with treehouse rules. For an important variant of this idea, see (Woods 2018).

Further, it is not clear whether or how such accounts could be extended to explain relative authoritativeness. For example, consider again McPherson's account of the authoritative practical ought. This account, recall, is tied to the functional role of settling a certain kind of (at least perceived) conflict that can arise in first-person deliberation. The "settling" idea is tailor-made to explain full authoritativeness, and it is not immediately clear how an account built on it can be used to explicate the idea of a norm being *more or less* authoritative.

The idea of relative authoritativeness raises questions beyond those explored in this paper. Consider two further examples. Is there just one kind of "relative authoritativeness", or are there diverse ways that some norms can "matter more" than others? How does the idea of relative authoritativeness relate to the idea of perfect (or full) authoritativeness? In short, taking seriously the phenomenon of relative authoritativeness promises to significantly complicate, or even transform, our thinking about authoritative normativity.

Conclusion

In this paper, we've provided a critical overview of recent discussions about the distinction between authoritative and merely generic normativity, and our thought and talk about them. As we hope to have shown, there are many complex, interesting issues here, which are far from settled. In closing, we want to make brief case for why these issues are worth tackling head on. We here focus on two different reasons why this is so.

First, consider that, across many different parts of philosophy, philosophers wonder about connections between different kinds of normativity (or different normative notions). For example, take the debate over "legal positivism". Put roughly, this debate (or at least one aspect of it) concerns whether legal facts are partly grounded in moral facts. Or consider the debate over "moral rationalism". Again, put roughly, this debate (or at least one aspect of it) concerns whether there is an intimate connection between moral requirements and normative reasons. Plausibly, these debates are often motivated by implicit assumptions about authoritativeness: the positivism debate by the assumption that morality is authoritative, and the rationalism debate by the assumption that normative reasons are.⁵⁶

There are, of course, any number of philosophically interesting issues one might consider about the relations between different sets of norms (or kinds of normativity, etc.) even if they aren't fully authoritative, or lack a deep tie to such normativity. For example: there are interesting questions about how legal norms are related to the norms of social conventions, even if neither of these norms are authoritative. But if what's at issue is really a (potential) link between some kind of normative fact (or concept, property, word, etc.) and authoritativeness as such, then

⁵⁶ See (Plunkett 2019) for further discussion, focused on the case of legal positivism.

we should just ask about that directly.⁵⁷ Furthermore, suppose that authoritative normativity is what we care about in many of these contexts. If so, this motivates seeking to better understand authoritative normativity as such, and how it differs from merely generic normativity.

Second, consider how discussion about authoritative normativity interacts with a range of other important debates in metaethics.⁵⁸ Our basic thought here is this: the idea of authoritative normativity can be used in helpful ways in evaluating metaethical views or arguments. For example, suppose we think that ethical facts are authoritatively normative (or else bear some intimate connection to authoritative normativity, such as entailing authoritative normative facts). We can then assess whether a proposed account of ethical facts is compatible with (or, more ambitiously, explains) this. Similarly, suppose we think that ethical thought and talk involves authoritatively normative concepts. We then assess whether a proposed account of ethical thought and talk is compatible with (or, more ambitiously, explains) this. In doing this, we can try to identify what about the view allows it to succeed here (if it does), or what might need to change about it in order for it to do so. These kinds of arguments arguably already play an important role in existing metaethical debates (even if not in couched in terms of “authoritative” normativity).⁵⁹ Thinking about them more carefully in terms of “authoritative” normativity can help us more carefully evaluate them, as well as develop new arguments.

To illustrate the kind of arguments we have in mind here, consider non-naturalistic realism. Put roughly, according to this kind of view, ethical thought and talk is fundamentally representational, and concerns certain non-naturalistic facts.⁶⁰ Now suppose one thinks that the relevant ethical facts are authoritatively normative. What explains why they are? Is it (for example) meant to be the sheer fact that they are “non-naturalistic”? That seems doubtful. Or at least that is so on many ways of understanding what that idea amounts to. Consider: suppose that what it means for a set of facts to be “non-naturalistic” (in the relevant sense) is for them to be fundamentally different in kind from, and metaphysically discontinuous with, the sorts of “naturalistic” facts that are at the center of scientific inquiry.⁶¹ It might well be that certain mathematical facts fit that description, if certain views of mathematical facts (e.g., mathematical Platonism) are correct. But such

⁵⁷ For further discussion, see (Plunkett 2019). One peril of not doing that, and instead relying on purportedly “proxy” terms such as ‘rational’ or ‘moral’, is that it can contribute to unreliable inferential patterns, given the multiple meanings (or at least inferential patterns) associated with such terms. For discussion, see (McPherson and Plunkett 2020).

⁵⁸ For our own views about what metaethics involves (and what metanormative inquiry more generally involves), see (McPherson and Plunkett 2017). For connected discussion, see also (Plunkett and Shapiro 2017) and (McPherson and Plunkett 2021).

⁵⁹ For example, this kind of argumentative strategy is (in significantly different ways) arguably a core strand of (Korsgaard 1996), (Gibbard 1990), and (Parfit 2011).

⁶⁰ See (Enoch 2017) for an overview of this kind of view.

⁶¹ For our own more developed views about what non-naturalism in metaethics involves, which are ways of sharpening this kind of characterization here, see (McPherson and Plunkett 2022), drawing on (McPherson 2015).

mathematical facts aren't authoritatively normative ones (let alone even normative ones of any kind): e.g., it's not as if the fact that the number two exists as an abstract object (on the Platonic conception of what that involves) is an authoritatively normative fact. So it doesn't seem that the "non-naturalistic" status of facts by itself guarantees anything about their being authoritatively normative. Something else needs to do the explanatory work here. Whatever that "something else" is might well be something that naturalists could appeal to as well (e.g., perhaps a kind of primitivism about authoritative normativity). If that's right, then it suggests that it's going to be a subtle issue of what role appeals to authoritative normativity can play in settling the debate over naturalism.

This example underscores an advantage of thinking about (or developing) metaethical arguments with an eye to issues about authoritative normativity: it can help us identify concerns about capturing or explaining authoritative normativity (or about authoritatively normative thought and talk) from other concerns that might otherwise be run together. Once concerns are framed in this way, it then reinforces the importance of the kind of work we have been discussing in this paper: namely, work aimed at explaining what "authoritative" normativity (and/or our thought and talk about it) amounts to, and whether it is something that is in good stead that should be used in our theorizing. On this front, consider that it might well turn out that some arguments in metaethics that accuse another view of failing to capture (or explain) "normativity" are appealing not only to the idea of authoritative normativity, but to an impoverished or mistaken account of what such normativity is.

Our work in this paper only scratches the surface of the issues that such work has involved, both in the contemporary literature as well as throughout the history of philosophy (even if not discussed using the terminology we are using here). Moreover, we are confident that there are many new issues that such work will take on in the future, as well as new positions and arguments that will be developed, which might push the discussions here in seriously novel directions. We hope that this paper can serve as a useful departure point for those whose work might well end up pushing our thinking forward in such ways.

Acknowledgments

Thanks to Derek Baker, John Bengson, Mitch Berman, Ruth Chang, David Copp, Lorraine Daston, Benjamin Kiesewetter, Sean Kim, Niko Kolodny, Zachary Lang, Adrian Russian, Thomas Schmidt, Scott Shapiro, Zoe Thierfelder, Ezekial Vergara, Daniel Wodak, and Fangzhou Yu for helpful feedback and discussion. An earlier draft of this paper was presented at the Wissenschaftskolleg zu Berlin (institute for advanced study). Thanks to everyone who participated in that discussion.

Works Cited

- Baker, Derek. 2017. The Varieties of Normativity. In *The Routledge Handbook of Metaethics*, edited by T. McPherson and D. Plunkett. New York City: Routledge. 567-581.
- . 2018. Skepticism about Ought Simpliciter. In *Oxford Studies in Metaethics*, Vol. 13, edited by R. Shafer-Landau. Oxford: Oxford University Press. 230-252.
- Bedke, Matthew. 2019. Practical Oomph: A Case for Subjectivism. *Philosophical Quarterly* 69 (277):657-677.
- Berman, Mitchell. 2019. Of Law and Other Artificial Normative Systems. In *Dimensions of Normativity: New Essays on Metaethics and Jurisprudence*, edited by D. Plunkett, S. Shapiro and K. Toh. Oxford: Oxford University Press.
- Camp, Elisabeth. 2017. Metaethical Expressivism. In *The Routledge Handbook of Metaethics*, edited by T. McPherson and D. Plunkett. New York City: Routledge. 87-101.
- Carroll, Lewis. 1871/2019. Jabbewocky. In *Through the Looking Glass*. Orinda, CA: SeaWolf Press.
- Chrisman, Matthew. 2016. *The Meaning of 'Ought': Beyond Descriptivism and Expressivism in Metaethics*. Oxford: Oxford University Press.
- Copp, David. 2004. Moral Naturalism and Three Grades of Normativity. In *Normativity and Naturalism*, edited by P. Schaber. Frankfurt: Ontos-Verlag. 7-45.
- . 2009. Toward a Pluralist and Teleological Theory of Normativity. *Philosophical Issues* 19 (1):21-37.
- Cullity, Garrett, and Berys Gaut, eds. 1997. *Ethics and Practical Reason*. Oxford: Oxford University Press.
- Cuneo, Terence. 2017. Quasi-realism. In *The Routledge Handbook of Metaethics*, edited by T. McPherson and D. Plunkett. New York: Routledge. 626-642.
- Dreier, James. 2004. Meta-Ethics and the Problem of Creeping Minimalism. *Philosophical Perspectives* 18:23-44.
- Egan, Andy. 2007. Quasi-realism and fundamental moral error. *Australasian Journal of Philosophy* 85 (2):205 – 219.
- Eklund, Matti. 2017. *Choosing Normative Concepts*. Oxford: Oxford University Press.
- Enoch, David. 2011. Reason-giving and the law. In *Oxford Studies in Philosophy of Law*, edited by L. Green and B. Leiter: Oxford University Press.
- . 2017. Non-Naturalistic Realism in Metaethics. In *The Routledge Handbook of Metaethics*, edited by T. McPherson and D. Plunkett. New York: Routledge.
- Finlay, Stephen. 2010. Recent work on normativity. *Analysis* 70 (2):331-346.
- Foot, Philippa. 1972. Morality as a System of Hypothetical Imperatives. *Philosophical Review* 81 (3):305-316.
- Gibbard, Allan. 1990. *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Cambridge, Mass.: Harvard University Press.
- . 2003. *Thinking How to Live*. Cambridge, Mass.: Harvard University Press.
- Herskovitz, Scott. 2015. The End of Jurisprudence. *The Yale Law Journal* 124 (4):1160-1204.

- Holton, Richard. 1996. Reason, value and the muggletonians. *Australasian Journal of Philosophy* 74 (3):484 – 487.
- Howard, Christopher, and R.A. Rowland. 2023. *Fittingness*. Oxford: Oxford University Press.
- Howard, Nathan, and N.G. Laskowski. 2024. Robust vs. Formal Normativity II: No Gods, No Masters, No Authoritative Normativity. In *Oxford Handbook of Metaethics*, edited by D. Copp and C. S. Rosati.
- Hussain, Nadeem J. Z. 2004. The Return of Moral Fictionalism. *Philosophical Perspectives* 18 (1):149–188.
- Jones, Karen, and Francois Schroeter. 2018. *The Many Moral Rationalisms*. Oxford: Oxford University Press.
- Joyce, Richard. 2006. *The Evolution of Morality*. Cambridge: MIT Press.
- . 2017. Fictionalism in Metaethics. In *The Routledge Handbook of Metaethics*, edited by T. McPherson and D. Plunkett. New York City: Routledge.72-86.
- Kant, Immanuel. 1785/2012. *Groundwork of the Metaphysics of Morals*. Cambridge: Cambridge University Press.
- Kolodny, Niko. 2005. Why Be Rational? *Mind: A Quarterly Review of Philosophy* 114 (455):509-563.
- Korsgaard, Christine M. 1996. *The Sources of Normativity*. New York: Cambridge University Press.
- Kratzer, Angelika. 2012. *Modals and Conditionals: New and Revised Perspectives*. Oxford: Oxford University Press.
- Kripke, Saul. 1982. *Wittgenstein on Rules and Private Language*. Cambridge, MA: Harvard University Press.
- Leary, Stephanie. 2020. Choosing normative properties: a reply to Eklund's Choosing Normative Concepts. *Inquiry: An Interdisciplinary Journal of Philosophy* 63 (5):455-474.
- Lord, Errol, and Kurt Sylvan. 2019. Reasons: Wrong, Right, Normative, Fundamental. *Journal of Ethics and Social Philosophy* 15 (1).
- Mackie, J. L. 1977. *Ethics: Inventing Right and Wrong*. New York: Penguin.
- McLeod, Owen. 2001. Just plain "ought". *The Journal of Ethics* 5 (4):269-291.
- McPherson, Tristram. 2011. Against Quietist Normative Realism. *Philosophical Studies* 154 (2):223-240.
- . 2015. What is at Stake in Debates among Normative Realists? *Noûs* 49 (1):123-146.
- . 2018. Authoritatively Normative Concepts. In *Oxford Studies in Metaethics Vol. 13*, edited by R. Shafer-Landau. Oxford: Oxford University Press.253-277.
- . 2020. Ardent Realism without Referential Normativity. *Inquiry*.
- 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.
- . 2020. Conceptual Ethics and The Methodology of Normative Inquiry. In *Conceptual Engineering and Conceptual Ethics*, edited by A. Burgess, H. Cappelen and D. Plunkett. Oxford: Oxford University Press.274-303.
- . 2021. Conceptual Ethics, Metaepistemology, and Normative Epistemology. *Inquiry*.

- . 2022. Ground, Essence, and the Metaphysics of Metanormative Non-Naturalism. *Ergo* 9 (26):674-701.
- . 2024. The Fragmentation of Authoritative Normativity. In *Oxford Studies in Metathics, Vol. 19*, edited by R. Shafer-Landau. Oxford: Oxford Studies in Meaethics.1-28.
- Olson, Jonas. 2017. Error Theory in Metaethics. In *The Routledge Handbook of Metaethics*, edited by T. McPherson and D. Plunkett. New York City: Routledge.58-71.
- Parfit, Derek. 2011. *On What Matters*. Oxford: Oxford University Press.
- Phillips, Jonathan, and Angelika Kratzer. Manuscript. Decomposing Modal Thought.
- Plunkett, David. 2019. Robust Normativity, Morality, and Legal Positivism. In *Dimensions of Normativity: New Essays on Metaethics and General Jurisprudence*, edited by D. Plunkett, S. Shapiro and K. Toh. New York City: Oxford University Press.
- Plunkett, David, and Scott Shapiro. 2017. Law, Morality, and Everything Else: General Jurisprudence as a Branch of Metanormative Inquiry. *Ethics* 128 (1):37-68.
- Radzik, Linda. 2002. A coherentist theory of normative authority. *The Journal of Ethics* 6 (1):21-42.
- Railton, Peter. 1999/2003. Normative Force and Normative Freedom: Hume and Kant, but not Hume Versus Kant. In *Facts, Values, and Norms : Essays Toward a Morality of Consequence*. New York: Cambridge University Press.322-352.
- Raz, Joseph. 1979/2002. *The Authority of Law: Essays on Law and Morality*. Oxford: Oxford University Press.
- Scanlon, T. M. 1998. *What We Owe To Each Other*. Cambridge: Harvard University Press.
- Schroeder, Mark. 2007. *Slaves of the Passions*. Oxford: Oxford University Press.
- Smith, Michael. 1994. *The Moral Problem*. Cambridge: Blackwell.
- Street, Sharon. 2012. Coming to Terms with Contingency : Humean Constructivism about Practical Reason. In *Constructivism in Practical Philosophy*, edited by J. Lenman and Y. Shemmer: Oxford University Press.
- Thomson, Judith Jarvis. 2001. *Goodness and Advice*. Princeton: Princeton University Press.
- Tiffany, Evan. 2007. Deflationary Normative Pluralism. *Canadian Journal of Philosophy* 37 (5):pp. 231-262.
- Wedgwood, Ralph. 2004. The metaethicists' mistake. *Philosophical Perspectives* 18 (1):405-426.
- Williams, Bernard Arthur Owen. 1965/1973. Ethical Consistency. In *Problems of the Self*. Cambridge Eng.: Cambridge University Press.166-186.
- Wittgenstein, Ludwig. 1953/1991. *Philosophical Investigations*: Wiley-Blackwell.
- Wodak, Daniel. 2018a. Fictional Normativity and Normative Authority. *Canadian Journal of Philosophy* 49 (6):828-850.
- . 2018b. What Does 'Legal Obligation' Mean? *Pacific Philosophical Quarterly* 99 (4):790-816.

Woods, Jack. 2018. The Authority of Formality. In *Oxford Studies in Metaethics Vol. 13*, edited by R. Shafer-Landau. Oxford: Oxford University Press. 207-229.