

Authoritatively Normative Concepts^{*}

Tristram McPherson
Ohio State
dr.tristram@gmail.com

MS. FOR SUBMISSION TO *OXFORD STUDIES IN METAETHICS*
November 27, 2016

Abstract:

This paper offers an analysis of the authoritatively normative concept PRACTICAL OUGHT that appeals to the constitutive norms for the activity of non-arbitrary selection. I argue that this analysis permits an attractive and substantive explanation of what the distinctive normative authority of this concept amounts to. I contrast my account with more familiar constitutivist theories, and briefly show how it answers ‘schmagency’-style objections to constitutivist explanations of normativity. Finally, I explain how the account offered here can be used to help realists, error theorist, and fictionalists address central challenges to their views.

^{*} I am grateful for illuminating discussion of this paper at the Melbourne Moral Rationalism Workshop, the Australasian Annual Workshop in Moral Philosophy, the Hong Kong Metaethics Workshop, University of Kentucky, the Virginia Tech philosophy faculty reading group, Ohio State, the CRNAP Varieties of Normativity Workshop at Princeton, the Varieties of Normativity Workshop at Uppsala, and the Chapel Hill Metaethics Workshop. Special thanks to Derek Baker, Garrett Cullity, David Faraci, Anandi Hattiangadi, Sarah McGrath, David Plunkett, and Michael Smith, whose comments and conversations dramatically improved my understanding of the issues in this paper.

Introduction

There are many species of norms. For example, an act can be an illegal chess move, and impolite to boot, but still morally required. And there are famous philosophical puzzles about the *generic* normativity shared by all of these species.¹ For example, Saul Kripke's Wittgenstein (1982) famously challenged our ability to explain how someone could count as following one norm rather than another. However, metaethicists have not typically focused on understanding normativity in this generic sense, because puzzles about generic normativity arguably do not cut to the heart of the metaethical dialectic.² One illustration of this is that many of the most influential and well-motivated classes of views in metaethics – non-naturalism, expressivism, error theory, and fictionalism – are not attractive views about merely generic norms like the norms of chess.³

Work in metaethics has instead focused on moral normativity, or – increasingly – on *practical normativity*. What is this latter notion? Quick glosses usually focus on two ideas. The first is a contrast with morality: sometimes one has a choice where morality is silent, but it seems nonetheless clear that one option is *better* than the other. For example, suppose that I have no morally significant options, but can choose whether to spend the afternoon in pleasant conversation or counting blades of grass. In this case I arguably practically ought to choose conversation. The second is the idea that practical normativity is distinctively *authoritative*. The practical norms are the ones that *settle what to do*, in a way that somehow contrasts with the deliverances of etiquette (for example).

These glosses orient us to the idea of authoritative normativity. This paper aims to help us to better understand this idea by providing an illuminating account of one authoritatively normative concept: PRACTICAL OUGHT. (I use small caps to denote concepts.) My account analyzes this concept in terms of the

¹ The 'generic' label is from Copp 2005a. I previously called this sort of normativity 'formal' in my 2011.

² In the text, I use 'metaethical' and related terms in a familiar, very loose sense. For a more careful taxonomy, see McPherson and Plunkett *forthcoming-b*.

³ For a discussion that usefully and sharply distinguishes the task of explaining practical normativity from that of explaining generic normativity, see FitzPatrick 2008. Gibbard 2012 takes the core issues at stake in metaethics to carry over at least to semantic norms. In this respect, the project of this paper sides with FitzPatrick against Gibbard. For an important challenge to Gibbard's attempt to extend his master argument for expressivism to semantic norms, see Baker *Ms*.

constitutive norms for the activity of non-arbitrarily selecting an option (§4). I argue that this analysis permits an attractive and substantive explanation of what the distinctive normative *authority* of this concept amounts to. I briefly show how my account can answer ‘schmagency’-style objections to constitutivist explanations of normativity (§5). Finally, I briefly sketch some of the central metaethical implications of my account, showing that it can be used to help realists, error theorists, and fictionalists address central challenges to their views (§6). I begin by more carefully introducing the concept that I intend to analyze, the style of account I will offer, and the constitutivist resources that this account deploys (§§1-3).

1. Elusive practical normativity

In the Introduction, I briefly gestured at the notion of authoritative normativity. Because this notion is central to my project in this paper, it will be helpful to re-introduce it more carefully. To begin, consider the following deliberative scenario:

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?*⁴

Sticky Situation concludes with an interesting question, about which you might agonize. Because of this, it is implausible to read ‘ought’ in this question as expressing⁵ either the concept MORALLY OUGHT, or the concept PRUDENTIALY OUGHT. For you already take yourself to know the answer to the question, understood in those ways. (Note that the answer you take yourself to know might be *incorrect*: perhaps morality and prudence *cannot* conflict. But this does not undercut the intelligibility of someone thinking that they can conflict, and posing this question in light of the perceived conflict.)

⁴ My initial characterization and discussion of Sticky Situation is indebted to Wedgwood 2004, 406, who in turn credits Cullity and Gaut 1997.

⁵ I use *expression* to pick out the relation between linguistic tokens and the mental states they are associated with in virtue of the meanings of those linguistic tokens. In doing so, I commit myself to the (plausible but controversial) idea that the meanings of (some) linguistic tokens entail such a connection to contentful speaker mental states. I do *not* intend the more controversial thesis that the meanings of linguistic entities are, or are grounded in, their conventional relation to the contentful mental states they express.

It is natural to step back from the specific question highlighted in Sticky Situation, to query the normative significance of conflict between morality and prudence more generally. Here, one familiar view is:

Moral Rationalism If moral requirements and prudence conflict, one *ought* to do the morally required thing.⁶

Whether or not it is true, Moral Rationalism appears to be an interesting and substantive thesis. The upshot is the same as in Sticky Situation: if we interpret ‘ought’ here as expressing MORALLY OUGHT, Moral Rationalism is a trivial claim; if we interpret it as expressing PRUDENTIALLY OUGHT, it is instead trivially false.

In both Moral Rationalism and Sticky Situation, it seems most plausible to read ‘ought’ as expressing a concept that purports to wear a distinctive normative authority on its sleeve, in a way that even moral and prudential ‘ought’s do not. For example, this explains why the truth of Moral Rationalism is seen by many as crucial to vindicating the normative significance of morality. It will be useful to have a label for this distinctively authoritatively normative concept: *PRACTICAL OUGHT*.

PRACTICAL OUGHT is part of a family of authoritative concepts, which range across various dimensions of normative structure. Just as we can talk of a range of narrowly *moral* concepts – MORAL REQUIREMENT, MORALLY BETTER, MORAL REASON, etc. – we can talk about their explicitly authoritative structural correlates: PRACTICAL REQUIREMENT, PRACTICALLY BETTER, PRACTICAL REASON, etc. The last of these is arguably the most infectious contemporary locution for gesturing at authoritative normativity.⁷ This paper focuses almost exclusively on PRACTICAL OUGHT, postponing the important question of how my account could be extended to other authoritatively normative concepts.

Authoritatively normative concepts are an especially natural locus for metaethical enquiry. Suppose on the one hand that Moral Rationalism (or

⁶ For discussion of related principles, see Smith 1994, Darwall 1997, van Roojen 2010, Lord and Plunkett *forthcoming*, and several of the papers in Jones and Schroeter *forthcoming*.

⁷ Contemporary philosophers very often deploy locutions that are plausibly intended to convey the distinctive authority characteristic of this family of concepts: compare Scanlon’s talk of reasons in the ‘standard normative sense,’ (1998, 17-19), Schroeder’s ‘normativity of the normative’ (2007, 79), and Hampton’s talk of ‘normative authority’ (1998, 85ff) which my talk of ‘authoritative’ norms echoes. In conversation – if less often in print – philosophers will sometimes speak of normative ‘oomph’.

something like it) is true. Then it would be very natural to expect that an account of the authoritatively normative concepts or properties will be crucial to understanding moral thought, talk, and reality. This will also hold true if Moral Rationalism is false, but a weaker distinctive connection between morality and practical normativity holds. Suppose finally that moral requirements turn out to be like the requirements of chess or etiquette, in being only contingently connected to what we practically ought to do. Then there would be little reason to think that the distinctive metaethical puzzles were raised by morality, and it would be natural for metaethicists to turn their attention to the authoritatively normative concepts. For these reasons it is not surprising that over the past generation, the focus of metaethical work has shifted significantly from morality to what I am calling authoritative normativity.⁸

The nature of the authoritatively normative concepts can seem elusive, however. This can be illustrated by considering four *unfruitful* strategies for illuminating them. First, notice that one cannot illuminate authoritatively normative concepts *generally* simply by analyzing one member of the family of such concepts in terms of another. For example, even if it were possible to do so, it would not suffice to analyze all other authoritatively normative concepts in terms of PRACTICAL REASON. This is because the core question here is about what is distinctive of the whole family of authoritatively normative concepts, or – if inter-normative conceptual analyses are possible – whichever of these concepts are conceptually basic.

Second, it would be a mistake to gloss authoritativeness as *categoricity*, where a norm is categorical if you cannot escape its application simply by changing your desires or intentions. Categoricity appears neither necessary nor sufficient for our understanding of authoritativeness. Its insufficiency was well-established by Philippa Foot (1997). The norms of etiquette are categorical in the sense just mentioned: indifference to these norms does not make their violation any less impolite. Its necessity is rendered doubtful by the Humean research program, according to which facts about authoritative normativity are grounded in

⁸ Representative examples include Bedke 2010, Gibbard 2003, Schroeder 2007, Street 2008, and Wedgwood 2007.

facts about agents' contingent desires. If categoricity were central to our understanding of authoritativeness, then familiar forms of Humeanism would be guilty of a transparent category mistake, which is hardly plausible.

Third, in order to mark a contrast with norms like etiquette, Derek Parfit contrasts normativity in the 'rule-implying' sense, with normativity in the 'reasons-implying' sense (2011, §88), where the latter is his way of adverting to what I am calling authoritative normativity. But this is potentially misleading: there are lots of ways of using 'reason', many of which fail to be transparently authoritative. One can talk about *moral* or *aesthetic* reasons, or *reasons of etiquette*, all of which are manifestly generically normative, but none of which is *obviously* authoritatively normative. In these cases, 'reason' is best understood as adverting to a certain kind of normative structural kind, which can have instances across both authoritative and merely generic normative systems.⁹

Fourth, Ralph Wedgwood (2004) glosses our target concept as the 'all-things-considered' ought, but this is potentially misleading in the same way. The locus classicus for 'all-things-considered' talk is Donald Davidson's discussion of weakness of will, which contrasted the judgments that X is better than Y *simpliciter*, with the judgment that it is better *prima facie*, and that it is better *all-things-considered* (2001 [1969]). The important point is that Davidson's contrast here is again *structural*: the *simpliciter/prima facie/all-things-considered* contrast will show up within moral judgments, prudential judgments, even within chess judgments. For example, one might say, (pedantically), "Weaknesses around white's king *prima facie* support mounting an attack there, but in light of the concentration of my pieces, and the open c-file, it is better *all-things-considered* to secure strategic advantages on the queenside instead." Here 'better *all-things-considered*' is most plausibly read as a claim internal to the norms of chess.

As these examples show, common attempts to provide an informative gloss on the notion of an authoritatively normative concept appear to fail. At this point,

⁹ I do not deny that there is a reading of 'aesthetic reason' (for example) as meaning something like *a practical reason simpliciter that has an aesthetic basis*. I only mean to insist that 'reason' is also used in careful philosophical contexts with the structural meaning (for one example, see the discussion of moral rationalism in Lord and Plunkett *forthcoming*).

an objector might claim that she has no idea what is being gestured at with talk of ‘distinctively authoritative’ normativity. She might point out that the term ‘authoritative’ is so far simply a label, and should in no way convince us that we have a grip on the alleged concept being deployed. She might continue: morality is distinctively *morally* authoritative, prudence is distinctively *prudentially* authoritative (etc.), and there is no other coherent notion of authority which can be used to give us purchase on PRACTICAL OUGHT. This might lead her to suggest that philosophers’ attempts to discuss authoritative normativity simply fail to latch onto a genuine concept. Call this view *deflationary pluralism* about normative concepts.¹⁰

The cost of deflationary pluralism is high, however. It appears to deny that Sticky Situation raises an interesting question and that Moral Rationalism is an interesting thesis. And it suggests that the range of central metaethical views mentioned in the Introduction – non-naturalism, expressivism, error theory, and hermeneutic fictionalism – are confused at a fairly fundamental level, if (as I suggested) they are often tacitly motivated in part by the thought that metaethics has authoritative normativity as its explanatory target, either directly or indirectly.

In light of this, one might claim instead that the elusiveness of authoritative normativity is explained by *concept primitivism*. On this view, perhaps some authoritatively normative concepts can be analyzed in terms of others, but the fundamental authoritative concept(s) are unanalyzable: nothing non-circular can be said to illuminate their distinctively nature (compare Scanlon 1998, 17 on reasons).

The strategy of positing a primitive concept here arguably pairs best with a non-naturalist metaphysics: on this view, what is distinctive of authoritatively normative thought is that it alone is *about* a sui generis, distinctively normative part of reality. But arguably, our only way of understanding the idea that this bit of reality is distinctively *normative* is that we talk about it using *these* concepts.¹¹ If so,

¹⁰ Compare Copp 2005a, 2005b, Tiffany 2007, and especially Baker *This Volume* for relevant discussion.

¹¹ For important discussion of the underexplored question of whether the normativity of our concepts is parasitic on the normativity of the reality they represent – or vice-versa – see Eklund *forthcoming*.

this metaphysical explanation seems unhelpful, and primitivism remains unsatisfying.

Against these views, I argue that it is possible to provide an informative account of the concept PRACTICAL OUGHT that vindicates its distinctive normative significance. The nature of the account will also suggest a plausible explanation of why the nature of authoritative normativity appears so elusive.

2. Conceptual analysis and the vindication of authority

My account will consist of an analysis of the concept PRACTICAL OUGHT paired with an explanation of how this analysis vindicates the distinctive normative authority that concept purports to have. Because the very possibility of providing informative analyses of philosophical concepts is controversial, this section sketches how I understand this project.

I have sought to locate the concept I am discussing – the PRACTICAL OUGHT – in part by adverting to a kind of thought that most of us recognize (for example, in Sticky Situation) and to familiar moves in ethical theorizing (for example debates about Moral Rationalism). There may be a word in non-philosophical English that uniquely picks out this concept, but for the purposes of this project, I set that question aside. My aim is to understand the concept – whatever it is – that plays these distinctive roles in deliberative and theoretical contexts. This means that the central desideratum on the analysis of this concept is that it permits an illuminating explanation of the concept's distinctive normative purport.

I take this desideratum to require in turn that the analysis be informative and non-circular. For otherwise the explanation it is paired with risks presupposing the very thing it seeks to explain. However, I do not take it to require analyzing the concept PRACTICAL OUGHT in wholly *non-normative* terms. I will make use of generic normativity in my analysis without precisely explicating the idea of a generically normative concept. This is because I aim to explain the *contrast* between authoritative and merely generic normativity, not to solve the puzzles that are shared by authoritative and merely generically normative concepts.

I take the project of this paper, as just glossed, to be compatible with a range of views about the methodology for investigating our concepts. For example, on the expansive conception to conceptual analysis championed by Frank Jackson (1998), my project could potentially be understood as an instance of conceptual analysis.¹² Alternatively my project could be understood as an instance of the sort of ‘reforming definition’ approach defended by Peter Railton (1989). Here the idea is, roughly, that the correct account of PRACTICAL OUGHT is the one that best articulates and refines the theoretically crucial elements of our conceptual practice. (For more discussion of relevant methodological complications, see McPherson and Plunkett, *forthcoming-a*.)

3. Constitutivist resources

The account that I will offer appeals to the contrast between the norms constitutive of certain activities, and merely generic norms. To see this contrast, consider the following norm, which I hereby introduce:

Touch Nose Move You must touch your nose while playing any chess move.

To be clear, this is neither an interpretation of the rules of chess, nor a proposed amendment to them. Nor is it an (absurd) claim about your moral, prudential, or authoritative obligations. It is introduced as an independent norm. It is generically normative: it is a standard that actions can clearly meet or violate. Indeed, it is clear what it takes to abide by it or to violate it: if you play chess moves without touching your nose, you are violating it.

What then is the contrast between merely generic norms and the constitutive norms for an activity? In my view, the constitutive norms for an activity fix whether one performs that activity correctly or successfully, not merely

¹² Given that my project may be well-understood as offering a sort of conceptual analysis, I should explain why I am unworried by the open question argument. In anything like its canonical form (à la Moore 1993 [1903]), I take the argument to simply be hopeless. Consider: suppose that I offer you a novel and surprising analysis A of some philosophical concept C. As a result of reading my argument, your credence in the analysis goes from negligible to 0.7. As philosophy goes: *extraordinary success!* But now you ask yourself: *I know that X is A; but is it C?* This is likely an open question given that your credence in the analysis is only 0.7. As far as I can tell, the openness of this question gives you no information you did not have before, and is no reason to lower your credence in the analysis.

relative to some norm or another, but relative to the very activity one is engaging in.¹³ Consider an example: if, while playing chess, you move a knight diagonally, you violate the rules for the movement of the pieces, and hence play incorrectly. If you play legal moves, but get checkmated, you have thereby played unsuccessfully. By contrast, because the Touch Nose Move norm is a merely generic norm, and not constitutive of the activity of playing chess, there is a natural sense in which one does not play *chess* incorrectly or unsuccessfully in virtue of violating it.

To bring out the significance of this contrast, suppose that I clearheadedly intend to play chess. Suppose next that while doing so, I routinely make moves without touching my nose, and you point out that by doing so, I violate the Touch Nose Move norm. It seems that without any confusion I might simply note that I don't care about that. Suppose by contrast that I move a knight diagonally, and you point out that this violates the rules of chess. I might reply in all sorts of intelligible ways: I might decide that I am not playing chess after all; I might evince confusion about the rules of chess; I might accept the correction to my play, etc. But it would be puzzling for me to say that I *am* playing chess, but I simply don't care about playing according to its rules. It would become tempting to impute some rational failing to me if I were to say this: perhaps a failure to understand what it is to play chess. The precise nature of the criticism that is warranted here is controversial. My aim is only to establish that, insofar as one is engaged in an activity, the constitutive norms for that activity appear to have a kind of grip on one that the merely generic norms lack. And this grip is naturally understood as marking an asymmetry in the *normative significance* of constitutive norms, compared to merely generic norms.

The apparent promise of this asymmetry has launched a thousand constitutivisms in ethics. But by itself this asymmetry is not enough to explain the distinctive normativity of the concept PRACTICAL OUGHT. After all, the asymmetry does not entail that if one engages in an activity like chess, one practically ought to follow its constitutive norms. To make this vivid, suppose that someone invented a game in which players compete to torture a puppy in the most awful way possible.

¹³ For a different analysis of what is distinctive of constitutive normativity, which focuses on constitutive *aims*, see Katsafanas 2013, 39.

Finding oneself playing this game surely does not entail that one ought to torture a puppy. At least intuitively, it does not even entail that one has a practical reason to torture a puppy.¹⁴ Thus, the constitutivist must identify the added ingredient that, when combined with the constitutive element, will yield authoritative normativity.

The most familiar proposal here is that the authoritative norms are constitutive norms for an activity that is in some sense *inescapable*.¹⁵ To see the appeal of this proposal, notice that one striking feature of chess or the puppy torturing game mentioned above is that one can stop playing them. When one does stop playing them, any sense that one is under normative pressure to abide by their constitutive norms evaporates. Identifying an inescapable activity seemingly promises to prevent such evaporation.

Despite this intuitive appeal, I am pessimistic about the inescapability approach, for reasons suggested by Matthew Silverstein (2015) and especially David Enoch (2006, 2011). On the one hand, it is unclear in what sense deliberation or agency (or whatever else the constitutivist points to as the relevant activity) is inescapable. On the other hand, it is unclear why the inescapability of an activity makes the norms of that activity *authoritatively* normative. In what follows, I propose a very different way of developing the constitutivist idea.

4. An account of the concept PRACTICAL OUGHT

This section develops an analysis of the concept PRACTICAL OUGHT that appeals to the constitutive success norms for a distinctive activity: *non-arbitrary selection*. After sketching a partial analysis, I explain how it permits an attractive informative account of what it is for a concept to be authoritatively normative. I then complete the analysis in light of challenges related to the scope of applicability of the concept, including the conditional fallacy. I begin by introducing the activity at the heart of my account.

¹⁴ Compare Enoch 2006, 185-6. Note, however, that Schroeder's case against the reliability of 'negative existential intuitions' about reasons (2007, 92-7) could be used to challenge this intuition.

¹⁵ As Korsgaard memorably puts the inescapability idea, "Human beings are *condemned* to choice and action" (2009, 1, emphasis hers). See also e.g. Ferrero 2009, 304, Velleman 2009, 137, and Katsafanas 2013, 47.

4.1 Non-arbitrary selection

Return to Sticky Situation, at the moment when you have concluded that morality and prudence conflict in your current circumstances. There are many ways you could react to this perceived conflict. For example, you might flip a coin: heads for morality, tails for prudence! Or you might *plump* for following morality. Or you might shrug your shoulders, ignore both prescriptions, and look for a beer. In each of these cases, you *select an option* in the face of the conflict. On the most natural reading of Sticky Situation, you do none of these things. Instead, you ask yourself what you ought to do in order to then select an option on the basis of your answer.

Focus on this natural reading of the case, and consider the distinctive nature of the activity you are engaged in. A straightforward gloss on this activity is that you are seeking to select the option that you practically ought to select. However, this gloss is also relevantly *uninformative*, if we assume that we do not yet have a grip on the concept PRACTICAL OUGHT. We do better by focusing on the contrast with coin-flipping and plumping. What are you trying to achieve in Sticky Situation, that you would not be guaranteed to achieve if you selected an option in one of these ways? It is striking that these look like entirely *arbitrary* ways of selecting an option. Arbitrary selection is an ordinary part of human life, but in Sticky Situation you take there to be a striking conflict between significant norms that bear on your action. This is exactly the sort of context where we might seek to avoid arbitrary selection. This in turn suggests a natural and informative characterization of the activity you are engaged in, in Sticky Situation: you are seeking to select an option in a *non-arbitrary* way.

Plumping and coin flipping are obvious cases of arbitrary selection. Consider a series of cases which illustrate less obvious arbitrariness. Suppose first that in Sticky Situation Zoe takes there to be a conflict between morality and prudence, and then selects an option on the basis that it is the most prudent. By selecting on the basis of a norm, Zoe avoids the most obvious sort of arbitrariness. However, suppose that Zoe had initially selected the prudence norm as the norm to guide her deliberation by flipping a coin. This would constitute arbitrariness in the etiology of reasoning that led to the resolution of the conflict. Suppose that

instead of flipping the coin, she had reasoned as follows: *always following the prudence norm is the prudent thing to do, so I shall do it*. Here there is no overt picking or coin-flipping in the background. But this reasoning ignores a troubling symmetry: perhaps always following the moral norm is the moral thing to do.¹⁶ Upon noticing this symmetry, a reasoner would either need to plump for resolving the conflict by deploying the prudential norm (which is arbitrary), or she would need to find some further basis for choice between the two norms. The same point applies to a reasoner who arbitrarily picks a third norm to adjudicate the conflict. Suppose, for example, that Yan plumps for a policy of appealing to etiquette to guide his choice in cases like Sticky Situation.

These examples show that selection can count as arbitrary in virtue of the etiology of that selection involving either (a) relevant arbitrary picking or (b) a failure to even consider a relevant normative conflict.¹⁷ This gloss on ARBITRARY SELECTION may be incomplete, but together with the examples, I take it to provide readers with a substantial grasp of the concept. For example, I think it should allow a careful reader to confidently categorize novel cases as involving arbitrary selection, or not.

I aim to analyze the concept PRACTICAL OUGHT in terms of ARBITRARY SELECTION. Because of this, it is important that my account of ARBITRARY SELECTION not covertly deploy the concept PRACTICAL OUGHT, on pain of circularity. One might worry that the account just sketched fails this condition: that our only grip on the concept NON-ARBITRARINESS is parasitic on our grasp on PRACTICAL OUGHT. In reply, consider how I introduced the concept ARBITRARY SELECTION. I pointed to recognizable paradigms, explained how to apply the concept to more complex cases, and proposed an informative tentative gloss on it. None of these elements advert to the concept PRACTICAL OUGHT, and I claim that they jointly put the reader in a position to understand and apply the concept ARBITRARY SELECTION. It

¹⁶ The symmetry suggested here is intended only as an illustration. If we think of *following a norm* as an intentional activity guided by a representation of the norm, it is unlikely that either morality or prudence always endorses following itself, due to familiar 'rational irrationality'-style phenomena.

¹⁷ The mention of relevance here gestures at a pattern familiar from other contexts. For example, anti-luck epistemologies do not object to knowledge acquired via luckily acquired evidence. Similarly, if I flip a coin, and on that basis decide to non-arbitrarily select an option, which I then go on to do, the coin flipping is outside of the relevant scope of the deliberation.

is thus unclear how the concept PRACTICAL OUGHT could covertly enter this account.

A second natural worry is that the *activity* of non-arbitrary selection that I aim to elucidate doesn't make sense. Distinguish two versions of this worry. First, one might worry that the activity makes no sense because we cannot make sense of the concept of ARBITRARINESS. This version of the worry can again be answered by pointing to my constructive gloss on the concept. The second version of the worry is that the concept NON-ARBITRARY SELECTION is coherent but necessarily empty. On this version, seeking to make a non-arbitrary selection is like seeking to identify the largest prime: you can engage in this activity, but it is wrongheaded to do so, because you will necessarily be unsuccessful. While I am optimistic that this worry can be answered, I do *not* aim to do so in this paper. My aim here is to analyze the *concept* PRACTICAL OUGHT, and such analysis does not require that successful non-arbitrary selection is possible.

4.2 A preliminary account of the concept PRACTICAL OUGHT

Recall that my aim here is not necessarily to provide an analysis of a folk concept. Rather it is to provide an account of the concept PRACTICAL OUGHT that vindicates its apparent philosophical significance in debates about Moral Rationalism, and especially in cases like Sticky Situation. I have just argued that in Sticky Situation one is engaged in the activity of *non-arbitrary selection*. My account is motivated by a pair of hypotheses. First, what it is for the concept PRACTICAL OUGHT to be authoritatively normative is (in part) for it to be the concept of a norm that is distinctively appropriate to appeal to in cases like Sticky Situation. Second, the constitutive norms for the activity of non-arbitrary selection satisfy this description.

The first step in developing an account from this intuitive idea is to return to the contrast between constitutive correctness conditions and constitutive success conditions. Consider addition as an example. If I seek to add 17 and 34, I add them *correctly* if, without error, I use an appropriate mathematical procedure to arrive at the sum. On the other hand, suppose that in adding the two numbers I

make a pair of errors, on the basis of which I am lucky to conclude that the sum is 51. I clearly did not add the numbers correctly here, in light of these errors. However, I did add them *successfully*: my goal was to identify the sum, and I did that.

In some cases (like simple addition) correctness ensures success: I cannot correctly add two numbers and get the wrong answer. There are other constitutive norms where correctness fails to ensure success. For example, on a toy constitutivist epistemic theory, to believe a proposition correctly is to believe it on the basis of sufficient evidence. And for this belief to be successful is just for the proposition to be true. Fallibilism about sufficient evidence allows that in cases of misleading evidence, one can believe correctly but unsuccessfully.

We can distinguish constitutive correctness and success for the activity of non-arbitrary selection. Consider a motivating example: suppose that in Sticky Situation you sought to make a non-arbitrary selection, and thereby engaged in some reasoning that led you to stay and help. Later, you might revisit your choice, and conclude (a) that your reasoning was fallacious, but (b) that the correct reasoning in Sticky Situation would have led you to the same conclusion. It is natural here to take yourself to have succeeded in your constitutive aims in Sticky Situation, despite having engaged in the activity incorrectly.¹⁸ The question of whether correctness entails success for practical norms is interesting and substantive, and I will not settle it here.¹⁹

I propose to analyze the concept PRACTICAL OUGHT in terms of the constitutive norms for the activity of non-arbitrary selection. To be plausible, this analysis must appeal to constitutive *success* as opposed to *correctness* conditions. For example, in the fallacious reasoning case just above, it would be perverse to insist that because you selected the right option on the wrong basis, you failed to

¹⁸ One implication of this clarification is that the name I have given this activity - *non-arbitrary selection* - is slightly misleading, because it references the correctness conditions of the activity, and not its success conditions.

¹⁹ One model for entailment failure would be if correctness (but not success) is information-relative. Note as well that in rational-irrationality style cases, success may be directly *incompatible* with correctness. For example, a modest assumption about the relevant success conditions will entail that if an evil demon will torture everyone forever if I do whatever is required for correct non-arbitrary selection in C, then successful selection in C will require incorrect selection.

do what you ought to do. This puts me in a position to propose the following conceptual truth:

Conditional when an instance of the activity of non-arbitrary selection concludes in selection of an option A, for the agent's judgment *I practically ought to do A* to be true is for the selection to have satisfied the constitutive success norms of that activity.

Because Conditional provides only a sufficient condition, it falls short of being a conceptual analysis. However, it allows me to now sketch how the analysis that builds on Conditional can vindicate the distinctive normative authority that the concept PRACTICAL OUGHT purports to have.

This vindication begins with the fact that cases like Sticky Situation are especially powerful tools for orienting us to authoritative concepts like PRACTICAL OUGHT. A compelling diagnosis of this fact is that Sticky Situation is a case of perceived conflict among intuitively important norms. And conflict between such norms makes the need for a *non-arbitrary* basis for selection especially vivid. I claim that this sort of case is crucial to orienting us to PRACTICAL OUGHT exactly because the distinctive normative authority of this concept is authority in the context of non-arbitrary selection. More precisely, the activity of non-arbitrary selection is the fundamental context in which the question of the relative authority of different normative standards arises. Our grip on relative normative authority in other contexts is *parasitic* on the relationship that normative standards bear to non-arbitrary selection. This suggests a job description for the concept PRACTICAL OUGHT: it is the concept of a norm that is *the norm to appeal to* in the context of non-arbitrary selection. The fact that the concept PRACTICAL OUGHT satisfies this job description is what *constitutes* its distinctive authoritativeness.

The deflationary pluralist introduced in §1 has a seemingly powerful retort to the initial gloss on the vindication just sketched: what is the relevant priority suggested by the talk of the 'one to appeal to' here? As in any context, morality arguably has *moral priority* in cases of non-arbitrary selection, prudence arguably has *prudential priority*, etc. So what could talk of a norm being the 'one to appeal to' be here, except a cheat?

In response, consider the constitutive success norm for the activity of non-arbitrary selection. If you are engaged in this activity, and your selection satisfies this norm, you will thereby have concluded your activity *successfully*. And this is a very natural gloss on what it is for a norm to be the *one to appeal to* in the context of this activity. Returning to Sticky Situation makes this vivid. I have argued that in this scenario you are engaged in the activity of non-arbitrary selection. Suppose for concreteness that you were convinced that the constitutive success norms for this activity require you to always privilege morality over prudence. And suppose that you knew that you are engaged in this activity. There would then be little to intelligibly deliberate about: you already take yourself to know both what morality requires, and that if you complete your deliberations successfully, you will do what morality requires. And this in turn shows why Conditional is so plausible as a conceptual truth. In Sticky Situation, you ask the question *what ought I to do?* to identify the selection that would constitute success in you activity. Given the role of that question in that context, it is hard to see how you could conclude that activity successfully without doing what you practically ought to do.

This concludes my provisional account of what it is for a concept to be authoritatively normative. One reason the account is provisional is that Conditional – the conceptual truth I proposed above – only identified a sufficient condition for the truth of the judgment schema *I practically ought to do A*. The next section takes up two central challenges that arise en route to completing the conceptual analysis.

4.3 Scope and the conditional fallacy

The aim of this section is to complete the analysis of the concept PRACTICAL OUGHT, explaining how it applies to agents not engaged in the activity of non-arbitrary selection. I begin by suggesting a natural counterfactual extension of my account. I then address two important objections to this extension. The first concerns the scope of application of constitutive norms, and the second is an instance of the ‘conditional fallacy’ style of objection to counterfactual accounts.

We do not always engage in the activity of non-arbitrary selection. For example, we plump or make unreflective decisions all of the time. (When I argued in §4.1 that a natural way to interpret Sticky Situation involved non-arbitrary selection, I did so in part by *contrasting* that activity with other things we might do.) It is plausible that in many cases of unreflective decision, there is something that one ought to do. In any case, no account of authoritatively normative concepts should rule this out. In order to satisfy this constraint, I must extend Conditional to provide an analysis that can apply to such decisions. I do so by offering a natural *counterfactual* extension of my account. An initial gloss on the extension is this: in cases where I act unreflectively, for it to be true that I practically ought to have done A, is for it to be that *were I to have successfully non-arbitrarily selected*, I would have done A.

An immediate objection to this proposal is that my *counterfactual* participation in an activity governed by constitutive norms does not make those norms *actually* apply to me. For example, *were I playing chess right now*, and I touched my Queen, the rules of chess would require me to move that piece. But because I am not playing chess right now, touching a chess piece entails no such requirement. So, the objection continues, why should facts about counterfactual non-arbitrary selection have any normative significance for me in a situation when I do not engage in this activity?

I reply to this objection by emphasizing the contrast between the account of normative authority that I offer, and the more standard form of constitutivist explanation. To simplify brutally, this standard explanation is that normative authority *just is* the grip that constitutive norms have on the participants of the activities they govern, for an activity we cannot escape. My account of normative authority is different. Put abstractly, the standard form of constitutivist explanation locates normative authority in the relation of a norm to an agent (inescapability), while my account locates normative authority in the nature of the activity itself, independently of its relationship to any actual agent. On my account, as sketched above, what makes a norm distinctively authoritative is that it is distinctively the *norm to appeal to* in the context of non-arbitrary selection. Crucially, the fact that a norm has this status is independent of whether anyone is

actually engaged in the activity of non-arbitrary selection. Thus, on my account, authoritative norms can lack any ‘grip’ on actual agents: they can apply to agents who are wholly and coherently indifferent to its prescriptions.

A second natural worry about my counterfactual proposal is that it is vulnerable to the *conditional fallacy* (e.g. Shope 1978, Johnson 1997). The basic worry is this. Suppose (à la Conditional) that any time someone actually successfully non-arbitrarily selects, she selects the option that she ought to perform. Now consider a case in which I do not seek to non-arbitrarily select. It seemingly might be that the nearest possible world in which I seek to do so is relevantly different from my actual circumstances, such that while I *ought* to perform a certain action in that world, I *ought not* to perform it in my actual circumstances. If this were possible, then the counterfactual extension of my account would deliver the wrong results.

A theoretically contentious example may make the worry more vivid. My view about the concept PRACTICAL OUGHT should be compatible with familiar *Humean* theories about the substance of what we practically ought to do. According to these theories, facts about what one ought to do in a context are a (non-constant) function of one’s desires in that context. Consider the thesis that I ought now to deliberate, supposing that I currently have no desire to do so. The closest worlds in which I seek to non-arbitrarily select concerning this question are worlds in which I deliberate, and hence, in which I have at least some desire to deliberate. My counterfactual account thus seemingly might entail that I ought now to deliberate, in light of these counterfactual desires, despite my actually lacking any relevant desire. This is inconsistent with the familiar Humean thesis.

Crucially, my account is only vulnerable to the conditional fallacy if the constitutive success norms for the activity of non-arbitrary selection have a certain structural feature: that success conditions can vary with features of the context of *assessment*. This feature is absent from some familiar constitutive norms. To see this, consider chess: holding fixed a certain chess position, every correct assessment of the legal moves available to black in that position will be the same, no matter the context of assessment. It is plausible that we are conceptually

committed to the constitutive success norms for non-arbitrary selection exhibiting a similar invariance.

To see this, focus on a case where an agent engages in non-arbitrary selection *prospectively*. For example, suppose that yesterday I attempted to non-arbitrarily select an option to perform now, and that I now recheck that reasoning. I would be confused if I believed now that I had successfully non-arbitrarily selected yesterday, but that I might get a different answer if I successfully non-arbitrarily select now. The same thing holds for possible cases of *retrospective* or *hypothetical* non-arbitrary selection. A retrospective example: sitting in jail after staying to help, you might ask yourself: *should I have taken the money and run?* Or, reading this paper you might engage in non-arbitrary selection *hypothetically*, considering what to do were you in Sticky Situation. When we imagine such cases, we notice that a form of invariance is built into the concept NON-ARBITRARY SELECTION: at any two contexts of evaluation where I successfully non-arbitrarily select an option for a single context of action, I *must* come to the same conclusion. Call this property of the concept NON-ARBITRARY SELECTION *intrapersonal invariance*.²⁰

This property is plausible when combined with my background account of the nature of normative authority. Suppose that you were convinced by the argument of §4.2 that it is a conceptual truth that successful non-arbitrary selection in the context of action entails that one does what one ought to do in that context. And now consider the sort of case just envisioned, where I had previously engaged in successful non-arbitrary selection prospectively, for my current circumstances. It would be perverse to think that the mere fact that I engaged in the activity prospectively prevents me from drawing from this reasoning a conclusion about what I practically ought to do now. The same goes for beliefs about successful retrospective and hypothetical non-arbitrary selection: these are well understood as conceptually entailing conclusions about what I hypothetically or retrospectively ought to have done. This fits in with my account

²⁰ Note that *intrapersonal invariance* is compatible with the constitutive success conditions for non-arbitrary selection varying across *evaluators*. I am attracted to *interpersonal invariance* here as well, but nothing in this paper depends on that.

of normative authority: in short, the *question* of distinctive normative authority arises in the context of the activity of non-arbitrary selection, but we can answer such questions *about* decisions reached outside of that context.

The conditional fallacy worry arises because of the apparent possibility that success conditions for the activity of non-arbitrary selection *concerning* a certain context of selection might vary with features of the counterfactual context of deliberation. But intrapersonal invariance blocks exactly this possibility, and so blocks the conditional fallacy from arising for my account.

This puts me in a position to refine my counterfactual extension into a conceptual analysis of PRACTICAL OUGHT:

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

In §1 I suggested that the thought that a concept is authoritatively normative is *elusive*. My account neatly explains this elusiveness. The connection between PRACTICAL OUGHT and NON-ARBITRARY SELECTION is not obvious; I have argued that it becomes vivid only when we think carefully about scenarios like Sticky Situation. Further, only with conditions like intrapersonal invariance on the table is it possible to see how an account of PRACTICAL OUGHT in terms of NON-ARBITRARY SELECTION can have adequate scope. If Constitutive is the correct analysis of PRACTICAL OUGHT, it is very natural to think that the concept would appear highly elusive, seemingly supporting primitivism or deflationary pluralism.

5. Schmagency

My account offers a kind of constitutivist explanation of the distinctive normative authority of the concept PRACTICAL OUGHT. The most powerful contemporary objection to this broad explanatory strategy is David Enoch's powerful 'schmagency' challenge. I thus now consider two ways of attempting to adapt Enoch's challenge to apply to my account.

As I read it, the core of Enoch's challenge begins with the constitutivist's insight that our various ordinary motivations are too arbitrary a basis to ground normativity. By appealing to features that are constitutive of action (for example),

the constitutivist aims to eliminate this objectionable arbitrariness (2006, 178). Against this, Enoch argues that the problematic arbitrariness cannot be eliminated by the inescapability of the relevant activity (2006, §6), or by the related idea that we (can't help but) care about agency or its constitutive norms (2011, 212-13), or by insisting that one cannot even raise doubts about normativity except in a context where one is already committed to certain constitutive norms (2011, §6).

Because I reject the mechanisms for arbitrariness reduction that Enoch considers on behalf of the constitutivist, I can simply embrace the arguments just mentioned. However, it is worth addressing two important ways of adapting Enochian ideas to challenge my account. Both challenges focus on the concept NON-ARBITRARINESS that is central to my analysis.

First, it may seem that I face a dilemma. The concept NON-ARBITRARINESS will either:

- (a) be too thin to generate facts about what we ought to do, or
- (b) be thick enough to prompt reasonable normative resistance: why care about non-arbitrariness? (This dilemma is inspired by Enoch 2010, 213.)

The first horn fails against my account, simply because my account is not committed to there being such facts. As I noted in §4.1, my aim is to analyze a concept, not to show that we can derive substantive norms from that analysis. The second horn is a variant of the central 'schmagency' idea that Enoch prosecutes: he imagines someone who is 'normatively indifferent' to the alleged constitutive features of agency. The possibility of such indifference is no challenge to my account, for, as I noted in the previous section, nothing in my account entails that *anyone* cares about what they ought to do (or non-arbitrariness). It is also true that a reasonable agent could resist the explanation of normative authority that I sketched in §4.2. But again, this is compatible with my account: I have argued that this account is *true*, not that it must be accepted by any reasonable agent.

A second schmagency-style challenge starts by supposing that there are a set of slightly different arbitrariness concepts: ARBITRARINESS_A, ARBITRARINESS_B etc., which agree on central cases but diverge at the penumbra. These could seemingly be referenced by the constitutive norms of a set of activities: *non-arbitrary*

selection_A, etc. A schmagent might thus ask: why go in for *non-arbitrary selection_A*, rather than another variant activity?

The reply to this challenge begins by observing that the aim to make a non-arbitrary selection is incompatible with the existence of such higher-order arbitrariness. If the choice between the activities *non-arbitrary selection_A* and *non-arbitrary selection_B* were itself arbitrary, then *neither* of these activities satisfies our concept of non-arbitrariness. This shows that it is a conceptual constraint on the concept of NON-ARBITRARINESS that this concept is *privileged* in some way over nearby notions. I favor a picture where the presupposed privileging is metaphysical: where we take the arbitrary/non-arbitrary distinction to be part of the genuine structure of the world. In adverting to non-arbitrariness, we aim to be latching on to that structure. If it turns out that there is no such structure – or that there are multiple candidate such structures that we might be talking about – then the uniqueness assumption of the activity is undercut, and the result will again likely be that the concept PRACTICAL OUGHT is necessarily empty. This would be an unappealing result, but it does not threaten the conceptual analysis itself.

6. Consequences of the analysis

This paper has aimed to illuminate the nature of authoritatively normative concepts. The centerpiece of my account is an analysis of the concept PRACTICAL OUGHT, which draws in a novel way on some resources from the constitutivist tradition:

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

I have argued that this analysis permits an elegant explanation of what the distinctive authoritativeness of the normativity of this concept amounts to. Roughly: what makes these concepts distinctively authoritatively normative is that they are (constitutively) the norms to appeal to when engaged in the activity of non-arbitrary selection. I also suggested that this account permits an attractive explanation for why the nature of this authoritativeness has seemed so elusive.

Besides these intrinsic merits, my account illustrates the flexibility of constitutivist metanormative accounts. For example, I have just argued that my account escapes Enoch's schmagency challenge to constitutivist accounts of normative authority. One might reasonably regret that this flexibility is purchased at the cost of reduced ambition: for example, there is no hint in my account of an attempt to derive any substantive normative conclusions – let alone the categorical imperative – from constitutivist materials. However, I take it to be a mistake to think that constitutivism must accomplish everything its proponents have dreamed of, in order to deserve our attention or assent.

I conclude by briefly sketching some implications of my account for broader metaethical theorizing. Consider first the dialectic between the metaethical cognitivist and the non-cognitivist (where, roughly, the cognitivist claims that, at the fundamental explanatory level, the thought that I ought to do A is a belief). One way to motivate the non-cognitivist view is to insist that primitivism about authoritatively normative concepts is unacceptable. Once this is done, the non-cognitivist conjectures that what marks off authoritatively normative thoughts as distinctive is that they are constituted by (for example) Gibbardian planning states. One dialectical strength of the non-cognitivist view, if we grant the objection to primitivism, is that plausible and non-trivial cognitivist accounts of fundamental authoritatively normative concepts are thin on the ground. However, with my account in hand the dialectic shifts. The cognitivist can point out that her view accommodates the plausible thought that authoritatively and merely generic concepts are species of a unified conceptual genus: the normative concepts. The non-cognitivist, by contrast, faces pressure to posit a deep and not obviously plausible discontinuity in the fundamental nature of authoritative and merely generic normative concepts, on pain of offering an implausible non-cognitivist gloss on chess norms, for example.

My account also helps the cognitivist with respect to the dialectic concerning normative disagreement. R. M. Hare (1952) and Terence Horgan and Mark Timmons (1992 etc.) have argued persuasively that genuine moral disagreement seems possible between members of linguistic communities whose moral thought seems to track different properties. Such disagreement phenomena

seems to carry over with even more force in the case of authoritatively normative thoughts. My account can potentially help to explain how this fact is compatible with cognitivism: my account analyzes the authoritatively normative concepts in terms of an extremely thin functional role. As long as a linguistic community's word behaves in accordance with this role, it will express the same concept.

My account can also help to defend particular cognitivist metaethical projects against some of the characteristic challenges that they face. Consider first the way it can be useful to the metaethical *error theorist* and the *fictionalist*. Schematically, the most straightforward way of arguing for error theory about the authoritatively normative is as follows:

Conceptual claim:	All authoritatively normative claims commit us to the existence of X
Metaphysical claim:	<u>X does not, or could not, exist</u>
Conclusion:	Authoritatively normative claims are systematically erroneous

One central difficulty for the error theorist is to provide a plausible version of the conceptual claim (Finlay 2008). Illustrative here is the derision rightly heaped on John Mackie's claim that it is part of the ordinary concept of objective value that such values somehow magically *make us pursue them* (1977, 40). My account can help the error theorist here. For if she accepts my conceptual claim, she can argue as follows:

Conceptual claim:	the concept of an authoritative norm is the concept of the constitutive success norm for the activity of non-arbitrary selection
Metaphysical claim:	<u>There are no constitutive success norms for this activity</u>
Conclusion:	Authoritatively normative claims are systematically erroneous

Metaethical fictionalists face an analogue of the error theorist's conceptual challenge. In this case, the task is to tell us what, at least roughly, makes a fiction a fiction about *authoritative normativity* (compare Hussain 2004). My account can again help. The fictionalist could treat the claim that there are success norms for the activity of non-arbitrary selection either as a characterization of the content of the fiction that we already implicitly accept (for the hermeneutic fictionalist), or as an account of the fiction that we should adopt (for the revolutionary fictionalist).

As I have been emphasizing, my account is neutral concerning whether there are any authoritatively normative facts. So it does not by itself promise to vindicate normative realism. And while the account is compatible with the strategy of arguing for a constitutivist derivation of normative content, I am pessimistic about the prospects for that strategy. Readers who join me in this assessment may think that the account of the concept PRACTICAL OUGHT offered here makes it hard to see how that concept could have a non-empty intension.

I will close with the barest sketch of the picture that makes me optimistic on this front. I am compelled by a thesis David Lewis (1983, 1984) proposed about reference determination: that a factor in determining reference is the *eligibility* of candidate referents. Adapted to content-determination, the idea is that the content of a concept is a function of two factors: how well the candidate content satisfies the functional role of the term, and how well the candidate content carves nature at its joints. A natural way of characterizing the metaphysical commitment of normative realism is this: that the authoritative norms are distinctive joints of nature among the space of possible normative systems. The functional role suggested by Constitutive is tremendously thin. Being the distinctive joint of nature among the space of normative systems is (a) a non-arbitrary basis for selection, and (b) an eligibility maker. On this view, the determinate and stable reference of the concept PRACTICAL OUGHT is a function of the thin conceptual role suggested by Constitutive, the realist's distinctive metaphysical claim that there is a distinctive normative joint of nature, and (an appropriately developed version of the sketched theory of content-determination. This is, of course, the barest of sketches of an ambitious view (for more detailed discussion of related metasemantic ideas in a metaethical context, see Dunaway and McPherson *forthcoming*). However, I think it is highly promising in broad outline, and it at least suffices to show why accepting Constitutive does not force the metaethical realist to abandon all hope.

Works Cited

- Baker, Derek. Ms. "Intuitions about Disagreement Do Not Support the Normativity of Meaning."
- . *This Volume*. "Skepticism about Ought Simpliciter."
- Bedke, Matt. 2010. "Might All Normativity be Queer?" *Australasian Journal of Philosophy* 88.1: 41-58.
- Copp, David. 2005a. "Moral Naturalism and Three Grades of Normativity." *Normativity and Naturalism*. Peter Schaber, ed. Frankfurt:Ontos-Verlag. 7-45.
- . 2005b. "The Normativity of Self-Grounded Reason." *Social philosophy and Policy* 22: 165-203.
- Cullity, Garrett, and Berys Gaut (Eds.) 1997. *Ethics and Practical Reason*. Oxford: Clarendon.
- Darwall, Stephen. 1997. "Reasons, Motives, and the Demands of Morality." *Moral Discourse and Practice*. Eds. Stephen Darwall, Allan Gibbard, and Peter Railton. New York: Oxford UP. 305-312.
- Davidson, Donald. 2001 [1969]. "How is Weakness of the Will Possible?" *Essays on Actions and Events*, second edition. Oxford: Oxford UP. 21-42.
- Dunaway, Billy and Tristram McPherson. *Forthcoming*. "Reference Magnetism as a Solution to the Moral Twin Earth Problem." *Ergo*.
- Eklund, Matti. *Forthcoming*. *Choosing Normative Concepts*. Oxford: Oxford UP.
- Enoch, David. 2006. "Agency, Schmagency." *Philosophical Review* 115.2, 169-198.
- . 2011. "Schmagency Revisited." *New Waves in Metaethics*. Ed. Michael Brady. Houndsmills: Palgrave. 208-233.
- Ferrero, Luca. 2009. "Constitutivism and the inescapability of agency." *Oxford Studies in Metaethics* Volume 4. 303-333.
- Finlay, Steven. 2008. "The Error on the Error Theory." *Australasian Journal of Philosophy* 86.3: 347-369.
- FitzPatrick, William. 2008. "Robust Ethical Realism, Non-Naturalism and Normativity." *Oxford Studies in Metaethics* Vol. 3. Ed. Russ Shafer-Landau. Oxford: Oxford UP.
- Foot, Philippa. 1997. "Morality as a System of Hypothetical Imperatives (including her 1994 'Recantation')." Darwall, Gibbard, and Railton 313-22.

- Gibbard, Allan. 2003. *Thinking how to Live*. Cambridge, MA: Harvard UP.
- . 2012. *Meaning and Normativity*. Oxford: Oxford UP, 2012.
- Hampton, Jean. 1998. *The Authority of Reason*. Cambridge: Cambridge UP.
- Hare, R. M. 1952. *Language of Morals*. Oxford: Oxford UP.
- Horgan, Terence and Mark Timmons. 1992. "Troubles for New Wave Moral Semantics: The Open Question Argument Revisited." *Philosophical Papers* 21.3: 153-175.
- Hussain, Nadeem. 2004. "The Return of Moral Fictionalism." *Philosophical Perspectives* 18: 149-187.
- Jackson, Frank. 1998. *From Metaphysics to Ethics*. Oxford: Oxford UP.
- Johnson, Robert. 1997. "Reasons and Advice for the Practically Rational." *Philosophy and Phenomenological Research* 57.3: 619-625.
- Jones, Karen and Francois Schroeter (eds.) *Forthcoming. The Many Moral Rationalisms*. Oxford: Oxford UP.
- Katsafanas, Paul. 2013. *Agency and the Foundations of Ethics*. Oxford: Oxford UP.
- Korsgaard, Christine. 2009. *Self-Constitution: Agency, Identity, and Integrity*. Oxford: Oxford UP.
- Kripke, Saul. 1982. *Wittgenstein on Rules and Private Language*. Cambridge, MA: Harvard UP.
- Lewis, David. 1983. "New Work for a Theory of Universals." *Australasian Journal of Philosophy* 61.4: 343-377.
- . 1984. "Putnam's Paradox." *Australasian Journal of Philosophy* 62.3: 221-236.
- Lord, Errol, and David Plunkett. *Forthcoming*. "Reasons Internalism." *Routledge Handbook of Metaethics*. Eds. Tristram McPherson and David Plunkett. Routledge.
- Mackie, J. L. 1977. *Ethics: Inventing Right and Wrong*. Harmondsworth: Penguin.
- McPherson, Tristram and David Plunkett. *Forthcoming-a*. "Morality, Justice, and Other Mid-Level Normative Categories." *Conceptual Ethics and Conceptual Engineering*. Eds. Alexis Burgess, Herman Cappelen, and David Plunkett.
- and ---. *Forthcoming-b*. "The Nature and Explanatory Ambitions of Metaethics." *Routledge Handbook of Metaethics*. Eds. Tristram McPherson and David Plunkett. Routledge.

- Moore, G.E. 1993 [1903]. *Principia Ethica*. Revised Edition. Ed. Thomas Baldwin. Cambridge: Cambridge UP.
- Parfit, Derek. 2011. *On What Matters*. Oxford: Oxford UP.
- Railton, Peter. 1989. "Naturalism and Prescriptivity." *Social Philosophy and Policy* 7: 151-174.
- Scanlon, Thomas. 1998. *What We Owe to Each Other*. Cambridge, MA: Belknap.
- Schroeder, Mark. 2007. *Slaves of the Passions*. Oxford: Clarendon.
- Shope, Robert. 1978. "The Conditional Fallacy in Contemporary Philosophy." *Journal of Philosophy* 75.8: 397-413.
- Silverstein, Matthew. 2015. "The Schmagency Question." *Philosophical Studies* 172: 1127-1142.
- Smith, Michael. 1994. *The Moral Problem*. Oxford: Blackwell.
- Street, Sharon. 2008. "Constructivism about Reasons." *Oxford Studies in Metaethics* Vol. 3. Ed. Russ Shafer-Landau. Oxford: Oxford UP. 207-246.
- Tiffany, Evan. 2007. "Deflationary Normative Pluralism." *Canadian Journal of Philosophy*, Supp. Vol. 33: 231-262.
- Van Roojen, Mark. 2010. "Moral Rationalism and Rational Amoralism." *Ethics* 120.3: 495-525.
- Velleman, J. David. 2009. *How We Get Along*. Cambridge; Cambridge UP.
- Wedgwood, Ralph. 2004. "The Metaethicist's Mistake." *Philosophical Perspectives* 18.1, 405-426.
- . 2007. *The Nature of Normativity*. Oxford: Oxford UP.