

Conceptual Ethics and the Methodology of Normative Inquiry

Tristram McPherson and David Plunkett

Introduction

One of the striking features of normative theorizing in philosophy (as well as in related fields, such as political theory) is the diversity of concepts that feature centrally in it. In particular, it is commonplace for different theorists to offer different glosses on the overarching normative questions they are interested in, using what appear to be distinct concepts.

Consider, for example, broadly *ethical* inquiry. Suppose we focus our attention on a specific agent in a completely determinate set of circumstances. Normative inquirers ask:

- what the agent *ought to do, all things considered*
- what the agent *has most normative reason* to do
- what it would be *immoral* for her to do
- what it would be *rational* for her to do
- whether one of her options would constitute *injustice, or exploitation, or betrayal*
- etc.

Something similar holds for broadly *epistemic* inquiry, where inquirers ask, of a given agent, in a given set of circumstances:

- what she *ought all things considered to believe*
- what she *knows*
- what she has *adequate epistemic justification* to believe
- what it is *rational* for her to believe
- whether her beliefs have been formed in an *epistemically responsible way*
- etc.

A similar variety of questions are posed by inquirers in aesthetics and political philosophy. In all of these cases, normative inquirers frame their investigations in terms of a range of (what appear to be) different normative concepts.

Even if we bracket the idea that normative inquirers *do* frame their investigations using different normative concepts, it is very plausible that they *could*. And given this possibility, normative inquirers face an interesting set of questions concerning what

normative concepts they *should* use. This is a central part of what we call the *conceptual ethics* of normative inquiry. We take *conceptual ethics* to encompass a range of issues about the normative and evaluative assessment of concepts and words. If we focus on a specific agent in a specific context, two central questions in conceptual ethics are: *which concepts* that agent should use, and *which words* she should use to express those concepts.¹

The aim of this chapter is to explore two main questions in the conceptual ethics of normative inquiry. The first question concerns whether to orient one's normative inquiry around *folk* normative concepts or around *theoretical* normative concepts. For example, KNOWLEDGE and IMMORAL are arguably folk normative concepts: non-philosophers engage in a rich range of thought and talk about what various people do and do not know, and when and why they have been immoral.² By contrast, ADEQUATE EPISTEMIC JUSTIFICATION and PRO TANTO PRACTICAL REASON are arguably technical concepts that have their home in a range of theoretical activities, most markedly systematic epistemic and ethical investigation.

The second question that we explore is whether to orient one's normative inquiry around concepts whose *normative authority* is especially accessible to us, or around concepts whose *extension* is especially accessible to us. For example, the normative authority of OUGHT TO DO ALL THINGS CONSIDERED might seem especially clear, while its extension will seem intensely controversial in many important cases. By contrast, it is relatively clearer what falls in the extension of BETRAYAL, but more controversial how normatively significant betrayal is in many important cases.

In this chapter we do not aim to adjudicate these two questions in conceptual ethics. Instead, we have two central goals. First, we aim to make vivid a range of possible positions that one might occupy with respect to these questions. Attention to this range of options can be valuable for properly grasping the relationships between many existing normative theories, and methodological approaches to normative inquiry. It can also help to reveal approaches to normative inquiry that might otherwise remain obscure. We often survey these options in a relatively neutral manner. However, we do not mean to imply that all of these positions are ultimately equally good. With this in mind, our second goal is to highlight a range of schematic arguments favoring certain options over alternatives. In so doing, we hope to contribute to the long-term goal of adjudicating among those options.³

¹ We here draw from Burgess and Plunkett (2013a) and (2013b). As emphasized in Burgess and Plunkett (2013a), in calling this area of inquiry conceptual "ethics", we do not mean to privilege the idea that broadly moral/political norms matter here more than others. It could be, for example, that we should (at least in certain contexts) use those concepts that best carve reality at its objective joints, regardless of any broadly moral/political norms (e.g., regardless of whether it makes our lives go better or worse).

² In this chapter, terms in small caps (e.g. CAT) pick out concepts. Single quotation marks (e.g. 'cat') are used strictly to mention linguistic items. Double quotation marks (e.g. "cat") are used for a variety of tasks including quoting others' words, scare quotes, and mixes of use and mention.

³ For simplicity, we tend to frame our questions *individualistically*. But normative inquiry is usually a collective endeavor. This raises two important possibilities that we do not discuss in the chapter: (i) that it might be best overall if some groups of individuals oriented their inquiry around one normative concept, while others oriented their inquiry around another; and (ii) that what concept it makes sense for an individual to orient her inquiry around might depend in important ways on what other inquirers are doing.

The question of whether to orient inquiry around folk or theoretical concepts arises in many different kinds of inquiry. By contrast, the question of whether to orient inquiry around *authority-revealing* or *extension-revealing* concepts lacks obvious parallels in other areas of inquiry. As we aim to show, these two issues can interact in interesting ways, making it fruitful to explore them together.

We organize our discussion into four sections. Section 1 focuses on the choice between folk and theoretical normative concepts. Section 2 introduces questions about which words to use to express the folk or theoretical concepts we deploy. Section 3 concerns the choice between authority-revealing and extension-revealing normative concepts. These discussions explore arguments *within* the conceptual ethics of normative inquiry. Section 4 steps back from this focus, to a more theoretical question about the very practice of engaging in reflection about conceptual ethics. In particular, it explores the question of how we should understand—as well as how we should *choose*—the normative standards and concepts that we deploy when engaging in conceptual ethics.

1. Folk vs. Theoretical Concepts

This section concerns the question of whether normative inquiry should be oriented around folk concepts or theoretical concepts. For example, in theorizing about norms for action, should we focus on trying to understand which actions satisfy the folk concept MORAL RIGHTNESS or instead Allan Gibbard's theoretical concept THE THING TO DO?⁴ We canvass several reasons to favor using folk normative concepts, and theoretical normative concepts, respectively. Before proceeding, however, we first clarify several key assumptions that guide our discussion.

First, the issue we are focusing on here is not whether to *use* folk or theoretical concepts in normative inquiry. We take it that we will likely want to use both such concepts in any reasonably developed inquiry. The question is rather: which such concepts should we use to *state* our central theoretical questions, and the answers to those questions that we seek? When a concept plays both of these roles in an inquiry, we will say that the inquiry is *oriented* around that concept.

Second, we are investigating a question in the conceptual ethics of *normative inquiry*. As we shall understand it, this encompasses inquiry into what agents should do, think, or feel, and into what should or ought to be the case. We also take it to encompass inquiry into *evaluative* questions, concerning what is good or bad, valuable or disvaluable; *deontic* questions, concerning what is permissible, required, or forbidden; and *aretaic* questions, concerning what is virtuous or vicious. Further, for our purposes here, we also treat it as encompassing inquiry oriented around concepts such as COURAGEOUS or POLITE, whose status as normative is more controversial than the sets of questions just canvassed.⁵ So understood, normative inquiry occurs in many parts of philosophy, including in ethics, political philosophy, epistemology, aesthetics, legal philosophy, and the philosophy of science, as well as

⁴ For Gibbard's articulation of the concept THE THING TO DO, see Gibbard (2003).

⁵ See Väyrynen (2013) for an important argument for the idea that evaluation might be only pragmatically implicated by use of thick concepts like COURAGEOUS.

outside of philosophy (e.g., in debates about public policy). Given that conceptual ethics is itself a normative topic, this chapter itself involves normative inquiry. (This is something that raises complications that we will return to in section 4.)

Questions of how to individuate concepts and words are highly relevant to the foundations of conceptual ethics.⁶ We cannot, however, responsibly argue for answers to these questions here. Instead we aim to clarify our own usage.

We will think of *concepts* as components of thought, and *words* as linguistic vehicles. Somewhat more specifically, we will take concepts to be individuated by something like their inferential role at the level of thought (for suitable concepts, we will assume that inferential roles determine intensions). And we will take concept-users' inferential dispositions to provide significant (although defeasible) evidence of the inferential roles of the concepts tokened. Turning to words: we will not take words to be individuated by their meanings. For example, we will allow that a single word 'bank' can have one meaning that is about the sides of a river, and another that is about a kind of financial institution. And we will understand the meanings of words (perhaps only at an instance of use) in terms of the concepts they are used to express.

These are regimenting assumptions intended to simplify discussion; much of what we say could be straightforwardly reframed given different assumptions on these matters. For example: if you think that there are two different English words written as 'bank', everything we say could be re-phrased in terms of groups of homophonous but distinct words. Another, perhaps more important example, concerns our use of the term 'concept'. For much of our discussion, what will really matter to us are patterns of inference that are closely associated with the use of a given word (e.g., the word 'moral'). This is because it is these patterns of inference that will be most significant to the effects of the use of a given word in the context of normative inquiry. When we treat these divergent patterns of inference as evidence for the presence of distinct concepts, much of what we want to say could instead be represented in terms of multiple systematic patterns of inference associated with the same concept.

We assume the following rough distinction between *folk* and *theoretical* concepts. Theoretical concepts have their home within a community dedicated to a certain relevant inquiry, and are used for the purposes of that inquiry. For example, consider the concepts one acquires when learning advanced physics, or advanced linguistics. Folk concepts have their home in thought and talk beyond the specialists in a given field of systematic inquiry. These characterizations surely admit of hard and borderline cases. However, they are often clear enough. For example, inquirers sometimes find it useful in their work to deploy the folk concept STRING (as in: "we tied the samples together with some string"), and the folk sometimes use the theoretical concept STRING THEORY. In these cases, we have no difficulty determining which side of the folk/theoretical divide is the natural home of the relevant concept.

⁶ For further discussion, see Burgess and Plunkett (2013b); Cappelen (2018); Braddon-Mitchell (Chapter 4, this volume); Greenough (Chapter 11, this volume); and Haslanger (Chapter 12, this volume).

It is also worth emphasizing that theoretical concepts are diverse. One standard function of distinctively theoretical concepts is to provide manifest precision: it can be theoretically useful to offer stipulations, explications, operative definitions, etc., in order to allow the assessment and transmission of relatively precise theses. Other theoretical concepts are introduced to allow us to aptly discuss and investigate certain important worldly phenomena. Where a worldly phenomenon is imperfectly understood, we may need to choose between explicit precision and tracking the important worldly phenomenon. In light of this, the virtues and vices that we explore below may not apply to all theoretical concepts.⁷

We now examine two classes of reasons to favor orienting normative inquiry around folk concepts. The first is motivated by the idea that it makes sense to inquire into questions that we care about. Consider the questions *do I know anything?* or *is it morally wrong to eat meat?*, as well as many other questions that lead people to become interested in normative inquiry in the first place. It is natural to think that such questions are framed using folk concepts. Given this, if one instead orients one's normative inquiry using theoretical concepts, there is a danger that one will simply have *changed the subject*, and failed to address the question one cared about in the first place.⁸

Notice that this argument starts with the issue of how the *questions* we address in inquiry are framed, rather than the materials that we use to answer those questions. Thus, it is compatible with using highly theoretical concepts in one's *answer* to the question of whether it is wrong to eat meat. However, many who are sympathetic to this line of thinking hold that the answers to questions must also be framed at least partly in terms of the same folk concepts (e.g. KNOWLEDGE, MORALLY WRONG) deployed in the question itself, in order to be an apt answer to *that* question.⁹

This sort of argument suggests an important concern to be attended to. However, there are at least two reasons to be cautious about its force. First, on at least some natural ways of individuating folk concepts, a folk concept may include elements that you do not care about in a given instance when you are deploying that concept. Consider the following example. Carrie Jenkins motivates her exploration of flirting by pointing to the practical significance of judgments about flirting in our lives. She then goes on to offer an account of flirting that proposes (among other things) that the flirter must proceed in a *playful* manner.¹⁰ Suppose that playfulness is part of the folk concept FLIRTING. In many contexts, this aspect of the concept may be irrelevant to what people care about when they deploy that concept. For example, if two people are behaving in ways disposed to raise the salience of romance between them, the question of whether they are doing so in a playful manner might well not be practically significant to them, or to other interested parties. In these contexts, a

⁷ The idea that there is an interesting distinction between folk and theoretical *concepts* might be controversial, especially given certain theories of content. For example, on some views, the content of folk natural kind concepts might be fixed by facts about expert usage. (See, for example, Burge 1979 and Schroeter 2008). But recall how we are using the term 'concept': we think that versions of many of the points here could be reformulated within the sort of theory of content just mentioned, by focusing on contrasting inferential implications of certain uses of words.

⁸ See Strawson (1963) and Jackson (1998) for articulations of this kind of worry.

⁹ See, again, both Strawson (1963) and Jackson (1998). ¹⁰ Jenkins (2006).

person may frame much of her thinking and discussion about this interaction in terms of FLIRTING despite the fact that this concept includes elements that are irrelevant to her interests in this case. She might do so simply because this is the most useful concept for her to deploy, among those concepts readily available in her current conceptual repertoire.

This example brings out a more general issue. One reason why someone, in a given context, might not care about certain aspects of a concept's precise inferential role is that she was primarily interested in discussing a certain *property*, which one could (at least in principle) think about using different concepts. It might be that one's interest in a certain folk concept is explained by the fact that it enables one to discuss or think about that property in a given context. Suppose that we can make sense of goals of inquiry identified at the object-level (e.g., in terms of properties or facts) rather than at the representational-level (e.g., in terms of concepts). This would allow for us to change which concepts we use without a "change in topic" in the sense relevant for that inquiry.¹¹

A further reason to be cautious about this reason to favor folk concepts is that inquiry itself may *change* what you care about. For example, Carlos might start out caring a great deal about MORAL RIGHTNESS. But inquiry might change this attitude: for example, he might become convinced by an account of the history of this folk concept that destroys his interest in deploying it in normative inquiry.¹² One simple example: he might conclude that MORAL RIGHTNESS is an ideological instrument for the ruling class, and that this fact infects either the intension of this concept or its fruitfulness as a locus for inquiry.¹³ This might then in turn motivate him to introduce a theoretical concept in the vicinity of MORAL RIGHTNESS that answers to his *theoretically informed* interests.¹⁴ In general terms, inquiry might reveal that our folk concepts are *defective* in a range of different ways. For example, they might lead us into irresolvable paradoxes.¹⁵ At the same time, inquiry might reveal that we might be able to replace those concepts with others, which preserve the core of what we should care about in the vicinity.

Theoretical inquiry may be misleading some or even much of the time. But, at least in principle, it seems fully capable of providing or revealing reasons for investigating properties that are not picked out by our folk concepts. This is particularly hard to deny given that the line of reasoning above was not exclusively focused on normative

¹¹ For further discussion about preserving "sameness of topic" despite shifts in concepts, see Cappelen (2018) and Thomasson (Chapter 21, this volume).

¹² See Plunkett (2016) for relevant discussion of the significance of conceptual history for the conceptual ethics of normative inquiry.

¹³ For a brief discussion of this sort of broadly Marxist-inspired take on morality, see Sinnott-Armstrong (2006: 208). For helpful critical discussion of Marx's own (more subtle) views on morality, see Wood (1999).

¹⁴ Notice that on some sophisticated theories of folk concepts (e.g., Jackson 1998), the content of our folk concepts might turn out not to vindicate all of ordinary speakers' inferential dispositions for using those concepts. On this sort of picture, even if folk *inferential dispositions* for using a concept are objectionable in some way, the intension of the concept might not be objectionable in this way. Jackson also grants that sometimes theoretical reflection can reasonably lead us to care about something other than the folk concept we start with Jackson (1998: 44–5).

¹⁵ For a clear example of this line of thought with respect to the concept TRUTH, see Scharp (2013).

inquiry as such, or even philosophical inquiry as such. As the toy example of conceptual history above underscores, what we care about in normative inquiry could be affected by *non-normative* theoretical inquiry from areas outside of philosophy. For example, discoveries in sociology, psychology, or climate science might create a context where certain normative questions become especially salient and pressing.

A second reason for orienting inquiry around folk concepts arises from the idea that our ability to learn about certain normative facts might be mediated by trained cognitive capacities, which may be (a) largely implicit and (b) largely accessible via deployment of our folk concepts. For example, decently-raised children internalize a torrent of context-relevant guidance about which acts and arrangements are *fair*. Suppose that such guidance is crucial to building the competence of inquirers to investigate certain broad normative topics. Then seeking to replace such concepts with precise theoretical alternatives risks depriving the inquirer of the ability to use her implicit knowledge.

Similarly, suppose that one's aim in developing a theoretical concept is to capture more precisely what we most care about in the relevant topic. In many cases, it can prove difficult to understand precisely *what* we care about in the topic at hand, let alone make it fully explicit. Given this, there is a substantial danger that in shifting our focus by using a new or "reformed" concept, we will lose the ability to focus on what we cared about in deploying the folk concept—some of which, of course, might also be things we *should* care about. This might in turn lead us to miss important aspects of the normative issues at hand. In short, the switch to theoretical concepts might leave us with worse tools for investigating our normative topics at hand (even if such a switch still preserved "sameness of topic" in relevant ways). Indeed, we might end up with worse tools here without even fully realizing that this is so; something which might well make the shift to the new "reformed" concepts particularly dangerous.¹⁶

Tied to this worry is the question of how good we are at judging the merits and dangers of attempts to depart from conceptual folkways. Perhaps, as a general matter, one takes a dim view of our abilities here. This might then be combined with a strong version of the thesis that our ability to learn about key normative facts is mediated by trained cognitive capacities, which are largely implicit and largely accessible via deployment of our folk concepts. This combination of views might be used to support a kind of intellectual analogue of Burkean conservatism in social/political philosophy, according to which we should largely defer to our current conceptual repertoire, or at least let it change slowly and gradually and be highly skeptical of attempts to radically change it quickly with "conceptual revolutions".¹⁷

The fact that folk concepts are central sites of social learning, however, can also provide grounds for caution about orienting one's normative inquiry around folk concepts. For the influences that shape our implicit grasp of normative concepts need not tend towards accuracy or reliability. For example, moral and political concepts, from IMMORAL to JUST to CHASTE, are exceptionally natural targets for

¹⁶ This sort of criticism is prominent in Velleman's (1988) criticism of Brandt's reforming definition of 'good' (as in Brandt 1979/1998).

¹⁷ Compare Burke (1790/1982).

ideological forces. Such ideological contamination might render our intuitive grasp of the relevant concept unreliable, even if it does not infect the inferential role of the concept itself.¹⁸

Next consider two types of reason supporting orienting normative inquiry around *theoretical* concepts. First, consider “explicated” theoretical concepts that can be explicitly and uncontroversially characterized in relatively precise terms. One reason to adopt such concepts is that explicitness and precision can be highly useful in normative investigation. And it may be extremely difficult (if possible at all) to achieve an explicit precise characterization of a relevant folk concept. Further, explications can be tailored to match what the inquirer cares about most in inquiry (to the extent this is clear to her).

There are at least two reasons to be cautious about this sort of consideration, however. First, at least on some theories of meaning or content, explication or stipulation provide no guarantee of meaning. For example, some theorists will argue that even if ‘atom’ was introduced to pick out the simple indivisible constituents of matter, the term functioned in important ways in physics independent of that definition, allowing it to turn out that atoms are neither simple nor indivisible.¹⁹ Second, precision might be positively *misleading*, if the subject matter being studied itself lacks precise boundaries. Aristotle famously emphasized this point, and took it to apply to certain normative inquiries.²⁰

A second reason to orient normative inquiry around theoretical concepts is that they, as a general *kind* of concept, have a track record of relevant usefulness. Inquiry in the natural and social sciences often involves the introduction of new theoretical concepts. For example, think of the concept QUARK in physics, or the concept IMPLICIT BIAS in the social sciences. Because the natural and social sciences include important paradigms of successful inquiry, this might suggest that orienting around theoretical concepts can contribute to the success or significance of an inquiry. Plausibly, theoretical words and concepts are often introduced because inquirers identify a need for new concepts in order to promote their aims as inquirers. One might hope that the theoretical concepts that emerge in normative inquiry can similarly help us to better achieve the aims of such inquiry.

One reason for caution about this reason is that its force plausibly depends greatly on the degree of similarity between normative inquiry (and what it investigates) and scientific inquiry (and what it investigates). And there is great controversy concerning this very question.²¹ We return to this issue in section 4.

At the start of this chapter, we claimed that, if we focus on a specific agent in a specific context, two central questions in conceptual ethics are: *which concepts* that agent should use, and *which words* she should use to express those concepts. In this section, we have canvassed some reasons for favoring orienting normative inquiry around folk concepts, or around theoretical concepts. This is an instance of the first

¹⁸ For connected discussion, see Railton (2003); Jones (2005); and Eklund (2017: chapter 7).

¹⁹ Compare Schroeter and Schroeter (2014).

²⁰ Aristotle (2002: 1094b13–b25).

²¹ Compare the important (although now dated) review of metaethics in Darwall, Gibbard, and Railton (1997), which marks a central division in metaethics as being between theories that posit *continuity* vs. *discontinuity* between ethics and science.

kind of issue in conceptual ethics. In the next section, we now turn to the second kind of issue, about concept/word pairing.

2. Concept/Word Pairing

There are a wide range of philosophically rich issues in conceptual ethics that bear on which words to use in order to express a given concept. For example, the use of certain words might be advocated because (i) their use in relevant contexts promotes certain broadly epistemic goals (such as the pursuit of knowledge of a given subject matter—for example, in physics or moral philosophy)²² and/or (ii) their use in relevant contexts promotes certain broadly practical goals (such as the political goal of helping foster a more just or free society, or the ethical goal of living a better life or avoiding causing unjust harms to others).²³

It is beyond the scope of this chapter to address the full range of such issues as they apply to the conceptual ethics of normative inquiry. Instead, our goal is to focus on a specific cluster of issues connected to our discussion in the previous section of whether to orient one's inquiry around folk or theoretical concepts. The specific cluster of issues we focus on arises from the fact that certain words that are central to existing normative inquiry (such as 'moral' or 'rational') are associated with several significantly different inferential patterns. As we will show, this entails that normative inquirers face interesting questions about which words to use to express either folk or theoretical concepts.

To begin, we distinguish folk from theoretical *words*. As we shall understand them, *folk words* are words used to express one or more folk concepts. Similarly, *theoretical words* are words used to express one or more theoretical concepts. Given this way of drawing the distinction, some words will be both folk words and theoretical words. This is true, for example, of 'rational' and 'moral'.

To make the idea of associated inferential patterns vivid, we will focus on the word 'moral'. Consider the range of competing philosophical accounts of what distinguishes moral thought and talk from other normative thought and talk. Some associate morality with attitude-independent categorical norms.²⁴ Others associate it with an impersonal point of view,²⁵ or a distinctively interpersonal one.²⁶ Still others associate it with the fittingness conditions for attitudes of certain emotions, such as guilt, resentment, and anger.²⁷ Others associate it with a list of supposed "platitudes" about morality that include both things about the subject matter of morality (e.g., which kinds of activities it regulates) as well as things about its (at least purported) relative normative import.²⁸ And many assume or defend a kind of *rationalism* about moral obligation, according to which, if you morally ought to perform an action you therefore ought to perform it, in the most authoritative sense of 'ought'.²⁹

²² See, for example, Carnap (1950/1962); Brandt (1979/1998); and Railton (1986).

²³ See, for example, Brandt (1979/1998); Railton (1986); and Haslanger (2000).

²⁴ Williams (1985). ²⁵ Sidgwick (1874). ²⁶ Scanlon (1998) and Darwall (2006).

²⁷ Gibbard (1990). ²⁸ See Smith (1994). Cf. Foot (1972). ²⁹ Korsgaard (1996).

Our aim here is not to adjudicate among these proposals. Rather, we want to emphasize that—at least in our broad social/historical context—each of these proposals captures an idea that can seem natural to associate with the word ‘moral’, at least in some circumstances. We can understand these associations as connected to inferential patterns: in at least some circumstances, it will seem natural to infer from the presence of a moral property to the presence of one of these associated properties, or vice-versa.

On some views about how to individuate folk concepts, it might be that *each* of these competing proposals correctly analyzes one among many folk concepts that are sometimes expressed by ‘moral’. Suppose that this were so. This would pose a clear danger that many exchanges among normative inquirers could end up being “merely verbal disputes”, where speakers “talk past” each other in their discussions.³⁰ Merely verbal disputes can stall the progress of inquiry, or hinder it in other ways (e.g., by leading to confusion on the part of participants).³¹ However, our focus here is *not* on these dangers that verbal disputes may pose for normative inquiry. Rather, it is with related (but distinct) issues: issues that we think are under-appreciated in the practice of contemporary normative inquiry (even if many practitioners of normative inquiry will, on reflection, agree these are issues to be reckoned with).

To see these issues, suppose for the moment that the word ‘moral’ picks out a single concept MORAL across the range of uses relevant to normative inquiry (i.e. setting aside obviously different uses like ‘the moral of the story’.) As we suggested above, there are many distinct inferential patterns associated with the word ‘moral’ in our social/historical context. Even if one such inferential pattern is no part of the concept MORAL, it could still play an important role in explaining how we in fact reason using that concept. This raises what we call the *unreliable inference* danger when using such a word in normative inquiry. This is the danger that, when attempting to use MORAL, or consulting one’s intuitions about it, the inquirer could be guided not by this concept, but by an associated idea that is not constitutive of the concept, and which in fact is inaccurate. This danger can arise even for the inquirer who intends to use the word ‘moral’ to express the (allegedly unique) concept MORAL.

There are several potential species of such unreliable inferences. One important species that we will return to below involves what we will call *implicit switching*. This occurs when different associated ideas guide inference in a single sequence of reasoning. For example, one might infer that an action was morally required from its being demanded from a distinctively interpersonal point of view. One might then

³⁰ See Chalmers (2011) and Jenkins (2014) for helpful general discussions of verbal disputes.

³¹ We should underscore that we do *not* think that every time there is variation in word meaning speakers are doomed to end up in a merely verbal dispute. In some cases, the speakers might still express genuine disagreements when they utter claims using a term ‘X’, despite this variation in meaning. For example: the parties might be involved in a *metalinguistic negotiation*. In cases of metalinguistic negotiation, a speaker uses (rather than mentions) a term to advocate for a view about how that very term should be used. Speakers in a metalinguistic negotiation might well express conflicting normative views about how a word should be used—views that will standardly be based on normative considerations about things *other* than words and concepts (e.g., how we should live, how we should organize our social/political institutions, or what objective joints there are in reality)—even if those views are expressed through pragmatic mechanisms (rather than in terms of literal semantic content). See Plunkett and Sundell (2013) for further discussion. See also Thomasson (2016) and Ludlow (2014) for connected discussion.

infer *from* the action's being morally required that one has an authoritative obligation to do it. Unless these associated ideas are both necessary conditions on the concept MORAL, this inference could lead one astray.

Against these dangers, one must weigh a straightforward reason to use familiar, as opposed to obscure or novel words: speech or text composed of such words will generally be *easier to understand*. To see this, imagine a book that begins by introducing a long list of explicitly defined novel terms, and then proceeds to use only the newly defined terms in the text that follows. Such a book will be much harder to understand than a book conveying the same ideas using familiar words. Insofar as folk words are typically more familiar than theoretical words, this issue about comprehensibility suggests a presumptive reason for using folk words.

So far, we have focused on the use of folk words to express folk concepts. However, it is also common for normative inquirers to appropriate existing folk words as vehicles to express their *theoretical* concepts. For example, explications take this form.³² One reason to use existing folk words to express an unfamiliar theoretical concept is that it can help to orient you (and your audience) to roughly what you want to be talking about. Another reason is that it might be an important element of a campaign to get a group of speakers (e.g., all ordinary speakers, or a specific subset of philosophers, or a group of political activists, etc.) to reform their usage to accord with your preferred usage.³³

The use of folk words to express theoretical concepts faces a clear form of the unreliable inference danger. For any folk term, a competent speaker will tend to find natural the inferences that they have come to associate with that term. And this may lead them to make these inferences even if they are not licensed by the theoretical concept they are using.

To make this vivid, consider Ronald Dworkin's use of the term 'morality' in *Justice for Hedgehogs*. Dworkin stipulates that there is a "distinction between ethics, which is the study of how to live well, and morality, which is the study of how we must treat other people".³⁴ This stipulated definition of 'morality' means that we won't have specifically "moral" reasons or obligations (etc.) that stem directly from the welfare of at least many non-human animals, or the status of the natural environment. (And perhaps we will also lack "moral" reasons that arise from the welfare of human infants, depending on how we cash out 'people'.) Because the stipulation is fully compatible with there being weighty 'non-moral' reasons arising from (e.g.) the welfare of infants or non-human animals, it is not clear whether this fact about the consequences of the stipulation is a problem. However, we think this sort of stipulation, and the fact that it excludes considerations about the entities we have highlighted, illustrates important worries in conceptual ethics.

³² See, for example, Carnap (1950/1962); Carnap (1947/1956); Railton (1986); and Brandt (1979/1998).

³³ This sort of campaign might be explicit. See, for example, the discussion of race and gender terms in Haslanger (2000) and the discussion of 'true' in Scharp (2013). Or it might occur via metalinguistic negotiation. For discussion of this latter option, in the context of the use of philosophical terms, see Plunkett (2015).

³⁴ Dworkin (2011: 13). It should be noted that others follow Dworkin in this stipulated usage when framing their discussions. For example, see Appiah (2005: xiii) drawing on Dworkin (2000).

We take Dworkin's stipulation to face an especially worrisome instance of the implicit switching danger. Many ordinary people take 'morality' to pick out something that is *particularly* normatively important. It will be hard for many to shake this association. This threatens to make illegitimate inferences more likely: e.g., inferring, without argument, from a claim about something being "moral" (in the stipulated sense) to a claim about its distinctive normative importance, relative to other kinds of considerations. This means that using the term 'morality' to refer to Dworkinian morality threatens to undercut giving certain normative considerations their due in normative reasoning. This is especially worrisome if the kinds of considerations being ruled out involve entities (e.g., many non-human animals, human infants, or humans that lack certain cognitive capacities) that are already objectionably marginalized in our actual social/political practices, normative inquiry, or both.³⁵

It is worth emphasizing that this is not intended as a point about *any particular philosopher's* psychology. Insofar as normative inquiry is a social endeavor, these issues will be important when they arise for members of the community of inquirers who attempt to engage with or use this explication.

Notice that a related danger can arise even for philosopher who is careful not to stipulate the meaning of 'moral' in this way. For example, consider T. M. Scanlon's *What We Owe to Each Other* (1998). Scanlon offers a contractualist account not of morality as a whole, but "... of a narrower domain of morality having to do with our duties to other people, including such things as requirements to aid them, and prohibitions against harming, killing, coercion, and deception".³⁶ This is roughly the same part of the normative that Dworkin stipulatively uses the term "morality" to refer to. Scanlon goes on to say the following: "It is not clear that this domain has a name. I have been referring to it as "the morality of right and wrong", and I will continue to use this label".³⁷ However, even here the narrower usage of 'morality of right and wrong' is so close to an ordinary way of talking about morality as a whole, that (e.g.) implicit switching concerns may still loom large.

Similar dangers show up in many sorts of normative inquiry. For comparison, consider the case of 'rational' in epistemology. Many epistemologists treat 'rational belief' and 'justified belief' as interchangeable. For example, Stewart Cohen claims that "'reasonable' and 'rational' are virtual synonyms for 'justified'", Michael Huemer claims that "another word for what is justified... is 'rational'", and Declan Smithies claims that "to say that one has justification to believe a proposition is to say that it is rational or reasonable for one to believe it".³⁸ This association with 'rational' may have important consequences.³⁹ Quite generally, criticism of someone as "irrational" suggests that something has gone wrong in them. Applied to belief, this feature of IRRATIONAL makes for extremely natural inferences from irrationality to failure of epistemic responsibility. This in turn means that the assumed association between 'justified' with 'rational' may lend unearned plausibility to responsibilist

³⁵ For connected discussion here about the case of 'justice', in particular with respect to the way it interacts with normative concern for non-human animals, see Plunkett (2016b).

³⁶ Scanlon (1998: 6). ³⁷ Scanlon (1998: 6).

³⁸ See Cohen (1984: 283); Huemer (2001: 22); and Smithies (2012: 274), respectively.

³⁹ For discussion and critique of equating justification and rationality, see Sylvan (MS).

theories of epistemic justification.⁴⁰ Even if philosophers stipulate a meaning for ‘rational’ that does not support such inferences, use of the word may nonetheless influence the inferences or intuitions that inquirers make when using the stipulated notion, rendering them less reliable. Note that these issues closely parallel the ones just discussed above about the Dworkinian notion of morality, even though much of this discussion doesn’t involve any specific *stipulation* of terminology.

Thus far, we have considered the use of folk words in normative inquiry. This has led us to say some things along the way about the virtues and vices of using theoretical words. We now turn to that topic more explicitly.

As the preceding discussion makes clear, one reason for using theoretical words is to *disambiguate*. If there are many ideas associated with ‘moral’, then using a novel theoretical term might help to focus one’s attention on the specific concept one has in mind. This can be true even if the concept you wish to deploy is a folk concept. Consider three cases.

First, suppose that one is convinced that the correct understanding of the folk concept MORAL concerns when guilt and resentment are fitting. One might introduce a novel term to talk about this concept, rather than using the word ‘moral’, because one is worried about triggering associations between the word ‘moral’ and ideas unconnected to the conditions under which these emotions are warranted.

Second, some terms relevant to normative inquiry are associated with *both* folk and theoretical concepts. (As with ‘moral’ and ‘rational’.) In light of this, one might want to explicitly *flag* that one is using a folk concept by using a theoretical word like ‘folk morality’ or ‘folk rationality’.

Third, if you believe there are multiple folk concepts expressed by a given term, you might want to explicitly flag *which* one you are using, via introducing a technical term. For example, suppose that you believe that the term ‘knowledge’ is associated with at least two folk concepts: a factive concept, which philosophers have traditionally privileged, and a non-factive concept.⁴¹ One might want to introduce a technical term, like ‘non-factive knowledge’, to focus attention on which folk concept you intend to pick out. *Mutatis mutandis*, these same considerations can favor the use of theoretical terms to express theoretical concepts.

Next consider potential dangers associated with using theoretical terms. The first thing to note is that many *theoretical* terms (especially those with a significant history of use) have a plurality of ideas associated with them in theoretical contexts. This is true for many prominent theoretical terms in philosophy, like ‘grounding’,

⁴⁰ For one place in contemporary epistemology that relies on this distinction, see Weatherston (2008) and Littlejohn (forthcoming), which both offer replies to the “New Evil Demon Problem” (see Cohen 1984). Note that our calling attention to this conceptual distinction is fully compatible with the possibility that the correct theory of epistemic justification turns out to vindicate a tight connection between epistemic justification and rationality, and hence certain patterns of inference involving them.

⁴¹ This non-factive sense is prominent in the history and sociology of science, where scholars will sometimes talk about the “production of knowledge” concerning claims we now know to be false (but which were thought to be true at the time, or which were relied on in certain ways in scientific practice), as in Shapin (1994), and in psychology, as in Gilovich (2001). See also Michel Foucault’s use of “knowledge” in his discussion of his idea of “power/knowledge”, as in Foucault (1980), and in earlier work such as Foucault (1966/2000). For an apparently non-factive folk use, see Seuss (1965).

‘epistemic’, or ‘metaethics’.⁴² On some semantic views, this divergence might undermine our ability to effectively refer to *anything* with these terms. Consider, for example, certain externalist theories according to which reference is determined in large part by expert usage. If it is indeterminate which (if any) experts one is deferring to in use of one of these terms, one might fail to refer to anything at all.⁴³

Even setting this worry about reference failure aside, using theoretical terms with a plurality of associated ideas invites worries about unreliable inference. For example, consider the term ‘epistemic’. This is arguably a theoretical word in its central uses by contemporary philosophers engaged in normative inquiry.⁴⁴ In some cases, philosophers explicitly claim that properly “epistemic” justification must be the sort of justification that bears an explanatory connection to *truth*.⁴⁵ Others use the term ‘epistemic’ to pick out norms or values that are tied to the *constitutive* standards that govern beliefs, where it is then a further question whether or not those standards are truth-related or not.⁴⁶ Finally, some take the epistemic standards to be particularly normatively important or weighty with respect to all-things-considered normative theorizing about beliefs.⁴⁷

For many expert epistemic inquirers, each such association of the word ‘epistemic’ will have become psychologically entrenched and intuitive. And this means that such experts may become vulnerable to versions of the unreliable inference dangers that we discussed for the use of folk words.⁴⁸ Examples like this show that switching to theoretical terminology does not guarantee that one will thereby avoid vulnerability to implicit switching and unreliable inference worries. However, switching to such theoretical terminology may nonetheless help to mitigate those worries. For example, the associations of a theoretical term with a certain range of ideas will often be less psychologically entrenched (even in an expert), and easier to make explicit than associations with folk terms. And use of a theoretical term at least sometimes generates pressure to be explicit about what one intends to communicate via the term.

As we have sought to make clear, the question of which words to use remains complex, even given a decision concerning whether to use folk or theoretical concepts. Many central terms in normative inquiry—both folk and theoretical—are associated with multiple theoretically significant ideas. We have argued that this raises significant dangers for normative inquiry.

These dangers are exemplary of a cluster of general issues in conceptual ethics concerning concept/word pairing. In the next section, we discuss whether to orient normative inquiry around extension-revealing or authority-revealing concepts.

⁴² For extensive discussion relevant to ‘metaethics’, see McPherson and Plunkett (2017).

⁴³ See Cappelen (2013). ⁴⁴ On this point, see Cohen (2016).

⁴⁵ We endorse this explanatory connection in McPherson and Plunkett (2015). See Berker (2013: section 3) for references to epistemologists endorsing a range of similar theses. For a view that denies this kind of explanatory connection, see Enoch and Schechter (2008).

⁴⁶ See Nolfi (2014). Note that Nolfi herself goes on to deny that belief aims at truth. See also Nolfi (MS).

⁴⁷ One unusually explicit example: Schroeder (2018) presupposes that practical and epistemic reasons are of a kind, and explores potential explanatory priority relations between them. Elsewhere, Schroeder suggests that the normative is “all about reasons” (Schroeder 2007: 81), in (at least) the sense that the core normative facts that really matter can all be reductively explained in terms of normative reasons.

⁴⁸ For connected discussion, see Cohen (2016).

Many of the complications and dangers highlighted here about concept/word pairing also arise in connection with that issue. Having illustrated the general structure of this cluster issues about concept/word pairing here, we will merely touch on a couple of examples of this in the next section.

3. Authority-Revealing vs. Extension-Revealing Concepts

Consider two things that someone engaged in a normative inquiry might want from a normative concept. On the one hand, we care about which actions, states of affairs, etc., our normative concepts apply to. Given this, it would be attractive if there were a clear way of discovering these sorts of facts about a concept. On the other hand, it would be attractive if we were confident that normative conclusions framed in terms of that concept *really mattered*. Ideally of course, one could find a concept that possessed both of these features to a high degree. But it often seems that one in fact needs to weigh these features against each other.⁴⁹

To illustrate: an inquirer interested in norms that govern our practical lives might orient her inquiry using the concept POLITENESS. Or, she might focus on the concept MORAL WRONGNESS. Or, she might introduce a novel concept that she stipulates to be maximally normatively significant for deliberation or action. Competent users arguably typically have a pretty strong grip on what is polite, but how much politeness normatively matters is highly contested. The situation might be almost reversed for the introduced theoretical concept. If forced to choose between orienting her inquiry around one of these normative concepts, which should she choose? In this section, we consider what can be said about this question.

It will be useful to have names for the two attractive features we contrasted above, as well as to say a bit more about what each of these features involve. We will call a concept *extension-revealing* to the extent that ordinary use of that concept tends to make facts about its extension accessible in a way that is usable to the person employing that concept.⁵⁰ Extension-revealingness is gradable along a number of different dimensions. For example, one dimension is how quickly or easily someone can identify what's in the extension. Another dimension is how much of the extension she can identify (e.g., just paradigm cases or much larger swaths of the extension). Moreover, concepts can be extension-revealing for different reasons. In some cases, it might be that users can identify part of the extension of a concept by grasping analytic *a priori* truths about the concept. In other cases, users can identify part of the extension of a concept because of how the relevant community has been

⁴⁹ This might be *necessarily* so, given the nature of what we below call 'authoritatively normative' concepts. (For example: if things in the vicinity of Moorean "open-question" qualities of the sort discussed by Gibbard (2003), drawing on Moore (1903/1993), are true of authoritative normative concepts, but not other concepts.) For our purposes here, we leave this issue aside.

⁵⁰ We use 'extension' here because it is a familiar and useful label. We intend the point we are making to apply to normative terms even if they lack extensions. For example, suppose that some normative terms function as operators. For such cases the relevant question is how much is revealed about, roughly, the truth-conditions of sentences containing the operator.

trained to use the concept over time, even if the relevant knowledge is neither analytic nor *a priori*.

The relevant idea of *normative authority* is less straightforward, so we will introduce it via an example. Suppose that Priya takes herself to face a conflict between what morality and prudence demand of her. She might ask: given this conflict, *what ought I to do?* Notice that, when Priya asks this question, it does not make sense to interpret her as seeking just *some normative standard or other* that can adjudicate the perceived conflict. For example, if etiquette joins with morality against prudence in this case, this hardly answers her question. Rather, Priya is deploying a normative concept here that aspires to wear a distinctive normative *authority* on its sleeve: a concept such that, when answers are framed using it, those answers purport to settle conflicts such as one between morality and prudence.⁵¹ Or at least she is *trying* to deploy such a concept (e.g., perhaps she fails to token a determinate and coherent concept.) A concept is *authority-revealing* to the extent that competence with that concept tends to make its authority accessible in a way that is usable to an actual competent person employing that concept.⁵² For brevity, call the apparently authority-revealing concept that Priya tokens in her deliberation **AUTHORITATIVE OUGHT**.

It is crucial that the properties of concepts that we are focused on concern what inferences it is *relatively* straightforward for competent users of these concepts to make. For example, we want to leave it open that the concept **POLITENESS** is in fact maximally authoritative. We insist only that such authority is not transparent to competent users: it would not be particularly puzzling or surprising for a competent user of **POLITENESS** to claim that this concept simply picks out certain social relations, with little authoritative normative significance. By contrast, it would be puzzling and surprising for a normal competent user of **AUTHORITATIVE OUGHT** to take this concept to be non-defective, and yet to deny that it has authoritative normative significance.

Conversely, it would be puzzling for the ordinary user of **POLITENESS** to simply deny that they had any idea what sorts of behaviors were polite in their community. **POLITENESS** is arguably highly extension-revealing. By contrast, **AUTHORITATIVE OUGHT** is arguably not extension-revealing. This does not imply that it lacks a determinate extension. Rather, it simply signals that relatively few facts about the extension of this concept are transparent to ordinary competent users of this concept, or at least such users in our social/historical context. For example, it would not be

⁵¹ See McPherson (2018). Note that an authoritative resolution of such a conflict could take many forms. For example, it could involve deferring entirely to the dictates of morality. Or it could involve weighing the verdicts of morality and prudence in any number of more complicated ways. For example: generally favoring the dictates of morality, but not in cases when prudence *very strongly* suggests you should Φ , and morality only weakly recommends you not Φ . Or, for example, it could involve completely ignoring the final dictates of morality and prudence and balancing contributory factors (perhaps including factors that these norms take as morally or prudentially significant).

⁵² Notice that if you think there is a hierarchy of degrees of normative authority, then it would be important to distinguish: (i) *what degree of authority* is revealed to the competent user, and (ii) *how manifest* that authority is to the competent user. We will ignore this complication in the text.

puzzling or surprising for a competent user of this concept to deny that she had any idea what is in its extension.

Consider next claims about *genuine* reasons to perform a certain action, or the *genuine value* of certain states of affairs. On at least some ways of understanding these sorts of claims, they purport to deploy authority-revealing *contributory* notions. For example, one might think it is a conceptual truth that what one ought to do is just what one has most reason to do.⁵³ If this line of thinking is correct, it suggests that there might be whole classes of authoritative concepts, cutting across categories like the evaluative, deontic, aretaic, etc.⁵⁴

The contrast between authority-revealing and extension-revealing concepts might be taken to be equivalent with, or at least deeply tied to, the contrast between thin and thick concepts.⁵⁵ In rough terms, thick concepts (e.g. BRAVE, COWARD, KIND, JERK, etc.) involve a mixture of descriptive and normative application-conditions. The kind of mixture they involve is meant to contrast with thin ethical concepts (e.g. OUGHT, BAD, etc.). It is a matter of much debate how to understand the contrast between thin and thick concepts, as well as what exactly each of them *are*. For example, there is much debate about whether the descriptive and normative aspects of thick concepts can be separated from each other, as well as whether the normative aspects of thick concepts are part of their content, or whether they are instead pragmatically associated with their use.⁵⁶ We should not treat the distinction between thick and thin concepts as equivalent to the distinction between authority-revealing and extension-revealing concepts, even if it turns out that certain thin normative concepts are most authority-revealing. The basic reason is this: for both thick and thin concepts, there is the question of how authoritative the normativity associated with that concept is. This might vary across thick concepts: for example, contrast TREACHEROUS with BANAL. The same is arguably true for thin normative concepts: famously, the truth-conditions of 'ought' can vary widely across contexts of use, sometimes picking out intuitively non-authoritative standards.⁵⁷

With this orientation in mind, let's return to the question of whether to orient one's normative inquiry around authority-revealing concepts or extension-revealing concepts. This question presupposes that we cannot have a concept that is maximal on both dimensions. This might be denied. Consider, for example, one particularly optimistic constitutivist project, which involves two goals. The first is to identify a theoretical concept of a constitutive norm, perhaps CONSTITUTIVE NORM FOR ACTION, which is maximally authoritative. The second is to argue that we can derive the complete set of facts about what we ought to do in a relatively straightforward way on the basis of reflection on this concept. The rationales we consider below

⁵³ Douglas Portmore calls this the "least controversial normative principle concerning action" (Portmore 2013: 437). (Whether it is then a *conceptual* truth is a separate matter.) For an argument against understanding 'ought' in terms of 'most reason', see Broome (2015).

⁵⁴ See McPherson (2018) for this proposal.

⁵⁵ For example, parts of Eklund (2017) seem to presuppose this equivalence. See, for example, Eklund (2017: 18–19).

⁵⁶ See Roberts (2017) for an overview of thick concepts.

⁵⁷ On this point, see Kratzer (2012) and Finlay (2014).

assume that such a project is not highly promising.⁵⁸ It is instructive to contrast this sort of ambitious constitutivist project with an attempt to introduce a concept by stipulating both that it is maximally authoritative, and that it has a certain extension. It is plausible that this fails: we cannot stipulate our way to making anything we like maximally authoritative.⁵⁹

We begin by considering a natural rationale for focusing on authority-revealing concepts. Normative inquiry, like any inquiry, can be motivated by the desire to know (or better understand, etc.) facts about our world. However, it is very natural to think that normative inquiry can also be centrally motivated by the desire to discover answers that can (in some sense) *directly* guide our decisions about (e.g.) how to live, what to believe, or which social/political arrangements to support, promote, or protest. It seems plausible that answers framed in terms of authority-revealing concepts will be most suitable to play this role of direct guidance.

Further, many philosophers who investigate normative topics without using authority-revealing concepts arguably do so because they assume a connection between these topics and authoritative norms. For example, in political philosophy, many philosophers who investigate issues about justice, freedom, and/or equality do so in part because they take considerations about these things to bear in important ways on how we authoritatively ought to arrange our social/political institutions.⁶⁰ Similarly, in ethics, many philosophers who investigate morality assume that morality is closely connected to the authoritative norms for action.⁶¹ And within epistemology, it is widely assumed that *knowledge* is intimately tied to the most authoritative norms for the regulation of belief.⁶²

If we assume these commitments, one central reason to orient normative investigation around maximally authority-revealing concepts is that orienting it instead around EXPLOITATION OF MORALITY OR KNOWLEDGE may seem like a puzzling and potentially dangerous detour. At best, one investigates the topic of ultimate interest less directly than one could. But there are more significant dangers, which can be illustrated by considering the example of MORALITY. The fact that many inquirers focus on MORALITY is partly explained by the fact that normative authority is saliently associated with the word ‘moral’. One danger is that the association with

⁵⁸ Notice here that we are not setting aside the possibility of a slightly less optimistic constitutivism, according to which there is a concept whose authoritativeness and extension are both derivable via difficult constitutivist philosophical reasoning.

⁵⁹ For connected discussion, see Prior (1960).

⁶⁰ See, for example, Rawls (1971/1999); Nozick (1974); Anderson (1999); Sen (2009); Satz (2010); Dworkin (2011); and Pettit (2012).

⁶¹ See, for example, Smith (1994) and Korsgaard (1996).

⁶² On this front, consider the vast amounts of effort spent on thinking about the nature of knowledge in the wake of Gettier (1963). Some epistemologists might well be content to think of engaging in this enterprise as just investigating the contours of a concept we happen to employ, which picks out something with little connection to authoritatively normative facts about what we should actually believe (or how we should proportion our credences, etc.). However, surely part of why so many epistemologists have been concerned with the nature of knowledge is that they assume that issues about knowledge (e.g., who has it when) are *normatively important*. In a similar vein, consider the rise of so-called “knowledge-first” epistemology in the vein of Williamson (2000), which uses knowledge as the basis for many things we take to have normative import (e.g., evidence).

normative authority might not be vindicated by the correct theory of the inferential role of MORALITY. Even if the association is vindicated, the plurality of ideas associated with ‘moral’ raises the danger that framing one’s inquiry in terms of this concept will make that inquiry vulnerable to unreliable inference dangers.

What can be said against this rationale for orienting one’s theorizing around authority-revealing concepts? To begin, consider the idea that MORAL is not itself authority-revealing, but that certain moral facts *ground* the facts about what you authoritatively ought to do. If this were true, it might seem to make sense to investigate moral facts as a way of discovering what one ought to do. One thing to be said for this approach is that—assuming that all of its presuppositions are correct—such investigation might contribute to providing *understanding* of authoritative normativity, by allowing us to discover not just what we ought to do, but *in virtue of what* we ought to do it.⁶³

Another rationale for investigating the grounds of normative facts is that one might think this is an especially promising way to discover what those facts are. One difficulty with this rationale is that priority in metaphysical explanation is not reliably correlated with epistemic access. To take a simple example, the fundamental micro-physical facts presumably ground facts about the observable features of our environment. But we evidently have more direct access to the latter than we do to the former. So this rationale would require an argument that shows why metaphysical priority is a good guide to epistemic access in the case of the normative. For example, it might be that, at least in certain epistemic contexts, we have no good epistemic access to the relevant facts that doesn’t proceed via learning about the facts that determine them.⁶⁴ But it is far from clear that this is so in the case of authoritatively normative facts, in the sorts of epistemic contexts that we standardly find ourselves in.

Consider a different epistemic argument for focusing on certain less authority-revealing concepts (e.g., perhaps MORALITY or EXPLOITATION), rather than directly on what one authoritatively ought to do. Recall from section 1 the suggestion that we all have a good deal of (perhaps implicit) trained moral knowledge, as a result of our moral education. One might think that most of us lack a parallel training that is framed in terms of the concept AUTHORITATIVE OUGHT, or associated contributory notions, such as GENUINE REASON FOR ACTION. This might be because AUTHORITATIVE OUGHT (e.g.) is a theoretical, and not a folk concept. But it need not be: the same point would hold even if it were simply a *relatively obscure* folk concept. If this is true, then, when one is considering a possible action, and asking oneself the question “authoritatively ought I to do this?” one might tend to come up blank. Or, one might tend to generate an answer that is based in one’s implicit understanding of some *other* concept, such as WHAT SERVES MY PRESENT AIMS, or WHAT IS GOOD FOR ME, or MORAL RIGHTNESS, without interrogating the implicit inference to what one ought to do.

⁶³ Note that one might view understanding as *one* goal of normative inquiry, or as *the distinctive* goal of normative inquiry. A case for the latter view might draw on parallel arguments that understanding is a constitutive goal of moral inquiry in Hills (2009).

⁶⁴ See Greenberg (2006) for discussion of this idea within legal philosophy, in the context of discussing facts about the content of the law (in a given jurisdiction, at a given time).

Suppose next that normative authority was part of the concept MORALITY, such that, at least ordinarily, if one morally ought to perform an action, one also authoritatively ought to perform it. If this were true, it might make sense to try to answer at least some questions about what we authoritatively ought to do by engaging in moral inquiry, and inferring conclusions about what we authoritatively ought to do from the moral conclusions found.

We can imagine parallel rationales for the other cases. For example, one might think that a similar argument could be made for orienting one's inquiry in social and political philosophy around JUSTICE, even if one's ultimate aim was to make claims about how we authoritatively ought, all things considered, to arrange our social and political institutions. Similarly, one might advocate for orienting epistemological inquiry around KNOWLEDGE, even if one's ultimate aim was to make claims about how we authoritatively ought, all things considered, to regulate our beliefs.

The force of this rationale interacts with the dialectic concerning folk vs. theoretical concepts discussed earlier, as well as with the issues we discussed about concept/word pairing. For example, this rationale may require inquiry that deploys *folk* moral concepts, as opposed to explicitly theoretically refined moral concepts (where it is much less clear that we would have the relevant sort of trained intuitive knowledge). As such, this rationale is also vulnerable to the various worries about using folk words and concepts discussed in sections 1 and 2 (including ideological contamination, implicit switching, etc.).

Note next that the idea that *any* concepts are authority-revealing might be challenged. Consider the way we tried to describe Priya's concept. It might be denied that anything could satisfy such a concept.⁶⁵ Or it might be denied that examples like Priya's deliberation even succeed in isolating a coherent concept.⁶⁶ Perhaps, as we suggested with 'morality', above, there are several features that philosophers associate with 'authoritative normativity', but no single property or concept could possess all of those features.⁶⁷ If this were the case, there might be a range of authoritative-ish concepts for the normative inquirer to choose among in her explication. And such a choice, once clarified, might reduce the force of the natural rationale for focusing on the authoritative concepts in one's normative inquiry.

Other cases raise analogous complications. For example, it is not clear whether there is a single way in which norms are authoritative *for belief*. One way to bring this out is to consider the variety of interesting properties that various philosophers have associated with *epistemic justification*.⁶⁸ One possible explanation for this diversity is that—given the importance that many place on the topic of "epistemic justification"—many of these properties are ones that (at least some) philosophers associate with authoritative norms for belief. Another way to bring out the issue is to consider possible competing proposals concerning what it is for a norm to be authoritative for belief. For example, a natural proposal is that being authoritative

⁶⁵ See Copp (1997) and Tiffany (2007). See McPherson (2018: section 6) for discussion.

⁶⁶ See Baker (2018).

⁶⁷ See McPherson and Plunkett (2017: section 2.3) for brief discussion of this possibility. For connected discussion, see Finlay (2019).

⁶⁸ See Alston (2005) for a survey of some of the key views here.

for belief is a matter of being related to truth in the right way.⁶⁹ Or perhaps it is a matter of being related to authoritative practical norms in the right way.⁷⁰ Or perhaps it is a matter of being the constitutive norm for the mental state of belief.⁷¹ One possibility is that each of these proposals captures something that we care about in believing, but there is no single property *being authoritative for belief*.

Let's take stock of where we are. This section has sought to illuminate the interest and complexity of the question of whether to orient normative inquiry around authority-revealing or extension-revealing concepts. As we have indicated in certain places, the plausibility of answers to this question interact with the plausibility of answers to certain questions about folk vs. theoretical concepts. This underscores some of the interest in investigating these two different topics in the conceptual ethics of normative inquiry together.

Importantly, conceptual ethics is itself a kind of normative inquiry. This means that the issues in conceptual ethics we have discussed thus far also interact with how we understand and approach the very normative questions we have been posing in this chapter. We now turn to this topic.

4. Evaluating the Norms and Concepts used in the Methodology of Normative Inquiry

So far in this chapter we have explored the choices between focusing normative inquiry around folk or theoretical concepts (and words), and between focusing on authority-revealing and extension-revealing normative concepts. In doing so, we have introduced a series of normative considerations and arguments that bear on these choices. In seeking to assess these considerations and arguments, we can ask: *what sorts of norms* should we be deploying in our methodological evaluation of normative inquiry? We can further ask: for any such answer, what would *explain* the aptness of that answer? In this section, we briefly consider three broad approaches to answering this latter question: (i) considering the significance of inquirer aims, (ii) appealing to the results of metanormative inquiry, and (iii) engaging in further conceptual ethics, of the kind we have been discussing in this chapter. We will also explore some of the ways that these approaches interact.

We begin by considering a natural proposal: that the normative facts about which normative concepts and words an agent should use in a given context are determined

⁶⁹ We defend a related thesis about *epistemic justification* in McPherson and Plunkett (2015). As we just discussed, epistemic justification might well be taken by some to be tightly associated with authoritative normativity. On this front, consider our earlier discussion of those who take justification and rationality to be virtually equivalent. In turn, many such philosophers who accept that (at least near) equivalency also take considerations of rationality to bear in significant ways on what we (authoritatively) ought to believe. (See, for example, Wedgwood 2012, in connection to Wedgwood 2007).

⁷⁰ This might be thought of as a kind of "pragmatist" explanation. See McPherson and Plunkett (2015: 111) for discussion of a related thesis about the epistemic. For connected discussion, see Enoch and Schechter (2008).

⁷¹ For example, Shah and Velleman (2005) defend a constitutivist account of epistemic normativity. For connected discussion, see Nolfi (2014).

by facts about what *promotes the aims* that an agent has.⁷² One way to motivate this *instrumentalist* idea is to notice that it seems to explain certain obvious data points. For example: if an agent aims to do work in sociology, then the concepts she should deploy are arguably different than those she should deploy if seeking to do work in physics. Or, to take another example: the concepts she should deploy when engaged in public political advocacy are arguably different than those she should deploy if seeking to do advanced theoretical work in political philosophy. In both of these cases, it seems that what matters is what promotes the aims the agent has in each case. One elegant hypothesis is that, if the aims of inquirers explain normative facts about concept choice in *these* cases, perhaps they also do so in *all* cases.⁷³ Or more modestly, perhaps the aims of inquirers play this role in the relevant range of cases about conceptual ethics that we are considering in this chapter. This would suggest a clear way in which (at least in principle) one could assess the choices we have highlighted thus far in this chapter: get clear about one's aims are, and investigate how best to serve those aims. (Note that, given that mature normative inquiry is often a collective endeavor, it may be more appropriate to speak of *our* or *their* aims here, rather than the aims of a given individual agent.)

It is not difficult to see that this sort of instrumentalist picture requires substantial defense. Instrumentalism is intensely controversial as a foundational theory of most significant types of norms and values (e.g., epistemic norms, reasons for action, well-being, etc.). Nor is it obvious that it has special credibility for the conceptual ethics of normative inquiry in particular. One way of making this vivid is to consider inquirers with substantively bizarre or awful aims; we may resist the idea that achieving those aims constitutes their doing normative inquiry *well*.⁷⁴

It is possible to draw a different lesson from the example just used to motivate instrumentalism: perhaps our aims function to determine *what* we are investigating (e.g., *sociology* or *normative ethics*). Once we have settled on a topic, it is something about the topic itself that determines how we ought to inquire into it.

Consider an example to see how this idea might apply to a *normative* inquiry. Suppose that our aims make it the case that we are engaged in *moral* inquiry, and that the word 'moral' is univocal enough for this aim to fix a topic. We might then ask the following: what is the nature of moral thought and talk? And what, if anything, is that thought and talk distinctively about? These questions form the heart of what we understand to be *metamoral inquiry*. This inquiry, as we understand it, aims to explain how actual moral thought and talk—and what (if anything) such thought and talk is distinctively about (e.g., moral facts, properties, etc.)—fit into reality.⁷⁵ On a prominent class of metamoral theories, moral concepts are fundamentally *practical* in nature, such that they must be sharply distinguished from standard, descriptive concepts.⁷⁶ Recall

⁷² See Haslanger (2000) and Anderson (2001) for two places where this sort of idea is advanced. See Burgess and Plunkett (2013b) for both sympathetic and critical discussion of this idea.

⁷³ The motivation we have suggested here for instrumentalism about the norms of inquiry is adapted from the motivation that frames Schroeder's defense of the Humean theory of reasons in Schroeder (2007).

⁷⁴ For further discussion, see Burgess and Plunkett (2013b: 1105).

⁷⁵ See McPherson and Plunkett (2017).

⁷⁶ See, for example, Korsgaard (1996), Blackburn (1998), and Gibbard (2003) for versions of this thought.

that one rationale proposed in section 1 for orienting inquiry around *theoretical* concepts appealed to an analogy to scientific inquiry. If moral concepts are of a radically different kind than scientific concepts, this might undermine that rationale.

By contrast, now consider different views about the nature of the reality that moral thought and talk is distinctively about. (Call this part of reality “moral reality.”) On certain *naturalistic* views of moral reality, it is metaphysically continuous with—or a part of—the reality studied in the natural and social sciences.⁷⁷ This may support the idea that the study of moral reality *should* be modeled on the methods of the sciences, arguably strengthening the case for using theoretical concepts. By contrast, consider a non-naturalistic realist metamoral view, according to which moral reality is metaphysically discontinuous with the reality studied in the natural and social sciences.⁷⁸ Perhaps *that* reality can only be accessed using specific kinds of *a priori* reasoning that, at its core, must involve the deployment of folk concepts all of us have prior to any particular theoretical training.

The possibilities sketched here for moral inquiry extend to other kinds of normative inquiry. In general, how we should proceed in normative inquiry may in significant part be determined by our progress in metanormative inquiry.⁷⁹ We understand *metanormative inquiry* as aimed at explaining how actual normative thought and talk—and what (if anything) that talk is distinctively about (e.g., normative facts, properties, etc.)—fit into reality.⁸⁰ As these brief examples suggest, the connections between overall metanormative views and the topics in conceptual ethics we have been discussing in this chapter are rich and worth exploring further. However, it is also worth flagging some of the limitations of this approach for informing the methodology of normative inquiry.

First, one might reasonably be more confident in certain relevant methodological claims than one is in any metanormative theory determinate enough to have the sorts of methodological implications sketched. This might make it unreasonable to change the former in light of the latter.⁸¹

Second, the significance of metanormative inquiry for the sorts of normative questions we have been asking in the first three sections of the chapter is arguably limited. For example, if one has decided to focus on *morality*, some of the metamoral theses just canvassed may help to determine whether to deploy folk vs. theoretical concepts. But it is harder to see how these sorts of theses could tell you whether to orient your inquiry around MORALITY as opposed to EXPLOITATION or AUTHORITATIVE OUGHT. These theses might well help us to better understand (a) the sort of thing(s) that we investigate when we investigate morality and (b) how to investigate morality were we to choose to do so. But understanding those things can't by itself settle our decisions

⁷⁷ For example, see Railton (1986) and Boyd (1997). Note that moral thought and talk might have some features that made it quite distinctive, despite the metaphysics and epistemology of morality being quite continuous with that of some sciences. See, for example, Copp (2001).

⁷⁸ For example, see Shafer-Landau (2003); Fitzpatrick (2008); and Enoch (2011).

⁷⁹ See McPherson (2012) for defense of this idea.

⁸⁰ For further discussion of this characterization of metanormative inquiry, see McPherson and Plunkett (2017) and Plunkett and Shapiro (2017).

⁸¹ See McPherson (2012) for further discussion of this idea.

about *what* to investigate, even if it might be helpful in understanding what possible things we might investigate.

One way to put the basic point here is as follows. Metamoral inquiry is a hermeneutic enterprise, aimed at explaining our *actual* moral thought and talk (and what, if anything, it is distinctively about). It can't, by itself, tell us that we should be investigating moral thought, talk, and reality, or, more generally, make our decisions for us about what to investigate.⁸² Given that we could imagine *that* decision process in instrumental terms, it is not clear how deep the contrast with the instrumental approach really goes here.

To push this line of thought further, consider the choice about which concepts to orient normative inquiry around. We can directly ask a normative question here: *how ought we to make this choice?* One possible answer is: *look to your aims*. But this is only one possible answer. The answer might instead be grounded in facts *independent* of the interests of the inquirer. For example, perhaps some normative truths are intrinsically more valuable to know than others (or some topics intrinsically better to investigate).⁸³ Or perhaps some kinds of knowledge are more *morally or politically valuable* than others, and thus should be privileged in normative inquiry. For example, perhaps we should orient normative inquiry around certain concepts that help *reveal the basis for an unforced political consensus* among people with widely divergent ethical and political views. Such a consensus might allow such people to live relatively harmoniously together (rather than killing each other, as in the history of religious warfare), and perhaps also allow for a certain kind of freedom.⁸⁴ Or perhaps we should aim to uncover theories and practices of reasoning that are "plausible to and usable by moral agents in the case at hand, nonabusive of social power or vulnerability, and capable of delivering feasible conclusions".⁸⁵ Or perhaps we should aim to uncover theories that help give "agents a kind of knowledge inherently productive of enlightenment and emancipation".⁸⁶

Of course, we can now apply the dialectic of this chapter to the norms that apply to the question *how ought we to choose the normative concepts to orient normative inquiry around?* For what sort of OUGHT concept should be deployed in this question? If there is a most authoritative concept that applies to normative inquiry, then perhaps it makes sense to structure inquiry around that (modulo the sorts of competing concerns discussed in section 3). On the other hand, suppose there is *not* a most authoritative norm that applies here. Then it appears that, if one is guided by a norm in selecting among one's options, this norm will have no distinctive

⁸² We here echo what we say at the end of McPherson and Plunkett (2017). For connected discussion, see Eklund (2017).

⁸³ For a systematic articulation of this idea for inquiry in general, and not just in the case of normative inquiry in particular, see Sider (2012).

⁸⁴ This kind of goal underwrites key parts of the liberal tradition in political theory, as well some of the ideas behind the ideal of "public reason". For relevant discussion, see Rawls (1996); Gaus (2011); and Quong (2013).

⁸⁵ Jaggar and Tobin (2013: 413). They articulate this in the context of developing an overall account of moral epistemology, which they take to be rooted in core political ideals of feminism.

⁸⁶ Geuss (1981: 2). Geuss takes this to be the central goal of *critical theory* as such, in the vein of the so-called "Frankfurt School" of critical theory. For connected discussion, see Horkheimer (1937/1975) and Habermas (1968/1987).

authority relative to competing norms you might have used for the purposes of selection. On this hypothesis, it would thus seem that arbitrary choice must play some role in the foundations of the conceptual ethics of normative inquiry.⁸⁷

5. Conclusions

Our central aim in this chapter has been to explore two central questions in the conceptual ethics of normative inquiry. First: should normative inquiry be oriented around *folk* or *theoretical* normative concepts? Second: should it be oriented around concepts that are *authority-revealing*, or ones that are *extension-revealing*? Along the way, we have also addressed questions about concept/word pairing. As we have tried to show, there is an important and interesting class of considerations that can be used to support different answers to these questions in conceptual ethics, as well as a complicated further set of questions about what sorts of normative standards we should be using to answer these questions in conceptual ethics. We take this to warrant attention: these questions are philosophically rich, relatively unexplored, and potentially have striking implications for how we should conduct normative inquiry.

Does this mean that normative inquiry needs to halt while we turn our attention to these methodological questions? Not at all. Some methodological inquiry in philosophy appears motivated by the idea that we cannot legitimately proceed in our other inquiries without first rebutting skeptical challenges, or otherwise using methodology to put our inquiry on a more secure footing. Tied to this, the boldest methodologists have Cartesian aspirations of finding indubitable methods and starting points to replace our shabby-seeming ordinary attempts to answer philosophical questions. Others dream of a method that—even if not indubitable—can usefully be followed by all, no matter how benighted their initial opinions. We take such motivations to be misguided. It might well be impossible for us to achieve either of these goals, at least in the near future. It is also far from clear how much we *need* such foundations to have reasonable confidence in the products of our philosophical theorizing. Finally, methodology is itself philosophy, so it is hard to see how it could provide a decisive antidote to worries about the legitimacy of philosophy.

Our approach is rooted in a different understanding of the value of methodology. Being *more philosophy*, methodological inquiry can be interesting and worthwhile for the same reasons other areas of philosophy can be. For example, it might be interesting because we seek understanding of ourselves, the world we live in, and our thought and talk about that world. In this spirit, we find it natural for (at least some) philosophers engaged in normative inquiry to be curious about the activity they are engaged in, and the arguments that can be made for pursuing that activity in one way or another. With this in mind, we hope that this chapter spurs more interest in the conceptual ethics of normative inquiry.

Further, even without Cartesian aspirations, we can hope that methodological reflection will help us to improve our inquiry. We think that this chapter can

⁸⁷ Wrestling with this hypothesis is one of the central themes of Eklund (2017). See also Burgess (Chapter 6, this volume) for connected discussion of the charge of ‘hypocrisy’ in conceptual ethics.

potentially help normative inquirers to do better in at least two ways. First, we take it that normative inquirers not only *could* make a wide variety of choices concerning which concepts to orient their normative inquiry around, but that they are in fact *actually making* different choices here. This diversity of practice is not always clearly signaled (including in some of our own previous work). We suspect that this is partly because the range of relevant options is rarely salient to those inquirers, with the result that there is no felt need to clarify one's target. We think the lack of clear signaling is also partly explained by the lack of a clear vocabulary for communicating the relevant orientation. The discussion in this chapter can help us to understand each other's work better, by providing (what we hope is) a useful framework within which to locate distinct projects. Second, we hope that this chapter paves the way for further methodological reflection that helps to adjudicate some of the central questions that we have asked in this chapter. If so, this may enable some normative inquirers to focus their efforts more successfully on the questions it is most worthwhile to investigate, whatever those turn out to be.

Acknowledgements

Earlier versions of this chapter were presented at University of Oslo (ConceptLab), University of Birmingham, University of Sydney, the NYC Philosophy of Language Workshop, University of Pennsylvania, University of Colorado Boulder, The California Philosophy Workshop, University of Leeds, Hebrew University of Jerusalem, University of Oxford, University of Edinburgh, Central European University, Stanford University, University of Virginia, the Logos Research Group at the University of Barcelona, and the Arché research center at University of St. Andrews. Thanks to everyone who participated in those sessions. Thanks also to Mark Budolfson, Alexis Burgess, Sarah Buss, Herman Cappelen, Matthew Chrisman, David Enoch, Max Etchemendy, David Faraci, Chris Heathwood, Nadeem Hussain, Andrej Kristan, Frank Jackson, Barry Maguire, Eliot Michaelson, Daniel Nolan, Bob Pasnau, Peter Railton, Raul Saucedo, Karl Schafer, Kevin Scharp, Sharon Street, Tim Sundell, Amie Thomsson, and Daniel Wodak.

References

- Anderson, Elizabeth. 1999. What Is the Point of Equality? *Ethics* 109 (2):287–337.
- Anderson, Elizabeth. 2001. Unstrapping the Straitjacket of 'Preference': A Comment on Amartya Sen's Contributions to Philosophy and Economics. *Economics and Philosophy* 17 (1):21–38.
- Appiah, Anthony. 2005. *The Ethics of Identity*. Princeton, NJ: Princeton University Press.
- Aristotle. 2002. *Nicomachean Ethics*. Oxford: Oxford University Press.
- Baker, Derek. 2018. Skepticism about Ought Simpliciter. In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics*, vol 13. Oxford: Oxford University Press.
- Berker, Selim. 2013. Epistemic Teleology and the Separateness of Propositions. *Philosophical Review* 122 (3):337–93.
- Blackburn, Simon. 1998. *Ruling Passions*. Oxford: Clarendon Press.
- Boyd, Richard. 1997. How to be a Moral Realist. In S. Darwall, A. Gibbard, and P. Railton (eds.), *Moral Discourse and Practice*. New York: Oxford University Press.

- Braddon-Mitchell, David. Chapter 4, this volume. Reactive Concepts: Engineering the Concept CONCEPT. In A. Burgess, H. Cappelen and D. Plunkett (eds.), *Conceptual Engineering and Conceptual Ethics*. Oxford: Oxford University Press.
- Brandt, Richard B. 1979/1998. *A Theory of the Good and the Right*. New York City: Prometheus Books.
- Broome, John. 2015. Reason versus Ought. *Philosophical Issues* 25 (1):80–97.
- Burge, Tyler. 1979. Individualism and the Mental. *Midwest Studies in Philosophy* 4 (1):73–121.
- Burgess, Alexis. Forthcoming. Never Say ‘Never Say “Never”’? In A. Burgess, H. Cappelen and D. Plunkett (eds.), *Conceptual Engineering and Conceptual Ethics*. Oxford: Oxford University Press.
- Burgess, Alexis, and Plunkett, David. 2013a. Conceptual Ethics I. *Philosophy Compass* 8 (12):1091–101.
- Burgess, Alexis, and Plunkett, David. 2013b. Conceptual Ethics II. *Philosophy Compass* 8 (12):1102–10.
- Burke, Edmund. 1790/1982. *Reflections on the Revolution in France*. New York City: Penguin.
- Cappelen, Herman. 2013. Nonsense and Illusions of Thought. *Philosophical Perspectives* 27 (1):22–50.
- Cappelen, Herman. 2018. *Fixing Language: An Essay on Conceptual Engineering*. Oxford: Oxford University Press.
- Carnap, Rudolf. 1947/1956. *Meaning and Necessity: A Study in Semantics and Modal Logic*. Chicago: University of Chicago Press.
- Carnap, Rudolf. 1950/1962. *Logical Foundations of Probability*. Chicago: University of Chicago Press.
- Chalmers, David J. 2011. Verbal Disputes. *Philosophical Review* 120 (4):515–66.
- Cohen, Stewart. 1984. Justification and Truth. *Philosophical Studies* 46 (3):279–95.
- Cohen, Stewart. 2016. Theorizing about the Epistemic. *Inquiry* 59 (7–8):839–57.
- Copp, David. 1997. The Ring of Gyges: Overridingness and the Unity of Reason. *Social Philosophy and Policy* 14 (1):86–106.
- Copp, David. 2001. Realist-Expressivism: A Neglected Option for Moral Realism. *Social Philosophy and Policy* 18 (02):1–43.
- Darwall, Stephen L. 2006. *The Second-person Standpoint: Morality, Respect, and Accountability*. Cambridge, MA: Harvard University Press.
- Darwall, Stephen L., Gibbard, Allan, and Railton, Peter Albert. 1997. Toward Fin de Siecle Ethics: Some Trends. *Moral Discourse and Practice: Some Philosophical Approaches*. New York: Oxford University Press.
- Dworkin, Ronald. 2000. *Sovereign Virtue: The Theory and Practice of Equality*. Cambridge, MA: Harvard University Press.
- Dworkin, Ronald. 2011. *Justice for Hedgehogs*. Cambridge, MA: Harvard University Press.
- Eklund, Matti. 2017. *Choosing Normative Concepts*. Oxford: Oxford University Press.
- Enoch, David. 2011. *Taking Morality Seriously: A Defense of Robust Realism*. Oxford: Oxford University Press.
- Enoch, David, and Schechter, Joshua. 2008. How Are Basic Belief-forming Methods Justified? *Philosophy and Phenomenological Research* 76 (3):547–79.
- Finlay, Stephen. 2014. *Confusion of Tongues: A Theory of Normative Language*. Oxford: Oxford University Press.
- Finlay, Stephen. 2019. Defining Normativity. In D. Plunkett, S. Shapiro and K. Toh (eds.), *Dimensions of Normativity: New Essays on Metaethics and Jurisprudence*. Oxford: Oxford University Press.
- Fitzpatrick, William. 2008. Robust Ethical Realism, Non-naturalism, and Normativity. *Oxford Studies in Metaethics* 3:159–205.

- Foot, Philippa. 1972. Morality as a System of Hypothetical Imperatives. *Philosophical Review* 81 (3):305–16.
- Foucault, Michel. 1966/2000. *The Order of Things: An Archeology of the Human Sciences*. London: Routledge.
- Foucault, Michel. 1980. *Power/Knowledge: Selected Interviews and Other Writings 1972–1977*. New York: Pantheon.
- Gaus, Gerald. 2011. *The Order of Public Reason: A Theory of Freedom and Morality in a Diverse and Bounded World*: Cambridge University Press.
- Gettier, Edmund L. 1963. Is Justified True Belief Knowledge? *Analysis* 23:121–3.
- Geuss, Raymond. 1981. *The Idea of a Critical Theory: Habermas and the Frankfurt School*. Cambridge: Cambridge University Press.
- Gibbard, Allan. 1990. *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Cambridge, MA: Harvard University Press.
- Gibbard, Allan. 2003. *Thinking How to Live*. Cambridge, MA: Harvard University Press.
- Gilovich, Thomas. 2001. *How We Know What Isn't So: The Fallibility of Human Reason in Everyday Life*. New York: The Free Press.
- Greenberg, Mark. 2006. How Facts Make Law. In S. Herskovitz (ed.), *Exploring Law's Empire: The Jurisprudence of Ronald Dworkin*. New York: Oxford University Press.
- Greenough, Patrick. Chapter 11, this volume. Neutralism and Conceptual Engineering. In A. Burgess, H. Cappelen and D. Plunkett (eds.), *Conceptual Engineering and Conceptual Ethics*. Oxford: Oxford University Press.
- Habermas, Jürgen. 1968/1987. *Knowledge and Human Interests*. Trans. J. J. Shapiro. Cambridge: Polity.
- Haslanger, Sally. 2000. Gender and Race: (What) Are They? (What) Do We Want Them to Be? *Nous* 34 (1):31–55.
- Haslanger, Sally. Chapter 12, this volume. Going On, Not in the Same Way. In A. Burgess, H. Cappelen and D. Plunkett (eds.), *Conceptual Engineering and Conceptual Ethics*. Oxford: Oxford University Press.
- Hills, Alison. 2009. Moral Testimony and Moral Epistemology. *Ethics* 120 (1):94–127.
- Horkheimer, Max. 1937/1975. Traditional and Critical Theory. In *Critical Theory: Selected Essays*. New York City: Continuum
- Huemer, Michael. 2001. *Skepticism and the Veil of Perception*. Lanham: Rowman & Littlefield.
- Jackson, Frank. 1998. *From Metaphysics to Ethics: A Defence of Conceptual Analysis*. Oxford: Clarendon.
- Jaggar, Alison M., and Tobin, Theresa W. 2013. Situating Moral Justification: Rethinking the Mission of Moral Epistemology. *Metaphilosophy* 44 (4):383–408.
- Jenkins, C. S. I. 2014. Merely Verbal Disputes. *Erkenntnis* 79 (1):11–30.
- Jenkins, C. S. 2006. The Rules of Flirtation. *The Philosophers Magazine* 36:37–40.
- Jones, Karen. 2005. Moral Epistemology. In F. Jackson and M. Smith (eds.), *The Oxford Handbook of Contemporary Philosophy*. Oxford: Oxford University Press.
- 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.
- Littlejohn, Clayton. forthcoming. A Plea for Epistemic Excuses. In J. Dutant and F. Dorsch (eds.), *The New Evil Demon Problem*. Oxford: Oxford University Press.
- Ludlow, Peter. 2014. *Living Words: Meaning Undetermination and the Dynamic Lexicon*. Oxford: Oxford University Press.
- McPherson, Tristram. 2012. Unifying Moral Methodology. *Pacific Philosophical Quarterly* 93 (4):523–49.

- McPherson, Tristram. 2018. Authoritatively Normative Concepts. In R. Shafer-Landau (eds.), *Oxford Studies in Metaethics*, vol. 13. Oxford: Oxford University Press.
- McPherson, Tristram, and Plunkett, David. 2015. Deliberative Indispensability and Epistemic Justification. In R. Shafer-Landau (ed.), *Oxford Studies in Metaethics*, vol. 10. Oxford: Oxford University Press.
- McPherson, Tristram, and Plunkett, David. 2017. The Nature and Explanatory Ambitions of Metaethics. In T. McPherson and D. Plunkett (eds.), *The Routledge Handbook of Metaethics*. New York: Routledge.
- Moore, G. E. 1903/1993. *Principia Ethica*. Cambridge: Cambridge University Press.
- Nolfi, Kate. 2014. Why is Epistemic Evaluation Prescriptive? *Inquiry* 57 (1): 97–121.
- Nolfi, Kate. MS. Epistemically Flawless False Beliefs.
- Nozick, Robert. 1974. *Anarchy, State, and Utopia*. New York: Basic Books.
- Pettit, Philip. 2012. *On the People's Terms: A Republican Theory and Model of Democracy*. Cambridge: Cambridge University Press.
- Plunkett, David. 2015. Which Concepts Should We Use? Metalinguistic Negotiations and the Methodology of Philosophy. *Inquiry* 58 (7–8):828–74.
- Plunkett, David. 2016. Conceptual History, Conceptual Ethics, and the Aims of Inquiry: A Framework for Thinking about the Relevance of the History/Genealogy of Concepts to Normative Inquiry. *Ergo* 3 (2):27–64.
- Plunkett, David, and Shapiro, Scott. 2017. Law, Morality, and Everything Else: General Jurisprudence as a Branch of Metanormative Inquiry. *Ethics* 128 (1):37–68.
- Plunkett, David, and Sundell, Timothy. 2013. Disagreement and the Semantics of Normative and Evaluative Terms. *Philosophers' Imprint* 13 (23):1–37.
- Portmore, Douglas W. 2013. Perform Your Best Option. *Journal of Philosophy* 110 (8):436–59.
- Prior, A. N. 1960. The Runabout Inference-Ticket. *Analysis* 21 (2):38.
- Quong, Jonathan. 2013. Public Reason. In E. N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy* (summer edn). <https://plato.stanford.edu/archives/sum2013/entries/public-reason/>
- Railton, Peter. 1986. Moral Realism. *The Philosophical Review* 95:163–207.
- Railton, Peter. 2003. Morality, Ideology, and Reflection, or the Duck Sits Yet. *Facts, Values, and Norms: Essays Toward a Morality of Consequence*. New York: Cambridge University Press.
- Rawls, John. 1971/1999. *A Theory of Justice* (revised edn). Cambridge, MA: Harvard University Press.
- Rawls, John. 1996. *Political Liberalism* (expanded edn). New York: Columbia University Press.
- Roberts, Debbie. 2017. Thick Concepts. In T. McPherson and D. Plunkett (eds.), *The Routledge Handbook of Metaethics*. New York: Routledge.
- Satz, Debra. 2010. *Why Some Things Should Not Be for Sale: The Moral Limits of Markets*. New York: Oxford University Press.
- Scanlon, T. M. 1998. *What We Owe To Each Other*. Cambridge, MA: Harvard University Press.
- Scharp, Kevin. 2013. *Replacing Truth*. Oxford: Oxford University Press.
- Schroeder, Mark. 2007. *Slaves of the Passions*. Oxford: Oxford University Press.
- Schroeder, Mark. 2018. The Unity of Reasons. In D. Star (ed.), *The Oxford Handbook to Reasons and Normativity*. Oxford: Oxford University Press.
- Schroeter, Laura. 2008. Why Be an Anti-Individualist? *Philosophy and Phenomenological Research* 77 (1):105–41.
- Schroeter, Laura, and Schroeter, Francois. 2014. Normative Concepts: A Connectedness Model. *Philosophers' Imprint* 14 (25):1–26.
- Sen, Amartya. 2009. *The Idea of Justice*. Cambridge, MA: Belknap Press of Harvard University Press.
- Seuss, Dr. 1965. *I Had Trouble in Getting to Solla Sollew*. New York: Random House.
- Shafer-Landau, Russ. 2003. *Moral Realism: A Defence*. Oxford University Press.

- Shah, Nishi, and Velleman, David J. 2005. Doxastic Deliberation *Philosophical Review* 114 (4):497–534.
- Shapin, Steven. 1994. *A Social History of Truth: Civility and Science in Seventeenth-Century England*. Chicago: University of Chicago Press.
- Sider, Theodore. 2012. *Writing the Book of the World*. Oxford: Oxford University Press.
- Sidgwick, Henry. 1874. *The Methods of Ethics*.
- Sinnott-Armstrong, Walter. 2006. *Moral Scepticisms*. Oxford: Oxford University Press.
- Smith, Michael. 1994. *The Moral Problem*. Cambridge: Blackwell.
- Smithies, Declan. 2012. Moore's Paradox and the Accessibility of Justification. *Philosophy and Phenomenological Research* 85 (2):273–300.
- Strawson, Peter F. 1963. Carnap's Views on Conceptual Systems versus Natural Languages in Analytic Philosophy. In P. A. Schilpp (ed.), *The Philosophy of Rudolf Carnap*. La Salle: Open Court.
- Sylvan, Kurt. MS. On Divorcing the Rational and the Justified in Epistemology.
- Thomasson, Amie L. 2016. Metaphysical Disputes and Metalinguistic Negotiation. *Analytic Philosophy* 57 (4).
- Thomasson, Amie L. Chapter 21, this volume. A Pragmatic Method for Normative Conceptual Work. In A. Burgess, H. Cappelen and D. Plunkett (eds.), *Conceptual Engineering and Conceptual Ethics*. Oxford: Oxford University Press.
- Tiffany, Evan. 2007. Deflationary Normative Pluralism. *Canadian Journal of Philosophy* 37 (5):231–62.
- Väyrynen, Pekka. 2013. *The Lewd, the Rude and the Nasty: A Study of Thick Concepts in Ethics*. Oxford: Oxford University Press.
- Velleman, J. David. 1988. Brandt's Definition of "Good". *Philosophical Review* 97 (3):353–71.
- Weatherston, Brian. 2008. Deontology and Descartes's Demon. *Journal of Philosophy* 105 (9):540–69.
- Wedgwood, Ralph. 2007. *The Nature of Normativity*. Oxford: Oxford University Press.
- Wedgwood, Ralph. 2012. Justified Inference. *Synthese* 189 (2):1–23.
- Williams, Bernard. 1985. *Ethics and the Limits of Philosophy*. Cambridge, MA: Harvard University Press.
- Williamson, Timothy. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.
- Wood, Allen W. 1999. *Karl Marx*. London: Routledge.