JURISPRUDENCE
2020, VOL. 11, NO. 2, 169–212
https://doi.org/10.1080/20403313.2020.1715104

Routledge
Taylor & Francis Group

Check for updates

Conceptual truths, evolution, and reliability about authoritative normativity

David Plunkett

Department of Philosophy, Dartmouth College, Hanover, NH, USA

ABSTRACT

An important challenge for non-naturalistic moral realism is that it seems hard to reconcile it with the (purported) fact of our reliability in forming correct moral beliefs. Some philosophers (including Cuneo and Shafer-Landau) have argued that we can appeal to conceptual truths about our moral concepts in order to respond to this challenge. Call this "the conceptual strategy". The conceptual strategy faces a problem: it isn't clear that the relevant moral concepts are "extension-revealing" in the way that the conceptual strategy needs them to be. A further problem stems from the tradeoff between the "extension-revealing" and the "authority-revealing" aspects of normative concepts. To underscore the import of these issues, I discuss a version of the reliability challenge that concerns authoritative normative facts in ethics (rather than moral facts). The problems I identify for the conceptual strategy carry over to versions of it that are used in response to a range of epistemological arguments in ethics that (like the reliability challenge) are tied to Street's "Darwinian Dilemma" argument. These problems also bear on the prospects of the conceptual strategy for explaining our (purported) reliability in other domains (e.g., in epistemology, law, and politics).

KEYWORDS

Reliability challenge; Darwinian Dilemma; non-naturalism; metaethics; normativity

Introduction

Many people, including many philosophers and biologists, have appealed to considerations from evolutionary biology to attack well-known views about the nature of morality, or, relatedly, widespread moral beliefs. In this paper, I discuss an argument of this kind, which has received significant attention in recent years. The argument is a kind of *reliability challenge*. In short, the argument is that it is hard to reconcile certain views about the nature of morality (e.g., non-naturalistic realism) with the (purported) fact of our reliability in forming correct moral beliefs. This counts against the plausibility of the views in question (e.g., non-naturalistic realism). Considerations from evolutionary biology (purportedly) serve to make this reliability challenge vivid, as well as (purportedly) provide evidence against certain possible responses on behalf of the views under attack.[1]

CONTACT David Plunkett david.plunkett@dartmouth.edu

[1]My formulation of this challenge draws heavily on David Enoch, 'The Epistemological Challenge to Metanormative Realism: How Best to Understand It, and How To Cope With It' (2010) 148 Philosophical Studies 413 202, David Enoch, *Taking Morality Seriously: A Defense of Robust Realism* (Oxford University Press 2011) ch. 7, and Joshua Schechter, 'Explanatory Challenges in Metaethics' in Tristram McPherson and David Plunkett (eds), *The Routledge Handbook of Metaethics*

© 2020 Informa UK Limited, trading as Taylor & Francis Group

The reliability challenge that I focus on has been proposed as an interpretation of – or way of improving upon core strands of – Sharon Street's "Darwinian Dilemma" argument.[2] I am sympathetic to this reading: I think that a version of the reliability challenge is at the core of Street's argument. At the same time, I don't think that it covers *all* of the philosophically important things going on in that argument.[3] These interpretative issues, however, are not my focus in this paper. Rather, my focus is on the reliability challenge itself. In particular, in this paper, I consider a response to the reliability challenge (a response that I call *the conceptual strategy*) and raise a series of issues for it. Importantly, I argue that many of these issues carry over to versions of the conceptual strategy that could be (and, in some cases, have been) used in response to a range of arguments in the vicinity of Street's Darwinian Dilemma argument.

In order to respond to the reliability challenge – as well as to other arguments related to Street's Darwinian Dilemma argument – one idea is to appeal to conceptual truths about our moral concepts (e.g., MORALITY, MORAL REASON, or MORAL WRONGNESS).[4] The basic idea behind this kind of strategy is as follows. Start with the idea that the relevant conceptual truths help fix the *topic* or *subject matter* we are thinking about when we engage in moral thought. Because of this, and combined with an idea about what those conceptual truths *are*, some have argued that if we have sufficient mastery of the relevant moral concepts and deploy them in our thinking, we can then be reasonably confident that we are reliable in our core moral judgments. Or, more precisely, we can be reasonably confident given the assumption that there are moral facts in our world (e.g., that moral nihilism is wrong). Call this line of thought *the conceptual strategy*. Terence Cuneo and Russ Shafer-Landau have advanced what I take to be the most developed version of this strategy, in response to their own preferred reconstruction of Street's Darwinian Dilemma argument.[5] In what follows, I focus on a schematic version of the strategy as applied to the reliability challenge. I return later in the paper to some of the details of the particular ways in which Cuneo and Shafer-Landau develop their version of the conceptual strategy, and its application to other ways of reconstructing Street's core argument. The issues with the conceptual strategy I target are easiest to see when we think of the conceptual strategy in these

(Routledge 2017). Both Enoch and Schechter understand this challenge for non-naturalistic moral realism as closely parallel to the challenge to mathematical Platonism developed by Hartry Field, *Realism, Mathematics, and Modality* (Blackwell 1991), drawing on Paul Benacerraf, 'Mathematical Truth' (1973) 70 Journal of Philosophy 661. See also Kieran Setiya, *Knowing Right from Wrong* (Oxford University Press 2012), Justin Clarke-Doane, 'Morality and Mathematics: The Evolutionary Challenge' (2012) 122 Ethics 313, and Michael Klenk, 'Old Wine in New Bottles' (2017) 20 Ethical Theory and Moral Practice 781 for similar readings.

[2]Sharon Street, 'A Darwinian Dilemma for Realist Theories of Value' (2006) 127 Philosophical Studies 109. See also Sharon Street, 'Reply to Copp: Naturalism, Normativity, and the Varieties of Realism Worth Worrying About' (2008) 18 Philosophical Issues 207 and Sharon Street, 'Mind-Independence Without the Mystery: Why Quasi-Realists Can't Have It Both Ways' in Russ Shafer-Landau (ed), *Oxford Studies in Metaethics*, vol 6 (Oxford University Press 2011).

[3]Many who take the reliability challenge I focus on to be at the heart of Street's challenge also agree with this qualification. (For example, see Schechter (n 1)). See Selim Berker, 'Does Evolutionary Psychology Show That Normativity Is Mind-Dependent?' in Justin D'Arms and Daniel Jacobson (eds), *Moral Psychology and Human Agency: Philosophical Essays on the Science of Ethics* (Oxford University Press 2014) for what I take to be one of the most thorough and accurate reconstructions of the different facets of Street's Darwinian Dilemma argument.

[4]In this paper, I use small caps (e.g., MORALITY) to designate concepts. Single quotation marks (e.g., 'morality') are used strictly to mention linguistic items. Double quotation marks (e.g., "morality") are used for a variety of tasks, including quoting others' words, scare quotes, and mixes of use and mention.

[5]It should also be noted here that Cuneo and Shafer-Landau explicitly claim that their argument should generalise to other important versions of the Street-inspired argument they discuss (including, I take it, the reliability challenge that I focus on). See Terence Cuneo and Russ Shafer-Landau, 'The Moral Fixed Points: New Directions for Moral Nonnaturalism' (2014) 171 Philosophical Studies 399 (n 57) 425.

schematic terms, and when we focus on its application to one particular kind of argument (e.g., the reliability challenge).[6]

For reasons that I discuss below, the conceptual strategy relies on the idea that our relevant moral concepts are fairly *extension-revealing* in the following sense: competence with those concepts tends to make facts about their extensions accessible to an actual competent person employing those concepts.[7] Moreover, given the nature of the reliability challenge and the conceptual strategy, the proponent of the conceptual strategy requires (a) that these concepts are extension-revealing *because of* conceptual truths about the concepts themselves (rather than, for example, agents being good at seeing the extension of those concepts because of particularly good moral education) and (b) that these concepts are so revealing to actual ordinary people in our current social-historical context. But, as I discuss below, there is a core issue for the proponent of the conceptual strategy. This is that it is not clear that our moral concepts are in fact extension-revealing in this way.

This means that the proponent of the conceptual strategy needs to defend a controversial account of our relevant moral concepts, which vindicates the claim that these moral concepts are fairly extension-revealing. Furthermore, I argue, there is a serious connected issue that looms for the conceptual strategy. This is that, in general, there is a tradeoff between how extension-revealing a concept is and how revealing it is of facts that, at first blush, we might describe as facts about what *we really and truly should do* (or, relatedly, what we really and truly should believe, feel, etc.).[8] In this paper, I discuss the facts that I want to target here in terms of *authoritative normativity*.[9] The more extension-revealing a normative (or evaluative) concept is, the less confident we should be that it is revealing of authoritative normative facts. In other words, the less confident we should be that it is *authority-revealing*. In what follows, I take a normative (or evaluative) concept to be authority-revealing to the extent that competence with that concept tends to make its authority accessible to an actual competent person employing that

[6]In proceeding this way, I follow Schechter (n 1), who also discusses the conceptual strategy as a schematic type of response, in abstraction from the details of how it is developed by different philosophers (e.g., by Cuneo and Shafer-Landau). It should be noted that, in contrast, Cuneo and Shafer-Landau don't present their argument as a version of a more schematic type of argument (e.g., as a version of what I am calling "the conceptual strategy"). This raises the possibility that perhaps it is a mistake to read them this way at all. I don't think that is so, for the reasons I give in this paper. But, for those who are worried about it, note that I return to this worry at the end of the paper. In short, I there argue that even if I am incorrect in reading them as advancing a version of the conceptual strategy, my core argument still has upshots for assessing their argument.

[7]My language here of "extension-revealing" concepts draws from Tristram McPherson and David Plunkett, 'Conceptual Ethics and The Methodology of Normative Inquiry' in Alexis Burgess, Herman Cappelen and David Plunkett (eds), *Conceptual Engineering and Conceptual Ethics* (Oxford University Press 2020). It is not obvious that all concepts yield extensions in the strict sense. If not all of them do, then one can focus on whatever concepts contribute to thought, and then modify the definition of "extension-revealing" accordingly. Since this complication doesn't matter for my purposes in this paper, I stick with the simplifying assumption that concepts yield extensions in the definition of "extension-revealing" concepts that I work with in this paper.

[8]The use of 'truly' here might throw off some philosophers drawn to metanormative positions according to which there are not normative "truths", strictly speaking. This kind of metanormative issue, however, isn't what I am flagging with my use of 'truly' here. As I explain below, the core contrast I have in mind here is one about (roughly) normative significance, which doesn't rely on the metanormative view that there is, strictly speaking, normative "truth" as such.

[9]My use of the label 'authoritative normativity' draws from Tristram McPherson, 'Authoritatively Normative Concepts' in Russ Shafer-Landau (ed), *Oxford Studies in Metaethics*, vol 13 (Oxford University Press 2018). In other work (such as David Plunkett, 'Robust Normativity, Morality, and Legal Positivism' in David Plunkett, Scott Shapiro and Kevin Toh (eds), *Dimensions of Normativity: New Essays on Metaethics and General Jurisprudence* (Oxford University Press 2019)) I use the label 'robust normativity' as an alternative label for 'authoritative normativity', drawing on Tristram McPherson, 'Against Quietist Normative Realism' (2011) 154 Philosophical Studies 223. I use the label 'authoritative normativity' in this paper, given its better fit with the other terminology I employ in this paper (e.g., my discussion of "authority-revealing" concepts).

concept.[10] As I will understand it in this paper, this involves making accessible the authoritative nature of the normative facts that are in the extension of this concept.[11]

Importantly, a core reason why many people care about explaining our reliability about morality – as well as why they care about a range of other epistemological issues concerning morality – is because they think morality *matters* in significant ways. Put roughly, many think that facts about what morality requires, recommends, prohibits, etc. bear heavily on what we really should do, all-things-considered. In other words: they think there is a close connection between the dictates of morality and the authoritative normative facts about how we should act, all-things-considered. This is different than how many think about the significance of facts about what etiquette requires, recommends, prohibits, etc., or facts about what the rules of a randomly chosen social club require, recommend, prohibit, etc. So, if there is a tradeoff between the authority-revealing and extension-revealing aspects of normative (or evaluative) concepts, this matters for the conceptual strategy.[12] In short, the more extension-revealing the proponent of the conceptual strategy claims our relevant moral concepts are, the more difficult it looks to square the strategy with the reasons why many are interested in the reliability challenge about morality in the first place. Or at least so I argue.

In order to underscore the importance of this tradeoff between the extension-revealing and the authority-revealing aspects of normative concepts, I look at a version of the reliability challenge that is about *authoritatively normative* facts. The conceptual strategy has dim chances of success for responding to this argument, even if it can be made to work in responding to the parallel argument about morality. This is because authoritatively normative concepts are *not* extension-revealing in the right way. Consider, for example, the concept AUTHORITATIVE OUGHT, as well as the concept (call it GENUINE NORMATIVE REASON) about those considerations that "count in favor of" what an agent *authoritatively* ought to do (rather than, for example, count in favour of what she *legally* or *morally* ought to do). In short, competent users of those concepts are often not in a strong position to identify the extension of those concepts. Moreover, *if* they are in a strong position to do so it is *not* because there are conceptual truths about the concepts themselves that make this so. Indeed, we might fully analyse those concepts without that providing much help – and perhaps *no* help – in identifying which particular acts we authoritatively should perform (or what relative strengths authoritative normative reasons have, etc.). The upshot of this discussion is that the conceptual strategy has limited dialectical force against the version of the reliability challenge that matters most in ethics (or which, at the very least, matters *just as much as* the parallel argument about morality).

[10]For further, general discussion about the tradeoff between how extension-revealing a concept is and how authority-revealing it is, see McPherson and Plunkett, 'Conceptual Ethics and The Methodology of Normative Inquiry' (n 7). My discussion in this paper (including my glosses on what it is for a concept to be extension-revealing or authority-revealing) draws heavily from our work in that paper.

[11]This idea involves thinking that authoritative normativity is a property of facts, in addition to (or perhaps instead of) being a property of concepts. Embracing this idea makes it easier for me to discuss the core issues involved in this paper. But nothing I say depends on this thesis being true, strictly speaking. So, if one rejects this idea, then I invite one to rephrase the discussion below accordingly. For example: one could put things in terms of facts under a given description, and then think that the relevant authoritative normativity is a feature of the concept(s) used for that description. See McPherson, 'Authoritatively Normative Concepts' (n 9) and Matti Eklund, *Choosing Normative Concepts* (Oxford University Press 2017) chs. 4–5 for connected discussion.

[12]In what follows, I will often use 'normative' as a shorthand for covering both the evaluative (e.g., concerning value, better and worse, etc.) as well as the normative more narrowly construed (e.g., concerning should, ought, etc.).

As I discuss in more detail later on, much of what I say in this paper echoes core points that Street makes in previous work – especially in her reply to David Copp, who can plausibly be read as advancing something akin to (or perhaps even a version of) the conceptual strategy in response to Street's Darwinian Dilemma argument.[13] I think that Street's points that I draw attention to have not been sufficiently appreciated by those (such as Cuneo and Shafer-Landau) who have advocated for a version of the conceptual strategy following her reply to Copp. Yet my goal here is not simply to repeat what I take to be insightful points already made by Street, and show how they apply to Cuneo and Shafer-Landau or others to whom she has not explicitly responded. Instead, I aim to put forward my own response to the conceptual strategy, which uses my own preferred conceptual resources and ideas, and targets the strategy at a high level of abstraction. More specifically, I aim to start with (a) a highly schematic account of concepts, in combination with (b) a general way of thinking about authoritatively normative concepts and (c) an associated idea about the tradeoff between the extension-revealing and authority-revealing aspects of normative concepts, in order to deliver a response to the conceptual strategy that is both ecumenical and powerful. If I am on the right track, this way of proceeding should help make clear to a range of different philosophers, with a range of different views about the nature of concepts and of normativity, why the conceptual strategy (including the specific version of it that Cuneo and Shafer-Landau develop) faces some serious issues. Some philosophers might well be convinced by the argument I give, even if they are not convinced by Street's previous arguments (e.g., that she makes in responding to Copp) that draw on similar ideas.

Before I begin, I want to note two further things about my aims in this paper. First, I aim to make an argument about the conceptual strategy that has wider implications than just its application to the reliability challenge for non-naturalistic realism in ethics. More specifically, I argue that my core argument in this paper generalises in three important ways. First, I argue that it raises an issue for anyone who wants to appeal to conceptual facts to explain our reliability in moral or ethical judgments, regardless of whether they endorse a version of non-naturalistic realism or not. Second, I argue that it has implications for the idea of using the conceptual strategy to respond to parallel reliability challenges that arise in other areas: for example, challenges about our (purported) reliability about mathematical facts, epistemological facts, or legal facts. One particularly interesting case I focus on is the issue of what explains our reliability in normative issues about concepts themselves, including issues about which concepts we should use. Third, I argue that my core argument has implications for other epistemological challenges in ethics, including others involving evolutionary considerations that some have taken to be tied to Street's Darwinian Dilemma argument. I illustrate this by discussing Katia Vavova's preferred reading of the core of Street's Darwinian Dilemma argument. Vavova reads Street's argument as a debunking argument about *irrelevant influences,* rather than a challenge about explaining our reliability. Vavova argues that the conceptual strategy offers a promising

[13]Copp appeals to "the nature of morality and the truth conditions of moral propositions" in responding to Street's Darwinian Dilemma argument. See David Copp, 'Darwinian Skepticism About Moral Realism' (2008) 18 Philosophical Issues 186 202. Copp does not explicitly claim that those "truth conditions" come from our moral concepts as such, but it is plausible to read him that way. If that reading of him is correct, then Copp's response to Street (which, importantly, involves his defending a form of *naturalistic* moral realism, rather than a *non*-naturalistic form of it) can also be understood as involving a version of the conceptual strategy.

solution to that argument.[14] I argue that the conceptual strategy faces the same core issues when applied to the irrelevant influences argument as it does when applied to the reliability challenge.

A secondary aim I have in this paper is to underscore the importance of two connected methodological points that I have argued for in previous work. Both methodological points resonate in important ways with Street's own discussion of her Darwinian Dilemma argument, as I will discuss later in this paper. To see what these points are, consider that many philosophers use (or least aspire to use) authoritatively normative concepts in their normative theorising. This is true not only in ethics, but also in epistemology, political philosophy, legal philosophy, and other areas of philosophy. However, at the same time, in many contexts in normative (or metanormative) theorising, many proceed as if it doesn't really matter whether we focus on authoritatively normative concepts as such, or instead focus on other normative concepts (e.g., specifically *moral* concepts, whose relationship to authoritative normativity is contested among philosophers). In some cases, this idea is correct: we can make progress on the core issues at hand without being precise about exactly which normative concepts are being used or discussed. But in many other cases, this is a mistake. If I am right, the case at hand is one of them. For, if successful, my argument shows how switching from discussion about moral concepts to authoritatively normative concepts *as such* (or vice versa) can matter a great deal to our assessment of the relevant philosophical challenges at hand (e.g., the reliability challenge), and in our assessment of the strategies used to respond to them (e.g., the conceptual strategy). The idea that this sort of switch can often matter in this way, and that its importance is underappreciated by many, is the first methodological point. The second point is closely connected, but broader. This is that paying attention to which specific normative concepts we might deploy (or which we might attempt to understand) – and, in particular, paying attention to (a) a variety of different "moral-ish" concepts and (b) authoritatively normative concepts – can matter a great deal to our normative and metanormative theorising, in ways that are often underappreciated in work in these areas.[15]

I break up my work in what follows into eight main sections, followed by a conclusion. In section 1, I discuss Street's Darwinian Dilemma argument. In section 2, I discuss Non-Naturalistic Moral Realism, which is the metamoral view that Street's argument most clearly targets. In section 3, I explain the reliability challenge for Non-Naturalistic Moral Realism, underscoring its connections to Street's argument and to the sort of evolutionary considerations she appeals to. In section 4, I explain the conceptual strategy, conceived of here as a response to the reliability challenge for Non-Naturalistic Moral Realism. In section 5, I argue that the conceptual strategy faces a serious issue that

[14]Vavova expresses optimism about the conceptual strategy at the end of Katia Vavova, 'Evolutionary Debunking of Moral Realism' (2015) 10 Philosophy Compass 104, in which she also argues that Street's Darwinian Dilemma argument is best understood as a kind of irrelevant influences argument. It should be noted that her optimism has to do with the broad kind of strategy it is, rather than the details of Cuneo and Shafer-Landau's particular version of it.

[15]I argue for a version of both methodological points in McPherson and Plunkett, 'Conceptual Ethics and The Methodology of Normative Inquiry' (n 7), Tristram McPherson and David Plunkett, 'The Nature and Explanatory Ambitions of Metaethics' in Tristram McPherson and David Plunkett (eds), *The Routledge Handbook of Metaethics* (Routledge 2017), David Plunkett and Scott Shapiro, 'Law, Morality, and Everything Else: General Jurisprudence as a Branch of Metanormative Inquiry' (2017) 128 Ethics 37, and Plunkett, 'Robust Normativity, Morality, and Legal Positivism' (n 9). These methodological points also play an important role in the discussion in David Plunkett and Tim Sundell, 'Disagreement and the Semantics of Normative and Evaluative Terms' (2013) 13 Philosophers' Imprint 1 and David Plunkett, 'Justice, Non-Human Animals, and the Methodology of Political Philosophy' (2016) 7 Jurisprudence 1.

stems from the tradeoff between the extension-revealing and authority-revealing aspects of normative concepts. I use a discussion of "fitting attitude" accounts of relevant moral concepts to help illustrate my points. In section 6, I turn to the reliability challenge for authoritative normative facts in ethics, rather than for moral facts. I argue that the conceptual strategy is even less promising for this version of the reliability challenge, and, moreover, that this discussion matters for evaluating its promise in the moral case as well. In section 7, I consider two important objections to my argument about the conceptual strategy, one of which involves discussion of Cuneo and Shafer-Landau's specific theory of concepts. In section 8, I explain how my discussion matters for the prospects of using a version of the conceptual strategy to respond to a range of reliability challenges in ethics and beyond, as well to other arguments tied to Street's Darwinian Dilemma argument.

1. Street's Darwinian Dilemma

In order to help situate my discussion of the reliability challenge – as well as to help situate how my discussion relates to a range of arguments made by others about what I am calling "the conceptual strategy" – a bit of intellectual context will be helpful. As I noted above, the reliability challenge I focus on has been proposed as a way of understanding the core of Street's Darwinian Dilemma argument. Given this context, it makes sense to start by briefly glossing Street's own argument, in terms that are close to her own.

Street puts her basic challenge as follows. She begins with "the observation that evolutionary forces have played a tremendous role in shaping the content of human evaluative attitudes. For example, we tend to view our survival as good, our children's lives as valuable, and the fact that someone has helped us as a reason to help that person in return".[16] Consider the following proposed explanation for why we have evolved to tend to have such normative or evaluative attitudes – roughly, pro-attitudes, desires, valuings, and the like that underlie our normative and evaluative judgments – and make corresponding judgments. In short, those who had attitudes and made judgments that ran counter to this – such as, for example, those who didn't value their own survival, or who didn't value the lives of their children – tended to produce less offspring that survived. Or at least they would have done so, had they actually existed and made such judgments. Of course, we (or at least many of us) also think that the judgments of the sort that Street glosses are not only ones that evolution pushed us to make. We also think that they are justified and true, and that they form the basis for further justified (and true) normative judgments. This means the following: "[t]here is a striking coincidence between the normative judgments we human beings think are true, and the normative judgments that evolutionary forces pushed us in the direction of making".[17] This kind of striking coincidence "calls out" for explanation. The problem is that certain kinds of views have a difficult time explaining this coincidence. In particular, Street thinks that views that endorse (or entail) a kind of *attitude-independence* about the basic normative facts, according to which (roughly) the basic normative facts don't depend on facts about the nature or content of the attitudes we have, have a difficult time explaining this. Indeed, she

[16]Street, 'Reply to Copp' (n 2) ibid, in the context of summarizing her core argument from Street, 'A Darwinian Dilemma for Realist Theories of Value' (n 2).
[17]Street, 'Reply to Copp' (n 2).

argues that they might have no good explanation available to them at all. Street uses the term 'realism' to refer to the relevant kind of attitude-independence.[18] One such view that endorses the relevant kind of attitude-independence – a view I focus on below in my discussion of the reliability challenge – is a non-naturalistic realist view about normativity.

Street puts her challenge here in terms of *normative* judgments and truths, as my gloss of her argument thus far has done. However, Street's basic kind of challenge may also be focused specifically on *moral* judgments and truths. For reasons that emerge in what follows – including the fact that the conceptual strategy has been chiefly developed with respect to moral concepts in particular – I start with that focus here. This switch, we will see, is hardly without consequence. Indeed, part of my argument is that this switch makes a significant difference to the prospects of the conceptual strategy. But, to see this, we are best off by starting with a focus on morality and then returning to this issue later on.

Street argues that realists face a dilemma in confronting this challenge (again, for purposes of argument, conceived of for now as one about morality). On the one hand, they can hold that evolutionary forces tended to push our moral judgments (and connected evaluative attitudes) toward the moral truth, as conceived of by the realist. On the other hand, they can hold that evolutionary forces tended to *not* push our moral judgments (and connected evaluative attitudes) toward the moral truth, as conceived of by the realist, or, relatedly, tended to neither push them toward nor away from such truth. As Selim Berker puts it "[t]alk of pushing here is of course metaphorical: the crucial issue is whether evolutionary forces have tended to influence our judgments about reasons, values, duties, and other normative matters [in our case for now, *moral* matters] in such a way as to make them line up with the attitude-independent facts about such matters. The first horn holds that this is the case; the second horn holds that it is not".[19] According to Street, if the realist takes the first horn of the dilemma, she ends up committed to a false empirical claim, which we see by looking at the actual evolutionary facts. On the second horn of the dilemma, Street claims the realist is stuck with an unexplained mystery for why there is this sort of "striking coincidence" between our moral judgments and the (purported) moral truths. Moreover, she claims that the realist is forced to hold the "skeptical conclusion" that our moral judgments are "in all likelihood hopelessly off track".[20] If Street is correct that realists face this dilemma, then this would be a serious blow to the plausibility of realism. Both horns of the dilemma are bad for realism. But other competing views – in particular, *attitude-dependent* views that ground the fundamental moral facts in facts about our attitudes – do better, or at least so Street claims. So we have a powerful reason to reject moral realism. Or at least so the basic argument goes.[21]

[18]I have reservations about this choice of label here, so I don't follow her in using 'realism' in this way throughout the paper. For some of the reasons why I worry about this choice of label, see Billy Dunaway, 'Realism and Objectivity' in Tristram McPherson and David Plunkett (eds), *The Routledge Handbook of Metaethics* (Routledge 2017).

[19]Berker (n 3).

[20]Street, 'Reply to Copp' (n 2).

[21]One more brief note on context. Street's argument is often labelled as a "debunking argument" about our ethical (or moral) beliefs. See, for example, Vavova, 'Evolutionary Debunking of Moral Realism' (n 14), amongst many. I am skeptical about the usefulness of this label for her argument. This is because, even if Street's Darwinian Dilemma argument is successful, it is consistent with our maintaining confidence in *many* of the normative beliefs that are widespread amongst both ordinary people and philosophers in our social/historical context. The argument (as I am understanding it here) targets specific views about the nature of ethics (conceived of as incorporating the entirety of our normative judgments about how to live, what to do, etc.). (Or, as I have been glossing it above, it could, in theory, be run against specific views

I think there is a lot to be gained by paying careful attention to Street's Darwinian Dilemma argument on its own terms.[22] My aim here, however, is not to focus on Street's argument in particular. Rather, my goal is to look at the conceptual strategy as applied to the reliability challenge, which a number of philosophers (I think rightly) have taken to be a good way of understanding a core strand of Street's argument. Even if I am wrong that this is a good way of reading Street's core argument, proceeding in this way can allow us to clearly see the issues involved in the conceptual strategy that I aim to highlight. I return later in the paper to how my arguments generalise to a number of other ways of thinking about Street's argument, as well as to a range of connected arguments in ethics.

2. Non-naturalistic moral realism

As I explained in the last section, Street uses her Darwinian Dilemma argument to attack views that endorse a kind of *attitude-independence* about (the relevant kind of) normative facts. It might well be that Street's argument ultimately does show a problem for all views that endorse such a thesis. However, it is easiest to see how her argument raises an issue for a more particular kind of view that involves attitude-independence: namely, the view I call "non-naturalistic realism". This is also the kind of view that, I argue, the reliability challenge most clearly raises an issue for. In line with the initial focus I signalled in the last section, I start with understanding non-naturalistic realism as applied to *morality* in particular. Following one standard way of carving up the terrain, in this paper, I take *morality* to be a distinct topic from *ethics* as a whole, which (in this paper) I take to concern the full range of normative issues about how we should live, what we should do, etc. There is no fully theory-neutral way of explaining what that means. But, put roughly, the basic cut I have in mind can be illustrated as follows: ethics incorporates kinds of considerations (e.g., prudential considerations) that might be left out of morality, or deals with kinds of issues (e.g., self-regarding issues, or issues of minor importance) that might be left out of morality.[23] This distinction between morality and ethics is important in the context of this paper, as I will discuss in more detail below.

As I understand it, non-naturalistic moral realism is a kind of view within *metamoral* inquiry. Metamoral inquiry, according to the view that I favor, aims to explain how actual moral thought and talk – and what (if anything) it is distinctively about – fits into reality.[24]

about the nature of morality in particular, rather than of ethics). In particular, Street attacks the thesis that the relevant normative facts (e.g., the ethical facts) are *attitude-independent*. If her arguments work, we might well need to give up on the thesis of attitude-independence, or decrease the credence we have in it. For some people, that won't amount to any sort of "debunking" of their beliefs at all, since they didn't hold that view to start with (or give it much credence). And, moreover, it is hardly obvious that *most* people in our social/historical context endorse the relevant kind of attitude-independence about the kinds of ethical facts Street cares about. Parallel reasons make me skeptical about the usefulness of the label of "debunking argument" for the reliability challenge I discuss below.

[22]For one of the most careful discussions with respect to the details of Street's work on the Darwinian Dilemma argument, see Berker (n 3).

[23]For connected discussion, see McPherson and Plunkett, 'The Nature and Explanatory Ambitions of Metaethics' (n 15) and Stephen Darwall, 'Ethics and Morality' in Tristram McPherson and David Plunkett (eds), *The Routledge Handbook of Metaethics* (Routledge 2017).

[24]This way of thinking about metamoral inquiry draws from McPherson and Plunkett, 'The Nature and Explanatory Ambitions of Metaethics' (n 15) and Plunkett and Shapiro (n 15). I am here using the term 'metamoral' rather than the more common term 'metaethical' to emphasise the import of the distinction between morality and ethics in this context. As I understand it, metamoral inquiry focuses on specifically *moral* thought, talk, and reality, whereas metaethical inquiry focuses on *ethical* thought, talk, and reality. See McPherson and Plunkett, 'The Nature and Explanatory Ambitions of

For our purposes here, we can understand "reality" as the totality of what there is and what it is like. Non-naturalistic realism is a kind of research programme within this explanatory project.

For our purposes here, we can understand non-naturalistic moral realism as a research programme that stems from a basic commitment to the following four theses (or something close to them):

(1) *Cognitivism*: At the most fundamental explanatory level, moral judgments consist in ordinary beliefs (attitudes with a mind-to-world direction of fit). These beliefs are about moral reality (moral facts, properties, relations, etc.).

For our purposes here, we can take this to have the following upshot: moral judgments are capable of being true or false. They are true insofar as they correctly represent reality and false insofar as they do not.

(2) *Descriptivism*: Moral talk is explained in terms of moral beliefs.

Roughly, this means that the meaning of moral talk is explained by the moral beliefs we use that talk to express. This means that moral talk is also capable of being true or false, depending on which beliefs are expressed by that talk.

(3) *Non-Naturalism*: Moral reality is non-naturalistic.

For our purposes here, we can take this to mean that moral reality – the part of reality that moral thought and talk is distinctively about – is fundamentally discontinuous with the "naturalistic" part of reality studied by the natural and social sciences. Roughly, this means that moral reality is not identical to, fully grounded in, or definable in terms of that part of reality. Among other things, I take that to mean that moral reality is *causally* inert. I also take it to mean that the fundamental moral facts, properties, relations etc. – those that explain other such facts, properties, relations, etc. but which are not themselves explained by any further normative facts – are *constitutively independent* from facts about people, their attitudes, practices, etc. In particular, they don't depend on such facts about those of us who engage in moral thought and talk (i.e., "moral judges").[25]

(4) *Non-Error Theory*: Moral nihilism is false. This means that there really are some moral facts, properties, relations etc. in our world.

Metaethics' (n 15) for further discussion about my understanding of metaethical inquiry in particular, as well as connected discussion about metanormative inquiry more broadly construed.

[25]My characterization of the way in which moral reality is "non-naturalistic" closely parallels the sketch given in David Enoch, "Non-Naturalistic Realism in Metaethics" in Tristram McPherson and David Plunkett (eds), *The Routledge Handbook of Metaethics* (Routledge 2017). For further discussion of the distinction between naturalistic and non-naturalistic accounts of moral reality (or other kinds of normative reality), see Tristram McPherson, 'What is at Stake in Debates among Normative Realists?' (2015) 49 Noûs 123, Stephanie Leary, 'Non-naturalism and Normative Necessities' in Russ Shafer-Landau (ed), *Oxford Studies in Metaethics*, vol 12 (Oxford University Press 2017), and Gideon Rosen, 'Metaphysical Relations in Metaethics' in Tristram McPherson and David Plunkett (eds), *The Routledge Handbook of Metaethics* (Routledge 2017).

There are obviously *lots* of questions to ask about the details of these theses. But for our purposes here, this gives us enough to work with in formulating the relevant arguments.[26]

3. The reliability challenge

A number of philosophers, including David Enoch and Joshua Schechter, think that one of the most important epistemological challenges for non-naturalistic moral realism – as well as the core challenge that Street's work brings out – is how to explain our (apparent) reliability in moral judgment in a way that is consistent with non-naturalistic realism.[27] The basic challenge is as follows. Most of us want to believe – and, moreover, think we are justified in believing – that we are reliable in (at least core parts of) our moral judgments. Of course, many of our moral judgments might be wrong. But, for key parts of our moral judgments, we seem to be on track: e.g., our judgment that killing innocent people for fun is morally wrong, or that (other things being equal) causing people pain is morally bad, etc. Or, at the very least, even if we aren't fully on track here, we are certainly significantly *better than random* in these moral beliefs. We thus have a (purported) correlation between (a) the content of our moral beliefs (at least our core moral beliefs) and (b) the moral facts. This is a striking correlation. It is the kind of correlation that "calls out for explanation". Or, at the very least, it is the kind of thing that would be good to have an explanation of, other things being equal. The trouble is that non-naturalistic moral realists face a significant challenge in providing such an explanation. Moreover, those who endorse a range of well-known competing metamoral views (e.g., forms of naturalistic realism that endorse relevant kinds of judgment-dependence) are better equipped to provide a good explanation. Thus, it seems, non-naturalistic realism has a comparative disadvantage here.

The problem is this. The kind of correlation on the table is between the X facts and the contents of our beliefs about the X facts. One common way to explain this kind of correlation is by a causal story. For example: if the X facts caused the X beliefs, or if they caused other facts that caused those beliefs. Certain kinds of *naturalistic* views of moral facts might be able to provide such a story. But such a causal story is ruled out by the thesis of *Non-naturalism*.

Another common way to explain this kind of correlation is to appeal to a constitutive story. For example: if the X facts are partly constituted by the relevant X beliefs, or in some other (non-casual) way depend on them. But that is not true for the fundamental moral facts, according to non-naturalism. And without that sort of dependence, this kind of constitutive move isn't going to be helpful for the non-naturalist (although it might be helpful, of course, for those who favour competing views, e.g., views that posit a kind of response-dependence or judgment-dependence). The non-naturalist can't appeal to constitutive connections here about the *fundamental* moral facts: facts that help explain other moral facts, but which themselves do not depend on further moral facts. A familiar thought about such facts is that they include facts about what is morally right and wrong in

[26]My understanding of non-naturalistic moral realism above closely parallels the one given of non-naturalistic realism about ethics (rather than morality in particular) in Enoch, 'Non-naturalistic Realism in Metaethics' (n 25). See Enoch, *Taking Morality Seriously* (n 1) for a book-length defense of the kind of non-naturalistic moral realism I am discussing here.

[27]See Enoch, 'The Epistemological Challenge to Metanormative Realism' (n 1), Enoch, *Taking Morality Seriously* (n 1), and Schechter (n 1).

general, and not just in a specific social setting; or at least that they will be *pure* moral facts, not dependent on any descriptive facts. If we are seriously unreliable about *those* facts, then our ability to be reliable about downstream, derivative moral facts seems limited indeed.

The general challenge about reliability on the table can be made vivid by thinking about evolutionary facts. Our moral judgments – and the underlying ideas, emotions, and patterns of thought that are connected to them – have been heavily influenced by evolutionary forces. Are these judgments, which are the product of such evolutionary forces, going to be reliable at tracking the moral facts (as conceived of by the non-naturalistic moral realist)? We (at least those of us who aren't moral skeptics) want to think that they are. Suppose, for the sake of argument, that they in fact are. But, if that is the case, how do we explain this reliability? Given the actual causal story of how we came to have those judgments (and connected ideas, emotions, etc.), the basic worry here – the worry that we lack a good explanation of that reliability – can seem especially pressing. Thus understood, the (purported) facts from evolutionary biology aren't essential to the reliability challenge. As Schechter puts it, "The evolutionary story makes salient a worry the robust moral realist [i.e., the non-naturalistic moral realist] already should have had – namely, that there is no satisfying explanation of how it is that we are reliable about morality".[28]

The evolutionary story, however, plays another role as well. This is that it (purportedly) helps block certain possible responses on behalf of the non-naturalist. For example, suppose that our coming to have *correct* beliefs about the relevant moral facts (as conceived of by the non-naturalist) was itself part of the actual historical explanation of how we came to acquire the basic moral beliefs we have (e.g., the belief that it is wrong to kill one's children, other things being equal). Roughly, on this sort of view, our beliefs about moral facts (as conceived of by the non-naturalist) would then be akin to our beliefs about such things as the location of tigers or which plants are poisonous. This is in the following sense: learning about those facts was evolutionary advantageous, and that is why we formed beliefs about them. Or perhaps the evolutionary story here is something like the following: we came to have a general-purpose capacity for reasoning of the right kind (which we have reason to think was reliable for independent reasons), which we then started using for learning about the moral facts (as conceived of by the non-naturalistic realist). On certain versions of that kind of story, the fact that we successfully represented the moral facts (as conceived of by the non-naturalist) might again be part of the actual historical explanation for the basic moral beliefs we have. If either of those stories turned out to be correct, perhaps they would offer a way to defend non-naturalistic moral realism. The problem is that neither such story is correct – or at least so Street claims.[29] Whether that is right will, of course, depend on the actual evolutionary history. Thus, pointing to the (purported) evolutionary facts can do more than just make the reliability challenge psychologically gripping. Doing so also has the potential to block (or advance) certain kinds of responses on behalf of the non-naturalist.

One important issue about this kind of reliability challenge is how far it extends to other kinds of metamoral views, or other views about the nature and grounds of moral facts that

[28]Schechter (n 1) 454.
[29]See Street, 'A Darwinian Dilemma for Realist Theories of Value' (n 2) and Street, 'Reply to Copp' (n 2).

might or might not be part of a more general metamoral account. The issues here are delicate: for example, whether (and to what degree) this challenge extends to *quasi-realist* metamoral views, or different kinds of *naturalistic realist* views, or views that are neutral on many metamoral questions but that make certain kinds of claims about the way in which moral facts are grounded in facts about our attitudes.[30] Given my purposes here, I am not going to wade into those discussions. Instead, given my focus, I simply want to stick with discussing non-naturalistic realism as glossed above (starting initially about *morality*, and then eventually about authoritative normativity). This is because this kind of view is the one most straightforwardly targeted by the reliability challenge and thus makes the reliability challenge easiest to understand. Later in this paper, I argue that my core points about the conceptual strategy below are independent of the particular metamoral view the strategy is being used to defend.

4. The conceptual strategy

There are a number of different ways that non-naturalistic realists can respond to the reliability challenge. For example, one influential kind of response is a "third-factor" response, according to which there is a "third factor" that is appropriately connected to both our moral beliefs and our moral facts that explains the reliability of our moral beliefs. If that third factor is connected in the right way to each of these things – casually, constitutively, or in some other way – then this might be used to explain the reliability of our moral beliefs, even without a *direct* causal or constitutive connection between our moral beliefs and the moral facts.[31] In what follows, I won't be discussing this popular kind of response (which I take to be among the most promising ones for the non-naturalist to develop), or any number of other responses worth investigating. Instead, I focus solely on the conceptual strategy.[32]

The core line of thought involved in the conceptual strategy can be summed up as follows. There are certain things we know *must* be true about morality, given constraints that are built into our relevant moral concepts (e.g., MORALITY, MORAL RIGHTNESS, or MORALLY PROHIBITED). More precisely: *if* there are moral facts, then, given the nature and content of our moral concepts, we have some sense of what those facts are like. To illustrate, consider the concept MORALLY WRONG. If that concept picks out a non-empty extension in our world, then, given the application-conditions of that concept, we know something about which acts are morally wrong. This is because, in order for an act to be morally wrong, it must be that the concept MORALLY WRONG applies to it. Drawing on this, the idea is then that, given the kind of information about the actual world that most competent users of the relevant moral concepts know, we can be confident about

[30]There are also important issues here that I leave to the side about which metaphysical ideology we should use to state the core metaphysical views at the heart of non-naturalistic moral realism: e.g., whether these views concern the *real definition* of morality vs. *grounds* of moral facts. See Rosen (n 25) and Leary (n 25) for discussion of some of the relevant terrain here. This issue matters in this context for the following reason: if there are a variety of different ways of developing the core metaphysical ideas driving non-naturalistic moral realism, using different bits of metaphysical ideology, then perhaps certain resulting views would be better equipped to deal with the challenge than others.

[31]For examples of this sort of strategy, see Enoch, *Taking Morality Seriously* (n 1), Erik J Wielenberg, 'On the Evolutionary Debunking of Morality' (2010) 120 Ethics 441, Kevin Brosnan, 'Do the Evolutionary Origins of Our Moral Beliefs Undermine Moral Knowledge?' (2011) 26 Biology and Philosophy 51, Copp, 'Darwinian Skepticism About Moral Realism' (n 13), Knut Olav Skarsaune, 'Darwin and Moral Realism: Survival of the Iffiest' (2011) 152 Philosophical Studies 229, and Berker (n 3).

[32]For a good overview of some of the other strategies here worth considering, see Schechter (n 1).

their key moral judgments concerning foundational moral issues (e.g., whether, for creatures like us in worlds like ours, killing others for fun is morally permitted, whether we morally should prefer pain over pleasure, other things being equal, etc.).[33] We might not be able to *fully* explain our reliability using this strategy. One reason is the following: the relevant moral concepts might not have any extensions in our world (e.g., moral nihilism might be true). But *if* the relevant moral concepts do have extensions (in our world), then we can appeal to truths about these concepts to help explain our reliability about key moral judgments. Or at least so the basic thought goes.

In theory, the conceptual strategy could be run using any range of moral concepts; e.g., one could attempt to run it by appealing to truths about *all* of our moral concepts vs. appealing to some subset of them (e.g., MORAL WRONGNESS and MORAL REASON). The basic line of argument I advance below doesn't depend on the details of which moral concepts the proponent of the conceptual strategy focuses on. Thus, I generally talk about the "relevant" moral concepts when talking about the moral concepts that the proponent of the conceptual strategy uses, or illustrate my points using a particular moral concept. Although I won't discuss it further in this paper, there are of course interesting questions about which moral concepts are comparatively better to use when developing the conceptual strategy for defending non-naturalistic moral realism. The question of which ones should be used obviously interacts with the question of which moral judgments are targeted by the specific version of the reliability challenge. (For example, is the argument targeting *all* of our moral judgments, or just our judgments about moral wrongness in particular?) The question of which moral concepts to focus on also interacts with one's views about how to correctly analyze different moral concepts. And, as we will see in what follows, it also interacts with the issue of how extension-revealing and authority-revealing the particular moral concepts are.

Importantly, for the conceptual strategy to work, not everyone who employs the relevant moral concepts needs to be aware of the right analysis of those concepts. The issue is not whether they themselves understand what guides their application of the relevant concepts. Rather, what is crucial is this: that there are conceptual truths built into the relevant moral concepts that *do* guide people's application of those concepts, such that their moral judgments can only be conceptually coherent if they are disposed to make certain judgments in non-defective contexts. Those judgments, the proponent of the conceptual strategy maintains, include ones about which particular acts are morally right and wrong – e.g., the judgment that, for creatures like us in worlds like ours, killing others for fun is *not* morally permitted. So, barring moral nihilism (or simply high levels of ignorance about descriptive facts about the actual world), we actually aren't in horrible shape here with respect to the arguments on the table. Or at least that is so for those of us who have a certain level of basic mastery at using the relevant moral concepts.

The basic (purported) upshot with respect to the reliability challenge is this: if the strategy works, then we can know that we must be reliable in some of our core moral judgments. These include judgments about foundational moral issues (e.g., whether, for creatures like us in worlds like ours, killing others for fun is morally permitted, whether we morally should prefer pain over pleasure, other things being equal, etc.). And given

[33]In using the qualifying phrase "for creatures like us, in worlds like ours", I draw from Cuneo and Shafer-Landau (n 5), who use the phrase "for beings like us, in a worlds such as ours". See ibid 404.

the centrality of those judgments with respect to making other, derivative moral judgments – perhaps, for example, through the process of our making these judgments cohere with each other in a kind of "reflective equilibrium" and through using general-purpose reasoning capacities – we end up in good shape with respect to a whole range of moral judgments. Thus, the "striking coincidence" that needs explaining is explained in part by appealing to truths about our moral concepts.

There are obviously a variety of different ways in which the basic conceptual strategy outlined above can be further fleshed out. One important issue here concerns the nature of concepts. One way of thinking about the basic job description of concepts is that they are constituent components of thoughts – roughly, they are parallel to what words are in linguistic expressions. In turn, we can then think of concepts as individuated by something akin to their cognitive role. To employ the concept DOG, for example, is to think of things *as dogs*; in short, it is to represent them in a particular way. For concepts that involve a representational function, we might think of them as providing (or consisting in) a function from possible worlds to extensions.[34] The concepts themselves might be thought of as abstract objects, which we grasp and use in formulating thoughts. When concepts are thought of in this way, the conceptual strategy will not enable us to block *nihilism* about morality, according to which there are no moral facts. This is because conceptual truths – understood along the lines I am glossing here – do not guarantee anything about the way (non-conceptual) reality as such actually is, and hence do not guarantee that anything *in fact* falls under such concepts.[35] But *if* the extension of these concepts is non-empty in our world, then the conceptual truths can tell us something about what the moral facts in our world are like. Because of this, we can be reasonably confident that we are reliable in our core moral judgments (at least insofar as there are moral facts in our world). Or at least so the basic thought goes.[36]

I am putting the conceptual strategy in terms of "conceptual" truths. However, one can get the core idea across without that terminology. Schechter glosses the conceptual strategy (which he rejects) as resting on the following idea: "if one accepts claims that are far enough away from the genuine moral truths, one is no longer thinking or talking about morality but about something else entirely. So there is a limit to how unreliable thinkers can be about the moral".[37] In this passage, Schechter doesn't use the term 'concepts' at all. For many philosophers (including Schechter), talking about "concepts" is a helpful way of putting forward the idea that he is talking about. This is because many philosophers associate concepts with the following job description: concepts enable us to think certain kinds of thoughts as opposed to others, in part by restricting what something

[34]This function might be complicated in all sorts of ways. For example: it might incorporate the sort of "two-dimensionalist" approach advocated for in Frank Jackson, *From Metaphysics to Ethics: A Defence of Conceptual Analysis* (Clarendon 1998) or David J Chalmers, 'Two-Dimensional Semantics' in Ernie Lepore and Barry Smith (eds), *Oxford Handbook of the Philosophy of Language* (Oxford University Press 2006).

[35]For extended discussion of this general point, see Jackson (n 34).

[36]It should be underscored that, as I am putting it here, the conceptual strategy is compatible with the possibility that there are no moral facts in our world. In other words, it is compatible with a kind of *moral nihilism*. Thus, the version of the conceptual strategy is compatible with the idea that we are in fact deeply unreliable about the moral facts. This would be so if it turned out that moral nihilism is true. A more ambitious (and I think lesss plausible) form of the conceptual strategy might try to rule out the possibility of moral nihilism by appeal to conceptual truths. Since my core worries about the conceptual strategy don't depend on an attempt to use it to rule out moral nihilism, I focus in this paper on the version of the strategy that is compatible with moral nihilsm.

[37]Schechter (n 1) 455.

must be like in order to fall under the relevant concept. Put one way, a widespread thought is that to possess a concept C is to be capable of thinking of things in the C-way, such that some things fall in the extension of C and others do not.[38] But, as Schechter's gloss makes clear, one can still run the same basic strategy even if one wants to eschew talk of "concepts" for whatever reason.

5. Moral concepts

There are a number of things one might worry about with respect to the conceptual strategy. One important issue that Schechter draws attention to is this: "even if there is a limit to how unreliable someone can be about morality, one can still be pretty unreliable about the moral. So we still need an explanation of how we are as reliable as (we think) we are".[39] Another important issue worth flagging – which I briefly return to later – involves the purported role of moral concepts in fixing a topic. The conceptual strategy, at least in its most straightforward version, relies on a tight connection here. In short, the idea is that the relevant moral concepts (which we currently employ) fix corresponding moral topics (e.g., the topic of moral wrongness, moral obligation, etc.), and thus what it is to be literally thinking about *that topic* as such.[40] But perhaps that's wrong. It might be that we can continue to think about a given topic even if our concepts change in significant ways: e.g., that we can continue to think about the topic of moral wrongness even if we no longer employ the specific concept MORAL WRONGNESS.[41] If that is right, then we might worry that perhaps there are certain conceptual truths about our current concepts that aren't actually a reliable guide to morality (or to moral wrongness, etc.). Perhaps our current concepts are even concepts that mislead us in investigating the subject matter at hand, such that we would be better off reforming those concepts, or replacing them with concepts that are better for our purposes at hand. This is a normative view about which concepts we *should* use.[42] Perhaps a crucial discussion at the end of the day (in the context of discussing

[38]Recall that I am here speaking for convenience in terms of extensions (as noted in footnote 7). If some concepts don't yield extensions in the strict sense, we can then focus on whatever it is that they contribute to thought.

[39]Schechter (n 1) 455.

[40]See Cuneo and Shafer-Landau (n 5) 407.

[41]For some different models of how this might be possible, see Herman Cappelen, *Fixing Language: An Essay on Conceptual Engineering* (Oxford University Press 2018), Amie L Thomasson, 'A Pragmatic Method for Normative Conceptual Work' in Alexis Burgess, Herman Cappelen and David Plunkett (eds), *Conceptual Engineering and Conceptual Ethics* (2020), Eklund (n 11), David Plunkett, 'Which Concepts Should We Use?: Metalinguistic Negotiations and The Methodology of Philosophy' (2015) 58 Inquiry 828, and David Plunkett and Tim Sundell, 'Metalinguistic Negotiation and Speaker Error' (2019) Inquiry. Note that, as some of the models above underscore, it might be that we give different answers to the question "is this really about the same topic?" depending on what context we are asking this question in, relative to what purposes. This is because it might well be that we use (and perhaps *should* use) different ways of talking about "sameness of topic" (as well as perhaps different ways of talking about "concepts", "topic", etc.) in different contexts. For example: we might have something different in mind when concerned with "sameness" of topic in the context of ordinary speech vs. linguistics vs. history of science etc. For more on this point, see Plunkett, 'Which Concepts Should We Use?' (n 41) and Tim Sundell, 'Changing the Subject' (Forthcoming) Canadian Journal of Philosophy.

[42]It is thus a view in what Alexis Burgess and I have dubbed "conceptual ethics". See Alexis Burgess and David Plunkett, 'Conceptual Ethics I' (2013) 8 Philosophy Compass 1091 and Alexis Burgess and David Plunkett, 'Conceptual Ethics II' (2013) 8 Philosophy Compass 1102. Put roughly, as Burgess and I understand it, conceptual ethics concerns a range of connected normative and evaluative issues about thought and talk, including questions about which concepts we should use, and why, as well as what we should mean by our words, and why. Conceptual ethics, as Burgess and I understand it, is closely related to what some have called "conceptual engineering" (as in Cappelen (n 41), Kevin Scharp, *Replacing Truth* (Oxford University Press 2013), and Eklund (n 11)). See Herman Cappelen and David Plunkett, 'A Guided Tour of Conceptual Engineering and Conceptual Ethics' in Alexis Burgess, Herman Cappelen and David Plunkett (eds), (Oxford University Press 2020) for an overview of conceptual engineering and conceptual ethics, and the papers in Alexis

the issues in normative epistemology at hand) is really a normative one about which concepts we should use. I will briefly return to this thought later on. For now, however, I want to start with another way of pushing on the conceptual strategy. This concerns the relation between extension-revealing and authority-revealing normative concepts.

In order to appreciate this issue, let's start with the following point. The conceptual strategy relies on the following idea: if we have sufficient mastery of the relevant moral concepts (the sort of mastery that most competent users of those concepts have), then we know some important truths about *which* things are in the extensions of those concepts (insofar as anything is). Put another way, we will have some grip on such things as what morality requires, permits, prohibits, favours, disfavours, etc., what we have *moral* reason to do and what we do not, etc. The idea is that knowing *this* sort of thing is what allows us to know we are talking about the relevant topic at hand (e.g., morality, moral wrongness, etc.). Consider whatever the relevant moral concept is that the proponent of the conceptual strategy appeals to (or, similarly, the relevant cluster of moral concepts). The underlying thought is that the relevant moral concept is fairly *extension-revealing* in the following sense: competence with that concept tends to make facts about its extension accessible to an actual competent person employing that concept.

The notion of extension-revealing is gradational, along multiple dimensions, including *how* easy it is for competent users of the concept to grasp the extension, how *much* of the extension they can grasp, etc. The more extension-revealing the relevant moral concept is along these dimensions, the better off the proponent of the conceptual strategy will be. This is especially so if it is extension-revealing with respect to moral truths or facts that are important for deriving other ones (for example, if facts about moral wrongness are explained by facts about moral reasons). After all, reliability is gradational too, including in the following ways: we might be *more or less* reliable in *more or fewer* of our moral judgments.

The conceptual strategy relies on the idea that the relevant moral concepts *are* quite extension-revealing. In turn, this is based on the idea that the relevant moral concepts involve conceptual truths with direct implications for *which* acts are morally right, wrong, favoured, disfavoured, etc. (at least for creatures like us, in worlds like ours). This idea has a long pedigree. As Cuneo and Shafer-Landau note, when "Thomas Reid defended the claim that the 'first principles of morals' are constitutive of competent moral thought, he endorsed a similar view".[43] More recently, Phillipa Foot argues that, in short, "a consideration bears on moral right and wrong only if it boils down, in the end, to good or ill for someone".[44] Thus, claims Foot, someone who thinks that mere facts about someone clasping her hands or not (in non-extraordinary circumstances) present a serious moral issue does not *just* have a bad moral view. She is also *conceptually* mistaken. Either this person doesn't possess the relevant moral concepts, or else she is incompetent at deploying them in a deep way. On this way of thinking about our moral concepts, they are distinguished from other concepts (at least partly) in terms of

Burgess, Herman Cappelen and David Plunkett, *Conceptual Engineering and Conceptual Ethics* (Oxford University Press 2020) for a range of discussions about the topic.

[43]Cuneo and Shafer-Landau (n 5) 406, discussing Thomas Reid, *Essays on the Active Powers of Man* (Edinburgh University Press 1788/2010).

[44]Allan Gibbard, 'Moral Concepts: Substance and Sentiment' (1992) 6 Philosophical Perspectives 199 203, glossing the kind of view argued for in Philippa Foot, *Virtues and Vices* (University of California Press 1978).

the *substance* of what morality is about – e.g., that morality involves treating other people fairly or with equal concern. In turn, the proponent of the conceptual strategy can argue that these (purported) conceptual truths help explain why these concepts are extension-revealing.

At this point in the discussion, it is important to note an issue about the distinction between so-called "thick" moral concepts (such as, plausibly, JERK or BRAVE) and so-called "thin" ones (such as, plausibly, MORALLY RIGHT and MORAL OUGHT). It is far from clear how exactly this distinction between "thick" and "thin" moral concepts works, how deep it goes, or which concepts are on which side of the divide.[45] But the basic thought is that "thick" concepts involve a mixture of descriptive and normative application-conditions that "thin" ones do not. Thick concepts are good candidates for moral concepts that are fairly extension-revealing. For example, it seems plausible that competent users of JERK are in a pretty good position to know *which people* in their own social/historical context are jerks, given the sort of information about the actual world that most competent users of this concept know. This is partly what makes thick concepts helpful in normative theorising. However, Foot and Cuneo and Shafer-Landau aren't just making a claim about concepts that are widely categorized as paradigmatically "thick" ones. They are also making a claim about ones that are often thought of as "thin" concepts. They think that the relevant moral concepts involve constraints about the substance of morality – e.g., that it is about the *fair* treatment of individuals, that it concerns certain kinds of impartial assessment of the welfare of people, or that it is about certain kinds of behaviour (e.g., promise-keeping) and not others (e.g., hand-clasping, in ordinary circumstances). In virtue of this, these concepts turn out to also be fairly extension-revealing as well.

But should we buy that the relevant thin moral concepts *are* extension-revealing? Or, more precisely, should we think that they are extension-revealing *enough* for us to have much confidence in the conceptual strategy? Consider the very wide range of views that people have had about what is moral and immoral, including many moral views that are *mistaken* (e.g., moral views that condemn interracial marriage, that back slavery, or that regard masturbation as a grave moral failing). One could think that many of these mistaken views involve some kind of *conceptual* mistake; e.g., a confusion about the actual content of the relevant moral concepts being deployed, or else simply a deployment of some *other* "moral-ish" concepts entirely. But many have thought that attributing such a mistake to everyone who has such mistaken moral views comes at too high of a cost. In short, many have thought that this just isn't the kind of error that all (or even most) of these people seem to be making. Consider what is wrong with (for example) someone who thinks that her married brother is a bachelor. If she acknowledges that her brother is in fact married, and yet continues to count him as a bachelor, then a likely diagnosis is that she fails to sufficiently grasp the concept BACHELOR. But it's not that people with bad moral views all fail to get the very *idea* of morality, or of moral right and wrong, or of moral value, etc. For many of these people, it's just that they have *mistaken* or *false* moral beliefs. In this way, they are more like someone who grasps the concept CAT but just has a false *substantive* belief about which things around her actually are cats. Perhaps she makes such a mistake because of bad reasoning on her part, or perhaps

[45]See Debbie Roberts, 'Thick Concepts' in Tristram McPherson and David Plunkett (eds), *The Routledge Handbook of Metaethics* (Routledge 2017) for an overview of thick concepts.

because of being in a poor epistemic situation (e.g., being exposed to misleading evidence). The thought that many such mistakes in moral judgment are *not* conceptual is bolstered, or so many have claimed, by further points: including what the epistemology of moral inquiry looks like vs. that of conceptual analysis, how actual moral arguments unfold, how to best account for the nature of (and extent of) substantive moral disagreement, what the phenomenologies of moral deliberation and argument are like, and how to best make sense of a variety of normative attacks on morality (e.g., of the kind given by Friedrich Nietzsche, Bernard Williams, and Elizabeth Anscombe).[46]

We then face a challenge. The moral concepts that are relevant for the conceptual strategy (e.g., MORALITY, MORAL RIGHTNESS, MORAL WRONGNESS, etc.) seem different from other related concepts that also concern the regulation of behaviour, activity, etc. For example, MORAL RIGHTNESS seems different than LEGAL RIGHTNESS. So what, then, differentiates moral concepts from other normative concepts, if not the sort of *substance-oriented* constraints that Foot and others in her vein propose? One kind of proposal is to focus less on substance-oriented constraints and instead pursue a *fitting attitude* account. The basic idea of a fitting attitude account – whether of moral concepts or something else – is to start with an attitude of a certain kind. We can then ask when that attitude would be *fitting* or *warranted* to have. Moreover, we can focus on when it would be fitting or warranted for the *right* kinds of reasons; what we might call *fittingness* reasons. To get the basic idea in mind, compare it being good or beneficial for you to believe X vs. it being the case that you have good evidence that X is true. The first thing might be a pragmatic reason for you to believe X. But it is not a "reason of the right kind" (if, indeed, it really is a reason *to believe* as opposed to, say, a reason to *want* to believe or to *try to act* in a way that will make yourself believe, etc.).[47] Or at least so the thought here goes.

One kind of fitting attitude account of moral concepts focuses on the fittingness reasons for the emotion of *obligation* (or, in a related vein, a cluster of related emotions). John Stuart Mill characterises this emotion as an "internal sanction of duty … a feeling in our own mind … attendant on violation of duty, which in properly cultivated moral natures rises, in the more serious cases, into shrinking from it as an impossibility".[48] The "feelings of obligation" discussed here involve a sort of "guilt-tinged aversion"[49] to the thought of performing an act. Using such an emotion as the basis for a fitting attitude account, we can analyse the relevant moral judgments as being about when it would be fitting to have this emotion. Here is one such account that works this way, which is a modified version of the sort of account offered by Allan Gibbard:

> Fitting Attitude Analysis of Judgments about Moral Reasons: To judge that R is a moral reason for agent X to φ is to judge that R is a fittingness reason for X to feel obligated to φ.

[46]For further discussion of versions of many of these points, see Gibbard, 'Moral Concepts: Substance and Sentiment' (n 44), Allan Gibbard, *Wise Choices, Apt Feelings: A Theory of Normative Judgment* (Harvard University Press 1990), and Allan Gibbard, 'Moral Feelings and Moral Concepts' (2006) 1 Oxford Studies in Metaethics 195. The critiques of morality at issue here include those from Friedrich Wilhelm Nietzsche, *On the Genealogy of Morality* (Cambridge University Press 1887/1994), Bernard Williams, *Ethics and the Limits of Philosophy* (Harvard University Press 1985), and G.E.M. Anscombe, 'Modern Moral Philosophy' (1958) 33 Philosophy 1.

[47]My gloss of what is involved in a fitting attitude account of this kind draws from Howard Nye, David Plunkett and John Ku, 'Non-Consequentialism Demystified' (2015) 15 Philosophers' Imprint.

[48]John Stuart Mill, *Utilitarianism* (Parker, Son, and Bourn 1863).

[49]Nye, Plunkett and Ku (n 47) 20.

> Fitting Attitude Analysis of Judgments about Moral Wrongness: To judge that it is morally wrong for X to ψ is to judge that it is, on balance, fitting for X to feel obligated not to ψ.[50]

Focus on the account of judgments about moral wrongness. On this account, there are no *conceptual* constraints built into the concept MORAL WRONGNESS that tell us much at all about which particular acts are morally wrong. Rather, given the bare conceptual truths alone, all we know is that, whatever those acts are, they are ones that warrant the specified moral emotions. As far as the concepts go, the question of *which* acts warrant such emotions is left entirely open.

Some versions of a fitting attitude account of moral concepts might of course build *some* further conceptual constraints into when particular emotions are warranted, or else include further conceptual constraints about the relevant moral topic (e.g., that morality must concern issues about human welfare and suffering, or issues about fairness). If enough further substance-oriented constraints are built in through one of these (or other) ways, then perhaps a version of the conceptual strategy might still get off the ground. But the point here is that a fitting attitude account need not involve any such substance-oriented constraints. For example, if the sort of Gibbardian fitting attitude account of judgments about moral wrongness that I sketched above is on the right track, then we can know we are thinking and talking about *moral wrongness* without knowing much (if anything) informative about which particular things (e.g., acts) are morally wrong and which ones are not (insofar as some things are and some things are not). In other words, if a Gibbardian fitting attitude account is on the right track for this moral concept, then this concept isn't going to be very extension-revealing. The same applies for any moral concept for which something akin to this Gibbardian account is correct. So the conceptual strategy will fail if this version of a fitting attitude account is right about the relevant moral concepts.

This means that the proponent of the conceptual strategy has work to do. In short, she has to defend the idea that moral concepts are extension-revealing. And to do that, she is going to need to reject a fitting attitude account along the lines Gibbard proposes, and instead defend a view on which moral concepts are (to a large degree) substance-oriented.

Thus far, the proponent of the conceptual strategy might well be unfazed. After all, she *already* knows she has to argue for a certain view of the content of moral concepts for her strategy to work. That is precisely why proponents of this strategy – e.g., Cuneo and Shafer-Landau in particular – aim to give arguments on behalf of certain substance-oriented views of the relevant moral concepts. One question we can ask is whether these arguments really establish the superiority of substance-oriented ways of thinking about the relevant moral concepts vs. fitting attitude accounts of them (and, in particular, fitting attitude accounts that don't *also* build in many substance-oriented constraints). That is an important question about the analysis of our moral concepts. But my goal here is not to wade into the details of those debates. Rather, the point I want to make concerns a general issue that looms at this point in the dialectic for the proponent of the conceptual strategy. This issue concerns the tradeoff between focusing on normative concepts that are more extension-revealing and those that are more authority-revealing.

[50]Ibid 21, drawing on the account developed in Gibbard, *Wise Choices, Apt Feelings* (n 46). See also Gibbard, 'Moral Concepts: Substance and Sentiment' (n 44), Gibbard, 'Moral Feelings and Moral Concepts' (n 46), Allan Gibbard, *Reconciling our Aims: In Search of Bases for Ethics* (Oxford University Press 2008) for further discussion by Gibbard of his preferred way of developing this style of account, and also Howard Nye, 'Ethics, Fitting Attitudes, and Practical Reason: A Theory of Normative Facts' (PhD Dissertation, University of Michigan 2009) for connected discussion.

The basic issue is that the more extension-revealing a concept is, the *less clear* it is that this concept helps reveal the facts that (one might reasonably think) ultimately matter the most in the context of ethical inquiry: facts about *what we really and truly should do*. For example, take the concept BRAVE, which is arguably fairly extension-revealing. How normatively important are facts about bravery? More specifically: how much do facts about who is *brave* or not (or which acts are brave or not, etc.) matter in terms of settling what agents really and truly should do? Ordinary competent users of the concept BRAVE aren't in a strong position to answer that question without significant further normative inquiry. Basic competence with the concept (in combination with the kind of basic knowledge about the way the world is that ordinary competent users of the concept have) doesn't put them in a strong position to answer that question. Contrast that with the concept ALL-THINGS-CONSIDERED WHAT YOU AUTHORITATIVELY OUGHT TO DO. That concept certainly seems more authority-revealing. It's hard to deny that the facts in the extension of this concept (which, on a straightforward way of thinking about it, will be the facts about what you all-things-considered authoritatively ought to do) are *very* normatively significant. If one denies that, it seems that one just hasn't really grasped the concept at issue, or is incompetent at using it. But one might be quite a competent user of that concept and not know much about what, in fact, someone (either herself or another) all-things-considered really ought to do.

This tradeoff between a normative concept's being more extension-revealing vs. more authority-revealing might just be apparent. Perhaps we have (or at least *could* have) concepts that are *very* revealing in both ways. Perhaps, for example, some of our moral concepts (e.g., MORAL WRONGNESS or MORAL REASON) are such concepts. That, however, is far from clear. And, indeed, at least prima facie, there are important reasons to doubt it. Consider, for example, the seeming conceptual coherence of forms of moral skepticism that claim that morality isn't that normatively important.[51]

A further problem is that even if there are such concepts that are both highly extension-revealing and authority-revealing, then they are ones we should be wary about using in normative inquiry, at least when we are engaged in all-things-considered normative theorising about what we really and truly should do. After all, we want to be able to normatively question whether, for example, facts about bravery *really are* normatively important or not, as well as *how* normatively important they are. And we want to be able to – and indeed seem to have the ability to – ask such questions about the normative import of *many* particular facts. So concepts that bring together *too closely* these two different ways of being revealing may be ones we want to resist relying on too much in our thinking.[52] They might also well be ones that are more likely to yield empty extensions, or prevent us from seeing which things *really are* authoritatively normative for us. After all, one could, of course, stipulate that there is a concept that picks out facts that are both (a) maximally normatively authoritative and (b) are thus-and-so (e.g., God has commanded them, or they involve promoting happiness, or they are about what people would contractually agree to under certain idealised conditions). But reality might not cooperate.

[51]For example, consider the critiques of morality in Nietzsche (n 46), Williams (n 46), and Anscombe (n 46).
[52]For connected discussion, see Simon Blackburn, *Ruling Passions* (Clarendon Press 1998) and Christine M Korsgaard, *The Sources of Normativity* (Cambridge University Press 1996).

Perhaps facts that are thus-and-so are *not* authoritatively normative. Just because concepts represent things one way doesn't mean things actually *are* that way.[53]

So where does this leave the proponent of the conceptual strategy? In order to assess this question, we can start with the following observation: part of why many (perhaps most?) people care about morality in this context – the context of arguments such as the reliability challenge – is that they think morality *normatively matters* in significant ways. The idea, in short, is that the norms of morality (and what they require) matter in a way that standards of etiquette do not, or that the laws of unjust states do not, or that the rules of chess do not. One way this is sometimes put is that morality provides genuine *reasons for action* in a way that these other systems of rules, standards, etc. do not. Reasons for action are, to use the standard gloss, considerations that "count in favor" of certain actions as opposed to others. And not just count in favour of what one *morally* should do, or what one should do "from the point of view of morality". But count in favour of what one *really and truly* should do. Or, put another way, they count in favour of what one *authoritatively* should do.

According to *morality/reasons internalism* (or "moral rationalism"), the connection here is one of metaphysical necessity. Roughly, the idea is that it's part of the nature of morality as such that it necessarily helps determine what an agent authoritatively should do.[54] Perhaps that sort of connection is needed to capture what some might think of as the "objective bindingness" of morality.[55] That will obviously depend in large part on what one has in mind with "objective bindingness". In any case, we can sidestep that issue and simply note the following: one doesn't need to think that the connection is *that* strong in order to think that morality matters in normatively significant ways for us. We might think that it is metaphysically contingent that morality matters in significant ways, but still think that it *does* so matter for creatures like us in the actual world (and perhaps in *many* worlds that are like ours in relevant respects).[56] The core issue then, isn't about moral rationalism. It's about the normative import of morality, however that is (or is not) established.[57]

If one doesn't think morality matters in this way, one might of course still be interested in the above reliability challenge to morality. But this argument about morality (as well as a range of associated arguments about moral epistemology) would then lose much of its pre-

[53]For connected discussion here in the context of ethics and other parts of normative inquiry, see McPherson and Plunkett, 'Conceptual Ethics and The Methodology of Normative Inquiry' (n 7). For connected discussion of this point in general, see Jackson (n 34). It should be noted that it is of course the case that for certain kinds of facts (e.g., certain kinds of judgment-dependent or response-dependent ones), facts about how we (or other agents) represent things (using certain concepts) might help ground those relevant facts. But, in such cases, that is because of what those relevant facts are like, and not because of the nature of concepts as such. In short, for such facts, there is still the conceptual distinction between issues about how we (or other agents) think about reality, and issues about reality itself. In any case, for the non-naturalistic realist, the relevant facts are *not* meant to be response-dependent or judgment-dependent in any way.

[54]See Stephen L Darwall, 'Reasons, Motives, and the Demands of Morality: An Introduction' in Stephen L Darwall, Allan Gibbard and Peter Albert Railton (eds), *Moral Discourse and Practice: Some Philosophical Approaches* (Oxford University Press 1997) for an overview of "morality/reasons internalism" along the lines formulated here.

[55]Street's discussion of morality/reasons internalism in Street, 'Reply to Copp' (n 2) suggests that she endorses this idea.

[56]For this kind of view, consider the case of Peter Railton, who both affirms morality/reasons externalism and the normative import of morality. See Peter Railton, *Facts, Values, and Norms: Essays Toward a Morality of Consequence* (Cambridge University Press 2003).

[57]My discussion here thus differs from Street's focus on morality/reasons internalism vs. externalism as the key dividing line in her reply to Copp in Street, 'Reply to Copp' (n 2). However, I take the core line I am arguing for in this paragraph to be the same underlying point Street is making there (using the morality/reasons internalism vs. externalism distinction). This is that a crucial question for those (such as Copp) who want to focus on morality in this dialectical context concerns the *import of morality* for determining authoritatively normative facts (e.g., what an agent authoritatively ought to do).

theoretic interest for many of us. With that in mind, then, there is good reason for the proponent of the conceptual strategy to vindicate the idea that the relevant moral concepts are fairly authority-revealing concepts, such that it is clear that the things being discussed – e.g., moral wrongness, morality, etc. – are things that really normatively matter. But the more one thinks that the relevant moral concepts are extension-revealing, the harder it is to vindicate this thought.[58]

Before moving on, I want to raise a final issue connected to the tradeoff between the extension-revealing and authority-revealing aspects of normative concepts. The conceptual strategy aims to explain how *we* are reliable in our moral judgments. It does so by appealing to (purported) conceptual truths about relevant moral concepts that we actually currently employ. But what if not all of us share the relevant moral concepts? Conceptual variation among people certainly seems possible here. Consider the wide range of different theories that philosophers have offered about what is distinctive about morality as such. Some associate morality with (at least purportedly) "objectively binding", attitude-independent, categorical norms,[59] others with a distinctively impersonal point of view,[60] and others with certain ideas about fairness, or our ability to justify our actions to others.[61] Additionally, as we saw earlier, some associate it with the fittingness of certain emotions. Each of these ideas resonates with certain strands of thinking about morality in our current social/historical context. Perhaps, then, this is some evidence that people in our social/historical context in fact employ a range of different moral concepts. Depending on how one thinks we should categorize properly "moral" concepts, we might put this more precisely as follows: people employ a range of "moral-ish" concepts.[62] If there is this kind of conceptual variation, then it presents a challenge for the conceptual strategy. For if not all of us are guided in our moral thinking by the same concepts, then it's not clear how the (purported) conceptual truths about some of these concepts will explain how *all* of us are reliable in our moral judgments.

This issue about possible conceptual variation connects back to the tradeoff between the extension-revealing and authority-revealing aspects of normative concepts as follows. Suppose some people employ moral concepts that are more extension-revealing (but not very authority-revealing), but others employ moral concepts that are the opposite of that. We then need to make sure that, when consulting our intuitions about the

[58]It should be noted here that Cuneo and Shafer-Landau explicitly claim that they are (for the purposes of their paper in which they argue for a version of the conceptual strategy) neutral as how to vindicate the normative authority of morality (which they aim to do, as do most non-naturalistic realists about morality). See Cuneo and Shafer-Landau (n 5) 406–07. This explicitly leaves it open how someone who adopts their version of the conceptual strategy should aim to establish the normative authority of morality. Perhaps she should aim to do so through a conceptual link between moral concepts and authoritatively normative concepts, which is what I am focused on here. Or perhaps she should instead aim to do so through some other, non-conceptual way. On such a route, it could be *conceptually* entirely possible that morality is not normatively authoritative at all (or at least not *very* normatively authoritative), but that morality's normative authority could still be discovered through substantive normative inquiry. I return to this latter option toward the end of the paper.

[59]See, for example, Williams (n 46).

[60]See, for example, Henry Sidgwick, *The Methods of Ethics* (Hackett Pub. Co. 1874/1981) and Railton (n 56).

[61]See, for example, T.M. Scanlon, *What We Owe To Each Other* (Harvard University Press 1998) vol and Stephen L Darwall, *The Second-person Standpoint: Morality, Respect, and Accountability* (Harvard University Press 2006).

[62]The possibility of conceptual variation here is especially pressing if one thinks (as I do) that a number of prominent arguments against such conceptual variation are overstated. For example, one prominent kind of argument against such variation is that it makes it too difficult (or impossible) to explain how moral disagreement is possible in cases involving the relevant kinds of conceptual variation. But those arguments face some serious problems, and often don't work. Or at least so I argue. See Plunkett and Sundell, 'Disagreement and the Semantics of Normative and Evaluative Terms' (n 15).

nature of morality, or when making arguments about morality, we don't switch from one moral concept to another. If we did so tacitly, and were unaware of doing so, then doing so might give us the false impression that we were employing a single relevant moral concept that is at once *both* highly authority-revealing and highly extension-revealing. So, in assessing the prospects of a given proposed analysis of a relevant moral concept (a concept that will be deployed in the conceptual strategy), we need to make sure that analysis is really about a *single* given concept, rather than an analysis derived from tacitly switching from theorising about one moral concept to another.[63]

6. Explaining our reliability about authoritative normative facts

In the last section, I argued that there is a tradeoff between the extension-revealing and authority-revealing features of normative concepts, and that this is a problem for the proponent of the conceptual strategy. Now, to press this point further, I focus directly on the most authoritative normative facts in ethics, rather than on moral facts.

Consider the following case.[64] Suppose Olivia has been convinced that she is morally required to donate large sums of her money (much more than she does now) to charity. But she also isn't sure about how important morality is, relative to her own well-being. She wonders to herself: "what should I do?". Suppose she asks others for advice. It won't be very helpful for someone to tell her "well, morally you should donate the money, but prudentially you should keep it". She already thinks that. She wants to know what she should do in light of this perceived conflict. In answering her question, we can also easily imagine it won't settle things to tell her "well, according to the code of conduct X, you should donate the money". She isn't interested in just having *some* further normative system take sides here in the (at least perceived) conflict between prudence and morality. Rather, she wants to know what she *really and truly* should do. Or, as Street puts it about an essentially parallel case, Olivia wants to know what she should do *period*, or what she should do *simpliciter*.[65] In asking about this, Olivia seems to aim to deploy a normative concept that wears its normative authority on its sleeve. Call this concept AUTHORITATIVE PRACTICAL OUGHT.[66]

Perhaps Olivia is wrong in thinking that it is even possible for prudence and morality to conflict. Or perhaps she is mistaken to think that moral considerations could even possibly be overridden by other kinds of normative considerations. Even if so, the *very idea* of there being an *authoritative practical ought* here – something that is distinct from just whatever morality prescribes – doesn't seem *conceptually* mistaken. It seems, rather, that Olivia is asking a perfectly intelligible question. Moreover, it is the kind of question that much of normative ethics (as well as much of normative political philosophy) ultimately seeks to answer. To appreciate

[63]For connected discussion, see McPherson's and my discussion of the "unreliable inference" danger (one species of which involves what we call "implicit switching") in McPherson and Plunkett, 'Conceptual Ethics and The Methodology of Normative Inquiry' (n 7).

[64]The below case closely parallels the case discussed in McPherson, 'Authoritatively Normative Concepts' (n 9), as well as one discussed in Street, 'Reply to Copp' (n 2).

[65]Street, 'Reply to Copp' (n 2) 221.

[66]My gloss of the basic job description of authoritatively normative concepts – and AUTHORITATIVE PRACTICAL OUGHT in particular – draws from McPherson, 'Authoritatively Normative Concepts' (n 9), McPherson and Plunkett, 'The Nature and Explanatory Ambitions of Metaethics' (n 15), and Plunkett and Shapiro (n 15). My use of the concept AUTHORITATIVE PRACTICAL OUGHT for the particular concept at hand draws from McPherson, 'Authoritatively Normative Concepts' (n 9) who puts it in terms of the concept PRACTICAL OUGHT.

this, consider a core contrast here between the idea of authoritative (or "robust") normative facts and the idea of generic (or "formal") ones. Roughly, formally normative facts are facts about how things (e.g., actions) stand in relation to standards of any kind, ranging from standards of etiquette to the rules of a board game.[67] In many areas of philosophy in which we engage in normative inquiry (e.g., ethics, political philosophy, and epistemology), we don't take ourselves just to be interested in questions about merely formal normativity (e.g., whether things conform to the rules of a board game). Rather, we take ourselves to be interested in questions about how we *really and truly* should live, or how our social/political institutions *really and truly* should be run, or what laws we *really and truly* should have.

There are, of course, challenges here. Perhaps the concept AUTHORITATIVE PRACTICAL OUGHT is defective in some way. Perhaps it is incoherent, such that there really is no concept here that could even possibly have an extension – or perhaps there is not even a concept here that we do (or even could) possess.[68] Or perhaps the concept is fine, but it just turns out that there are no authoritative normative facts about what agents should do.[69] These are important challenges. But, for our purposes here, we can put them aside. Instead, we can focus on the question: *if* we suppose there is a coherent concept here, what more (if anything) can we say about it? In particular, how extension-revealing is it?

The basic answer is that this normative concept looks like it is *not* very extension-revealing at all. There are different possible explanations for why this is so (assuming it is). For example, one possible explanation would be this: the concept is ultimately primitive and unanalysable.[70] But, importantly, not all explanations need take that strong line. On this front, consider Tristram McPherson's proposed analysis of AUTHORITATIVE PRACTICAL OUGHT (which he calls PRACTICAL OUGHT):

> McPherson's Account: The definition of 1) S practically ought to do A in context of action C is 2) the constitutive success conditions for S's activity of non-arbitrary selection concerning C require doing A.[71]

McPherson's idea, in short, is that we can get a grip on what it would be to *non-arbitrarily* select among various options (e.g., Olivia selecting to donate her money or not). We can then think that there are constitutive standards that govern this activity of non-arbitrarily selecting, which include constitutive success conditions for succeeding in this activity. (Compare: it might be that the constitutive success condition of playing chess is to win by checkmating one's opponent). McPherson thinks we can then use those success

[67] See McPherson, 'Authoritatively Normative Concepts' (n 9), McPherson and Plunkett, 'The Nature and Explanatory Ambitions of Metaethics' (n 15), Plunkett and Shapiro (n 15), David Copp, 'Moral Naturalism and Three Grades of Normativity' in Peter Schaber (ed), *Normativity and Naturalism* (Ontos-Verlag 2005), David Enoch, 'Is General Jurisprudence Interesting?' in David Plunkett, Scott Shapiro and Kevin Toh (eds), *Dimensions of Normativity: New Essays on Metaethics and Jurisprudence* (Oxford University Press 2019), and Plunkett, 'Robust Normativity, Morality, and Legal Positivism' (n 9).

[68] See Derek Baker, 'Skepticism about Ought Simpliciter' in Russ Shafer-Landau (ed), *Oxford Studies in Metaethics*, vol 13 (Oxford University Press 2018). Parts of Evan Tiffany, 'Deflationary Normative Pluralism' (2007) 37 Canadian Journal of Philosophy 231 can also be read in this vein.

[69] See David Copp, 'The Ring of Gyges: Overridingness and The Unity of Reason' (1997) 14 Social Philosophy and Policy 86 and Tiffany (n 68).

[70] Street expresses sympathy for this idea about at least some authoritatively normative concepts. See Sharon Street, 'What is Constructivism in Ethics and Metaethics?' (2010) 5 Philosophy Compass 363 and Sharon Street, 'Constructivism about Reasons' in Russ Shafer-Landau (ed), *Oxford Studies in Metaethics: Volume 3*, vol 3 (Oxford University Press 2008).

[71] McPherson, 'Authoritatively Normative Concepts' (n 9). I slightly modify the presentation of this definition for ease of presentation (in short, to avoid using McPherson's shorthand label of '=def' as a stand in for the "definition" relation).

conditions to give a definition of AUTHORITATIVE PRACTICAL OUGHT, as glossed above in his proposed definition.

McPherson's account of AUTHORITATIVE PRACTICAL OUGHT might well be a correct (and informative) analysis. But notice how little that analysis helps us with knowing which particular acts (if any) are the ones agents authoritatively ought to do. We don't know, for instance, *what* the constitutive standards of success are for this activity of "non-arbitrarily selecting". These standards need to be discovered through further inquiry (which will no doubt be highly controversial). If McPherson is right, we can't discover those standards just by analyzing the concept AUTHORITATIVE PRACTICAL OUGHT. Thus, if McPherson's analysis is right, it would be entirely unsurprising that agents who are competent at deploying the concept AUTHORITATIVE PRACTICAL OUGHT (whether they have grasped the correct analysis of it or not) find themselves confused or at a loss about which facts are authoritatively normative and which ones are not, given just their knowledge of the conceptual truths about this concept and of basic descriptive facts about their world. The concept AUTHORITATIVE PRACTICAL OUGHT is not a very extension-revealing concept. Perhaps more importantly in the context of discussing the conceptual strategy, we can also note the following: even if it turns out that this concept is fairly extension-revealing for certain creatures in certain contexts, it won't be *because of* the conceptual truths about that concept. Rather, it would be because of something such as how those creatures had mastered deploying this concept, in connection with mastery of relevant non-conceptual information. The concept itself simply doesn't have the right sort of conceptual content to guarantee that it would be extension-revealing. And that is what the proponent of the conceptual strategy needs, given her appeal to *conceptual truths*.

To emphasise the way in which this concept is not extension-revealing, consider that it's not just that the concept doesn't tell us what the constitutive standards of success are for this activity of "non-arbitrarily selecting". Notice also that, if McPherson's analysis is right, the mere concept AUTHORITATIVE PRACTICAL OUGHT itself doesn't tell us much about *how* perceived practical conflicts are to be settled, let alone what the end results will be, in order to count as falling under the concept. Perhaps morality trumps prudence. Or perhaps a Nietzsche-inspired critique of morality is right and morality is less normatively important than many people think.[72] Or perhaps what matters is something that doesn't make reference to the final dictates of morality or prudence at all. All of these are conceptually coherent options, according to McPherson's account of AUTHORITATIVE PRACTICAL OUGHT.

At this juncture, it is worth considering the following point. Someone with sufficient mastery of AUTHORITATIVE PRACTICAL OUGHT might well, when faced with a given scenario or thought experiment, be in a strong position to quickly and easily identify the extension of that concept. For example, when faced with the question of whether to get her shoes muddy in order to save the life of a drowning child, an agent might quickly grasp not only what morality requires of her, but also what she authoritatively should do: namely, save the child.[73] Given that, one might wonder: isn't AUTHORITATIVE PRACTICAL OUGHT fairly extension-revealing, after all?

[72]See Nietzsche (n 46).

[73]This case of the drowning child is taken from Peter Singer, 'Famine, Affluence, and Morality' (1972) 1 Philosophy and Public Affairs 229.

There are two points worth making about this. First, as I have emphasised, extension-revealingness is gradational, along multiple dimensions. The contrast that matters in this argumentative context is about where things stand on a continuum. Second, the basic contrast I want to draw remains, even with this case in mind. Consider how you would react when confronted with someone who seemed to think you should not save the drowning child. Her view here seems conceptually coherent (even if badly mistaken). Her view might well *not* seem conceptually coherent if the disagreement instead involved more extension-revealing concepts applied to the case at hand; e.g., ones involving what would be polite to do, what would be brave to do, or (perhaps) what one is morally required to do. That contrast between how conceptually "open" these different normative issues seem is tied to how extension-revealing the different normative concepts are.[74] Moreover, in the current context, it is important to emphasise the following point (underscored by McPherson's account of AUTHORITATIVE PRACTICAL OUGHT): whatever explains your (purportedly) intuitive grasp of the extension of AUTHORITATIVE PRACTICAL OUGHT in the drowning child case and other similar cases, it doesn't seem to be conceptual truths about that concept itself. The parallel point applies to other authoritatively normative concepts, e.g., the concept (call it GENUINE NORMATIVE REASON) about those considerations that directly "count in favor" of what an agent *authoritatively* ought to do (rather than, for example, in favour of what she *legally* or *morally* ought to do). This is a problem for the proponent of the conceptual strategy. This is because she is claiming that the relevant conceptual truths are doing heavy-lifting in explaining our grasp of the extension(s) of the relevant concept(s).

With this in mind, now return to the reliability challenge. Suppose we run this challenge but focus it *not* on morality but on authoritative normativity. In that case, the parallel conceptual strategy isn't going to work. Given the nature of the authoritatively normative concepts involved, this is a doomed strategy. As McPherson's analysis of AUTHORITATIVE PRACTICAL OUGHT suggests, there might well be important conceptual constraints built into the application-conditions of these concepts. But they aren't constraints of the kind that will be useful to the proponent of the conceptual strategy. In sum: authoritatively normative concepts don't have the kind of constraints built into them that would make them extension-revealing.[75] Importantly, this is not just true of "all-things-considered" notions, such as AUTHORITATIVE PRACTICAL OUGHT. It is also true of the relevant *contributory* notions that are conceptually tied to it, such as concepts along the lines of AUTHORITATIVE NORMATIVE REASON or AUTHORITATIVE VALUE (which are the concepts I take

[74]The sort of "open-question"-esque feel of some normative issues can be seen as part of what motivates G.E. Moore's open-question argument in G.E. Moore, *Principia Ethica* (Cambridge University Press 1903/1993), as well as more recent arguments that draw on Moore (such as Gibbard's use of "what's at issue" tests) in Allan Gibbard, *Thinking How to Live* (Harvard University Press 2003). One can think (as I do) that Moore's open-question argument is a bad argument as originally deployed, and that it doesn't support non-naturalistic realism (as he thought), but still think his argument was sensitive to important aspects of some normative concepts.

[75]I suggested this basic line of argument about the conceptual strategy in communication to Josh Schechter about a draft of his paper Schechter (n 1). He agreed, and put the point here as follows in the final published version of his paper: "It is implausible that there are significant conceptual constraints built into our most general normative concepts, so the analogous strategy [one analogous to the conceptual strategy] would not seem to work for explaining our reliability about beliefs involving these concepts." Ibid. Unsurprisingly, I think that Schechter's way of putting this is on the right track. Part of my aim in this paper is to explain more about *why* a version of the conceptual strategy won't work for authoritatively normative concepts (which I here take Schechter to be getting at with his discussion of the "most general" normative concepts). Part of what I am arguing is that it's not just that there aren't significant conceptual constraints here in these concepts. It's that, insofar as there are constraints, they are going to be the wrong kind of constraints, given the challenge. In short, they won't be constraints of a kind that would make the concepts extension-revealing.

some to express when they talk about "genuine" normative reasons or "genuine" value). We might think of these concepts as forming a family of authoritatively normative concepts.[76] And none of them are extension-revealing in the way that the proponent of the conceptual strategy needs.

This helps bring out the following. First, the case of fully authoritatively normative concepts shows a sort of limit case for the philosopher focused on morality. In short, the more you express concepts like these with moral terms (e.g., 'morality'), and then focus on those concepts as the relevant moral concepts for use in the conceptual strategy, the less likely it is that the conceptual strategy will work. Second, if we suppose (as seems reasonable) that we are (or at least should be) ultimately after the authoritative normative facts when engaged in ethical inquiry, it is not at all clear how helpful answering the reliability challenge about specifically moral facts really is going to be. What we really need is an answer to the reliability challenge when it is run about authoritatively normative facts as such. And the conceptual strategy won't work there.

At this juncture, it is important to emphasise the following point. Much of the critical discussion of the reliability challenge, as well as the discussion of other arguments tied to Street's Darwinian Dilemma argument, has been about morality, or other specifically *moral* topics. But a number of people have pressed the idea that these arguments are best run about authoritatively normative facts (although without using that exact terminology). Indeed, as I underscored at the start of this paper, Street is one of them. And her reasons for *why* resonate with what I have been arguing. First, she (like me) thinks these are the facts that ultimately matter the most in ethics.[77] Second, she also, in effect, points to the fact that authoritatively normative concepts *lack* conceptual constraints of a certain kind. She writes:

> According to the normative realist, there are normative truths that hold independently of all of our evaluative attitudes. Moreover, as a purely conceptual matter, these independent normative truths might be anything. In other words, for all our bare normative concepts tell us, survival might be bad, our children's lives might be worthless, and the fact that someone has helped us might be a reason to hurt that person in return. Of course we think that these claims are false—perhaps even necessarily false—but the point is that if they are false, it is not our bare normative concepts that tell us so.[78]

Street is talking about "bare" normative concepts here. Someone might use that phrase to talk about a normative concept (e.g., OUGHT or REASON) as it could then be fleshed out in any number of different ways: for example, in the context of talking about what we legally should do, prudentially should do, etc.[79] But it is clear that Street is *not* talking about concepts that are "bare" in that way. Rather, she is talking about what I am calling "authoritatively normative concepts".[80]

[76]See McPherson, 'Authoritatively Normative Concepts' (n 9).
[77]See Street, 'Reply to Copp' (n 2).
[78]Ibid 208.
[79]For example, think of a Kratzer-style account of deontic modals as giving us the "bare" account of the relevant modals. See Angelika Kratzer, *Modals and Conditionals: New and Revised Perspectives* (Oxford University Press 2012). See Jennifer Carr, 'Deontic Modals' in Tristram McPherson and David Plunkett (eds), *The Routledge Handbook of Metaethics* (2017) for a helpful overview.
[80]Among other things, consider how Street uses the terms 'value', 'reason for action', and 'normative' in the relevant papers (as in Street, 'A Darwinian Dilemma for Realist Theories of Value' (n 2), Street, 'Constructivism about Reasons' (n 70), Street, 'Reply to Copp' (n 2), Street, 'Mind-Independence Without the Mystery' (n 2), Sharon Street, 'Coming to Terms with Contingency: Humean Constructivism about Practical Reason' in Jimmy Lenman and Yonatan Shemmer (eds),

Street is right to insist that her focus on authoritatively normative concepts – rather than on *moral* concepts – matters in this kind of argumentative context, in discussing epistemic challenges of the kind I have been discussing. And it thus means that when philosophers *switch* the arguments from being about authoritative normativity to morality in particular, they are changing the arguments in significant ways. On this front, consider Cuneo and Shafer-Landau's engagement with Street. Directly after discussing the quote from Street that I discussed above, in which she never mentions specifically moral concepts at all (and is instead focused on authoritatively normative ones), Cuneo and Shafer-Landau switch (without comment) to taking her view to have direct consequences for moral concepts. They write that "[i]f Street is right, then, as a conceptual matter, the reference of our moral concepts is wholly elastic in the sense that while they must pick out moral duties, moral virtues, or moral goods, just about anything could qualify as a moral duty, a moral virtue, or a moral good".[81] Cuneo and Shafer-Landau do this a number of years after Street's reply to Copp, in which Street argues at length for why this kind of switch from talking about authoritative normativity to morality can matter a great deal, and makes clear that her own focus is on authoritative normativity and *not* morality. A switch, such as Cuneo and Shafer-Landau's, to discussing morality instead of authoritative normativity can of course be *motivated*. After all, people (including Cuneo and Shafer-Landau) might be *specifically* interested in morality as such for any number of reasons.[82] My point here, however, is that this is a very consequential shift in this dialectical context. In making this shift, Cuneo and Shafer-Landau make the relevant epistemological argument they are targeting (the one that stems from Street) less powerful, and thereby make the conceptual strategy more promising as a response. But that doesn't change how good of a response it is to the argument stated in terms of authoritative normativity, which is how Street originally stated it (using slightly different terminology).

Importantly, Cuneo and Shafer-Landau are not alone in making this kind of switch from taking Street's style of argument to be one about authoritative normativity to one about morality. Copp also does this, which forms part of the basis for Street's response to him, as I have discussed.[83] Finally, to take one more example, consider Vavova. As I noted at the start of this paper, Vavova argues that a version of the conceptual strategy is a promising way to respond to a kind of "irrelevant influences" argument, which she takes to be the best way of reconstructing Street's Darwinian Dilemma argument. At the start of the paper in which she argues for this claim, Vavova writes the following: "I focus on moral, rather than evaluative realism because it allows for the most familiar and vivid presentation of the problem. All I say applies *mutatis mutandis* to evaluative

Constructivism in Practical Philosophy (Oxford University Press 2012)). Consider also her discussion of what's at issue in the case that essentially parallels the above case of Olivia in her reply to Copp (see Street, 'Reply to Copp' (n 2) 208).

[81] Cuneo and Shafer-Landau (n 5).

[82] Thus, there might be very good reasons why Cuneo and Shafer-Landau are themselves specifically interested in morality in ibid, given their own philosophical aims. Or, to take another example, consider Copp's work in Copp, 'Darwinian Skepticism About Moral Realism' (n 13). This is an interesting case, given Copp's background views about normativity. His views are subtle, and I won't try to reconstruct them here. The basic point is that his views involve a kind of skepticism about the "authoritative ought", as well as independent reasons to be interested in morality as such aside from its (perceived) normative import. See Copp, 'The Ring of Gyges' (n 69) and David Copp, *Morality in a Natural World: Selected Essays in Metaethics* (Cambridge University Press 2007) for further discussion, as well as Street, 'Reply to Copp' (n 2) for discussion of why his views on these issues matter in the context of discussing the Darwinian Dilemma argument.

[83] Street, 'Reply to Copp' (n 2)

realism (cf. Street 2006), which may be broader than moral realism".[84] This makes it sound as though Vavova's choice to focus on morality is innocuous with respect to her assessment of the conceptual strategy, such that this strategy will carry over smoothly to discussion of what Street is ultimately after (namely, authoritative normativity). This, however, is a mistake with respect to the conceptual strategy when run as a reply to the reliability challenge, for the reasons that I have been discussing. And, as I will argue later in this paper, it is also a mistake, for the same kinds of reasons, when a version of the conceptual strategy is used as a reply to the irrelevant influences argument that Vavova focuses on.

The reasons I have given for *why* it matters whether we focus on morality or authoritative normativity in this dialectical context connect in crucial ways to previous arguments from Street. As I understand it, my work in this paper resonates with core aspects of Street's thinking that run throughout her discussion of the Darwinian Dilemma, as well as key parts of her discussions of normativity and morality more generally. But the argument I am making here is not simply a repeat of what Street has already said, updated to illustrate problems with new versions of the conceptual strategy (e.g., Cuneo and Shafer-Landau's version). My response to the conceptual strategy (including the more recent iterations of it that Street hasn't explicitly engaged with) relies on (a) my own preferred way of carving up the relevant terrain about normativity, concepts, etc., (b) the discussion of fitting attitude accounts of relevant moral concepts to illustrate key points, and (c) the idea of a tension between the authority-revealing and extension-revealing aspects of normative concepts. My response also explicitly engages the conceptual strategy as deployed not as a response to the details of Street's Darwinian Dilemma argument, but rather to the more schematic kinds of arguments that Street's work brings out (and which have been the focus of much subsequent work in normative epistemology): including the reliability challenge (which I have been focusing on) and the irrelevant influences argument (which I discuss later). If successful, the result is a way of responding to the conceptual strategy when run either for moral concepts or authoritatively normative concepts, and when run for responding to either a reliability challenge or an irrelevant influences argument. And, if successful, it is a reply that is at once both ecumenical and powerful.

7. Two replies

I now consider and respond to two important possible replies to my argument about the conceptual strategy.

7.1. A non-conceptual way of vindicating the tie to authoritative normativity?

Suppose for the sake of argument that the conceptual strategy works for the reliability challenge when run about *morality*. But also suppose that I am right that a parallel strategy won't work for the parallel argument about authoritative normativity. In such a situation, the proponent of the conceptual strategy might then seek a way for the conceptual strategy to still do *some* work. In particular, perhaps one could establish – not through conceptual analysis, but through substantive normative argument – that morality normatively matters

[84]Vavova, 'Evolutionary Debunking of Moral Realism' (n 14) (2) 114.

in significant ways for settling what we authoritatively should do. Then perhaps the conceptual strategy is useful after all, even if we ultimately are targeting an epistemological issue about authoritative normativity, and not morality.

I think this is a possible response worth further exploring. But it faces three significant issues.

First, if this response is correct, it undercuts much of the seeming force and interest of the conceptual strategy. In short, if this response is correct, then what would turn out to be really doing a lot of the work (in terms of addressing the relevant epistemological issues about authoritative normativity) is the (purported) fact that morality is normatively significant in such-and-such ways. On the line of thought being considered here, that morality has such normative significance is not a conceptual truth. So the conceptual truths about MORALITY (or other related moral concepts) aren't doing the serious explanatory work here. If conceptual truths are peripheral to the (purported) explanation of our reliability about authoritative normative facts, then it's not clear how much the conceptual strategy as such is adding to the discussion of the reliability challenge.

Second, if the (purported) connection between morality and authoritative normativity is what really matters, we can (and should) just run the reliability challenge when considering the (purported) normative fact that there is a connection between morality and authoritative normativity. In other words: if we think we are reliable in our judgment that morality (as we conceive of it) *normatively matters* in the ways we think it does (roughly, that it is *normatively privileged* with respect to a range of other actual and possible normative systems), then we should ask what explains that reliability.[85] It is not at all clear what a good response for non-naturalistic realists is to this challenge. More to the point: whatever their response is, it's not going to be conceptual truths about MORALITY (or other related moral concepts) that are the basis for such a response. This is because, on the line of response we are considering here, the whole idea is that we establish this connection via a non-conceptual route.

Third, and tied to the previous point, consider that the (purported) connection between morality and authoritative normativity is supposed to be established by *normative* argument. If we lack a good response to the reliability challenge when run about authoritative normativity, it is hard to see why we should be confident in using *normative argument* about *authoritatively normative facts* to help us block the challenge on the table. Perhaps we are entitled to use *some* such normative argument of this kind, even though it is the sort of thing whose reliability we want to explain.[86] But there is at least prima

[85]This idea ties into the second kind of worry that Schechter raises for the conceptual strategy in Schechter (n 1) 454, about maintaining a commitment to what he thinks of as a "non-plenitudinous" form of non-naturalistic realism. He puts the issue there in terms of moral properties being "metaphysically privileged". I think that the issue is better put in terms of "normative" privilege, since I think that normative privilege might come apart from metaphysical privilege, for the kinds of reasons given in Eklund (n 11)).

[86]Whether or not this is so is a crucial point of debate in discussion of a range of arguments about normative epistemology, including the reliability challenge, and other arguments thought to be connected to Street's Darwinian Dilemma argument. For some of the relevant discussion here, see Katia Vavova, 'Irrelevant Influences' (2018) Philosophy and Phenomenological Research 134, Vavova, 'Evolutionary Debunking of Moral Realism' (n 14), Enoch, 'The Epistemological Challenge to Metanormative Realism' (n 1), Berker (n 3), Street, 'Reply to Copp' (n 2), and Dustin Locke, 'Darwinian Normative Skepticism' in Michael Bergmann and Patrick Kain (eds), *Challenges to Moral and Religious Belief: Disagreement and Evolution* (Oxford University Press 2014). This debate ties into more general discussions about what sort of "independence" is needed from one's current views on a subject matter, in cases of assessing higher-order evidence that your current assessment of evidence (relevant to the subject matter at hand) might be mistaken. This has been a crucial

facie reason to be concerned about relying too heavily on such arguments in this context. Given the overall dialectic, the more important point is this: *if* we are entitled to rely heavily on such an argument, then, again, it is not clear how much the conceptual strategy as such is really adding beyond the resources we already then have at our disposal. For example: if we are entitled to such an argument, why would one then not just run a version of a third-factor response?[87]

7.2. A specific theory of concepts?

I have discussed the conceptual strategy in schematic terms. But perhaps what matters are the details of how the strategy is developed. In particular, perhaps what matters is a background theory of concepts (and an associated view of conceptual truths) that is used in developing the conceptual strategy. One reason to take this possibility seriously stems from Cuneo and Shafer-Landau. They present what I am here calling "the conceptual strategy" as one application of a specific kind of moral non-naturalism that they develop.[88] Their version of moral non-naturalism relies heavily on a certain kind of (neo-Fregean) theory of concepts and conceptual truths. Moreover, they never discuss the conceptual strategy as a broad, schematic type of strategy, as I do in this paper. So perhaps the key to making the conceptual strategy work is a specific background theory of concepts and conceptual truths.

With that in mind, let's consider the view of concepts that Cuneo and Shafer-Landau defend. They start with three core claims:

(1) "Concepts are abstract, sharable, mind-independent ways of thinking about objects and their properties. As such, they are very much objective, 'out there' sorts of things, extra-mental items".[89]
(2) "Concepts are the building blocks or sub-components of propositions, which are themselves the objects of belief and other propositional attitudes".[90]
(3) "Concepts are not merely the subcomponents of propositions; they are also referential devices or ways of getting things in mind that enable thinkers to refer to things such as objects and properties".[91]

With this understanding of concepts in hand, Cuneo and Shafer-Landau then introduce the idea of *conceptual truths*. They state that "a proposition <that *x* is *F*> is a conceptual truth if it belongs to the essence of '*F*' that, necessarily, anything that satisfies '*x*' also

part of the debate over how to respond to cases of "peer disagreement", as in the discussion in David Christensen, 'Disagreement, Question-Begging, and Epistemic Self-Criticism' (2011) 11 Philosophers' Imprint and Thomas Kelly, 'Disagreement and the Burdens of Judgment' in David Phiroze Christensen and Jennifer Lackey (eds), *The Epistemology of Disagreement: New Essays* (Oxford University Press 2013).

[87] In answer to this specific question, one idea would be that the conceptual truths give us justification to make certain starting assumptions we otherwise would not be entitled to. See Vavova, 'Evolutionary Debunking of Moral Realism' (n 14) for sympathetic discussion on this point. The core issue with this idea is that it's far from clear why the truths being *conceptual* ones gives them a different epistemic status with respect to the relevant issues about independence, higher-order evidence, etc. that are at the heart of the debates over the third-factor response.

[88] Cuneo and Shafer-Landau (n 5).

[89] Ibid 409.

[90] Ibid 409.

[91] Ibid 410.

satisfies '*F*'".[92] On their view, conceptual truths "are true propositions; specifically, those that hold in virtue of the essences of their constituent concepts".[93]

According to Cuneo and Shafer-Landau, this way of understanding conceptual truths has an important consequence. In their own words: "conceptual truths do not require facts that correspond to them, worldly truth-makers that render them true. Whether such propositions are true depends on the natures of their constituent concepts and not solely on the existence of any correlative worldly facts that can serve as truthmakers".[94] Now consider how Cuneo and Shafer-Landau implement the conceptual strategy. First, start with the thesis (that they endorse) that moral concepts have substantive conceptual constraints that describe which acts are morally right and wrong in fairly direct terms (e.g., that killing for fun is morally bad, for creatures like us, in worlds like ours). This idea is central to the conceptual strategy as I have been considering it in this paper. Second, combine this idea with (a) Cuneo and Shafer-Landau's theory about conceptual truths and (b) the further thesis, which I won't discuss in detail here, that some conceptual truths, including the relevant ones about moral concepts, are *non-natural*. This combination yields the thesis that Cuneo and Shafer-Landau call the *moral fixed points*. As they put it:

> There is a battery of substantive moral propositions, such as <that it is wrong to torture others just because they have inconvenienced you> and <that it is wrong to rape a child solely to indulge one's lust>, that are also nonnaturalistic conceptual truths. These propositions, which we hold must find a place in any system of moral norms that applies to beings like us, in worlds similar to our own, we call the *moral fixed points*.[95]

Spelled out in a bit more detail, using the language of conceptual truths above, the core idea here is as follows:

> There are nonnatural moral truths. These truths include the moral fixed points, which are a species of conceptual truth, as they are propositions that are true in virtue of the essences of their constituent concepts.[96]

On this view, it is not only that the relevant conceptual truths tell us about how we (or other agents) represent reality when we (or they) represent it using those concepts. In addition, these conceptual truths *directly* establish the truth of certain moral propositions, namely, those that they call "the moral fixed points". These moral fixed points include substantive moral principles of the sort I have been discussing.[97]

A natural question to ask is whether the key to successfully pursuing the conceptual strategy involves an appeal to Cuneo and Shafer-Landau's theory of concepts (and of conceptual truths), or else a theory closely related to it in key respects. In thinking about this proposal, we can separate out two questions. One question is whether Cuneo and Shafer-Landau's is the right theory of concepts and conceptual truths. The second question grants for the sake of argument that it is. It then asks whether this fact helps the proponent of the conceptual strategy respond to the issues for that strategy that I have raised in this paper. I am skeptical that it does.

[92]Ibid (89) 410.
[93]Ibid 411.
[94]Ibid 411.
[95]Ibid 400.
[96]Ibid 411–12.
[97]For criticism of Cuneo and Shafer-Landau on these claims, see Stephen Ingram, 'The Moral Fixed Points: Reply to Cuneo and Shafer-Landau' (2015) Journal of Ethics and Social Philosophy 1 and Daan Evers and Bart Streumer, 'Are the Moral Fixed Points Conceptual Truths?' (2016) Journal of Ethics and Social Philosophy 1.

There are a number of distinct issues here. One issue concerns the *kind* of non-naturalistic realism that might be vindicated using this view of concepts. Cuneo and Shafer-Landau explicitly aim to establish a kind of non-naturalism involving these non-natural truths (some of which are the moral fixed points), but not necessarily the more "robust" kind of non-naturalistic realism that I introduced at the start of this paper.[98] One worry is that, in attempting to save "non-naturalistic realism", Cuneo and Shafer-Landau in fact generate a view that is different enough from the view we started with that their defense of it really no longer qualifies as a defense of the kind of view that I have been discussing in this paper, despite their use of the label "non-naturalistic realism". Compare: the issue on the table isn't whether *quasi*-realist expressivists have a good answer to the reliability challenge, but whether non-naturalistic realists do. Of course, it's not clear how deep this worry runs for Cuneo and Shafer-Landau, given their own ambitions. After all, their paper is an attempt to defend a certain kind of view about the nature of morality, and of moral thought and talk. They might succeed at defending that view, even if it is not one that is aptly labelled "non-naturalistic realism" at the end of the day. However, in our context, when considering the prospects of the conceptual strategy for defending non-naturalistic moral (or normative) realism *along the lines it has been standardly understood*, this issue matters.

A second issue is perhaps more worrisome. In short, it is not clear what extra resources one gets for carrying out the conceptual strategy by using the view of concepts Cuneo and Shafer-Landau propose, beyond the resources one already has using other views of concepts, such as the highly schematic view of concepts that I have been relying on in the paper. In particular, it is not clear what extra resources one gets for responding to the issues for the conceptual strategy that I have raised in this paper. Consider a view, for example, on which concepts are, roughly, ways of representing how things might be, such that concepts involve (or consist in) a function from possible worlds to extensions.[99] It seems that most things that one can say using Cuneo and Shafer-Landau's framework could be translated into such a view of concepts, with different parts of the overall picture of cognition and reality described using the label 'concepts'. The key point at this juncture in the dialectic is this: whatever theory of concepts one proposes, we need to see how it gives us particular resources for defending the conceptual strategy against the worries I have raised for it in this paper. It is not clear to me what extra resources Cuneo and Shafer-Landau's view provides.

Moreover, any theory of concepts might well incur additional costs – including some that tie directly into the issues I have been discussing in this paper. On that front, consider that we can run a version of the reliability challenge about our beliefs about the relevant conceptual truths. The more one thinks of conceptual truths as mind-independent, abstract objects – of a sort that look broadly "Platonic" – the more important such a reliability challenge is going to be.[100]

[98]See Cuneo and Shafer-Landau (n 5) 400–401, 417–418, in contrast to the kind of view advocated for in Enoch, *Taking Morality Seriously* (n 1).

[99]For an example of this kind of view of concepts, see Jackson (n 34).

[100]To press this point, consider an analogy. Suppose that one appeals to the Platonic Forms in answering the reliability challenge (either targeting non-naturalistic realism about morality, or authoritative normativity). If the Platonic Forms are themselves non-naturalistic entities, then the same kind of reliability challenge can be put forward with respect to our beliefs about those entities, and with just as much force. For connected discussion (about whether Cuneo and Shafer-Landau's *conceptual* version of moral non-naturalism helps them sidestep traditional metaphysical and

Cuneo and Shafer-Landau anticipate this line of response to their proposal. They claim that there is reason to think that conceptual truths are in a different boat when it comes to the relevant epistemological challenges on the table. In short, they argue that we have a general capacity to recognise some things as conceptual truths, and that this capacity is reliable (at least in core instances). They claim that this ability is rooted in general cognitive abilities and forms of reasoning that we have as thinkers.[101] I agree that this is a line of argument worth exploring at this juncture. But consider the broader philosophical context of the reliability challenge. Even if we grant for the sake of argument that we do have such an ability *and* that we are reliable in that ability, we still should want an *explanation* for why we are so reliable. Cuneo and Shafer-Landau have sketched an explanation, but much more would need to be said. Furthermore, given the specific, broadly "Platonic" view of conceptual truths that Cuneo and Shafer-Landau propose, it is far from obvious that generally plausible accounts of how we learn about conceptual truths (where those truths are understood in other ways) will be plausible here. Until we have a plausible explanation of our reliability in learning about the specific sort of "conceptual truths" that Cuneo and Shafer-Landau appeal to, I think there is good reason to think that a version of the reliability challenge remains unanswered.[102]

The general upshot of this discussion of Cuneo and Shafer-Landau's theory of concepts can be summed up as follows. It might be that a certain view of concepts (and a connected view of conceptual truths) could help bolster the conceptual strategy. That view might be a version of Cuneo and Shafer-Landau's view, or else some other entirely different kind of view. This is a possibility that is surely worth exploring further. For example, a view that might seem particularly promising here is a view such as Ronald Dworkin's, on which conceptual truths (at least about certain kinds of concepts) are partly determined by normative facts. Dworkin argues that this is so for "interpretative concepts" (which he argues our normative concepts are examples of).[103] If such a view is correct, then "conceptual truths" about our normative concepts have a normative significance that they do not have on other rival views about those conceptual truths. It is beyond the scope of this paper to discuss Dworkin's views, or any number of other potential theories of concepts (or connected theories in metasemantics) that one might incorporate into the conceptual strategy.[104] However, there are also general reasons for concern. In particular, a view of

epistemological worries about more standard versions of moral non-naturalism), see David Copp, 'Are There Substantive Moral Conceptual Truths?' in Diego E Machuca (ed), *Moral Skepticism: New Essays* (Routledge 2018).

[101]Cuneo and Shafer-Landau (n 5) 427.

[102]I've just argued that what we need here is a story about the epistemology of conceptual truths that works for those truths *as understood by Cuneo and Shafer-Landau* or, somewhat more modestly, at least a theory of the epistemology of the things that Cuneo and Shafer-Landau refer to when they use the term 'conceptual truth'. Here is another reason why this matters, given the overall dialectic about Street's Darwinian Dilemma argument. Cuneo and Shafer-Landau claim that Street appeals to conceptual truths, and doesn't seem to think a version of her own argument should give her cause for concern in doing so. (See ibid (59) 428). That's a point that matters in this argumentative context. But note that Street's appeal to conceptual truths doesn't involve her committing to conceptual truths *as Cuneo and Shafer-Landau understand them*, as opposed to something more minimal, and more in line with the view of concepts we started with earlier in the paper.

[103]See Ronald Dworkin, *Justice for Hedgehogs* (Harvard University Press 2011), drawing on his earlier discussion in Ronald Dworkin, *Law's Empire* (Belknap Press 1986) and other work.

[104]For critical discussion of Dworkin's theory of concepts, see David Plunkett and Timothy Sundell, 'Dworkin's Interpretivism and the Pragmatics of Legal Disputes' (2013) 19 Legal Theory 242.

concepts (or a view of a connected topic in metasemantics) that would have the resources necessary to help the proponent of the conceptual strategy might either (a) be implausible as a theory of concepts (or whatever it is a theory of) or (b) create a new version of the reliability challenge, which is not easily answered by an appeal to general cognitive capacities. To illustrate the second point, consider that if a theory (e.g., Dworkin's) holds that normative truths ground or constitute conceptual truths, then we will want to know how the view answers the reliability challenge about conceptual truths so understood. Once normative truths show up as grounds of (or as part of) the conceptual truths, the obvious worry is that appeal to conceptual truths won't solve the reliability challenge, but simply relocate the problem. Thus, while an appeal to a specific theory of conceptual truths might in principle help the proponent of the conceptual strategy, these two kinds of concerns would need to be addressed, among others. As far as I can see, we don't yet have a good theory of concepts that does the job of helping the proponents of the conceptual strategy that also successfully addresses these concerns.[105]

Before moving on, I want to note another potential line of argument open to a proponent of Cuneo and Shafer-Landau's theory of concepts. The line of argument is this: perhaps once we fully appreciate the details of their theory of concepts, we will then see that perhaps the schematic version of the conceptual strategy *as I have stated it in this paper* simply is a mistaken way to use their theory of concepts in responding to the reliability challenge, or to a range of connected arguments tied to Street's Darwinian Dilemma argument. Perhaps, instead, there is a better way to make use of their theory of concepts. Indeed, someone might want to insist that Cuneo and Shafer-Landau should never have been read as giving an argument that fits into the schematic version of the conceptual strategy as I have reconstructed it here, and were using it in some other way, which I have not given them credit for. It is obviously beyond the scope of this paper to address the myriad other ways one might want to draw on Cuneo and Shafer-Landau's theory of concepts in this dialectical context.[106] So too is it beyond the scope of this paper to engage in further exegetical arguments about how to best read the details of Cuneo and Shafer-Landau's work that I have been focusing on, and to ask whether my reading of them illuminates a general underlying line of argument that their argument shares with other possible ways of developing the conceptual strategy. With that in mind, we can then note a challenge that my work brings out for future discussion of their work. If there is some other, (purportedly) more effective way to make use of their theory of concepts, the key question will be how it is relevantly different from the conceptual strategy as I have stated it in this paper *such that* the issues I have raised for that strategy (tied to the issue of the seeming tradeoff between how extension-revealing vs. how authority-revealing a normative concept is) don't also apply to this different strategy. Indeed, in light of my discussion thus far, that is a question we should ask *whenever* a philosopher aims to appeal to purported conceptual truths to respond to a reliability challenge, or to any range of connected epistemological arguments in ethics.

[105]An important question for future discussion about the conceptual strategy is whether Cuneo and Shafer-Landau's forthcoming work with John Bengson can deliver the goods here with respect to this challenge. See John Bengson, Terence Cuneo and Russ Shafer-Landau, *Understanding Morality* (Oxford University Press Forthcoming).

[106]See ibid for further discussion.

8. Beyond the reliability challenge in ethics

At the start of this paper, I stated that I would argue that my core argument generalises in three important ways. In particular, I claimed I would argue that it has implications for using a version of the conceptual strategy to reply to (a) forms of the reliability challenge that don't target non-naturalistic realism in particular, (b) parallel reliability challenges that arise in other areas (e.g., about our reliability about mathematical facts, epistemological facts, or legal facts), and (c) many other epistemological challenges in ethics, including others involving evolutionary considerations that some have taken to be tied to Street's Darwinian Dilemma argument. I now turn to making these arguments.

8.1. Other metamoral views and the reliability challenge

Start with the idea of running a version of the reliability challenge against a metamoral view other than non-naturalistic moral realism. If the reliability challenge, as I have formulated it, is on the right track, then *every* metamoral view can gain or lose plausibility based on how well it deals with *some* version of this challenge. In short, insofar as a view is able to provide a good explanation of our reliability in moral judgment (whatever that reliability consists in), then that is a good thing for such a view, and it is a bad thing for such a view if it lacks a good explanation. So, with this in mind, we should ask: would the conceptual strategy fare better if it was used to support a rival metamoral view, e.g., some kind of *naturalistic* realism, or perhaps a form of quasi-realist expressivism, or some form of "quietist" non-naturalistic realism, or any range of anti-realist views? If what I have argued in this paper is on the right track, there is no reason for serious optimism here. Recall the core issues I have pointed to about our moral concepts in this paper: namely, the possibility of an analysis of them on which they are not very extension-revealing (as on a Gibbardian fitting attitude account of them) and the general tradeoff between the extension-revealing and authority-revealing aspects of normative concepts. These are general points about our moral concepts, and appreciating these points causes trouble for *anyone* who wants to appeal to conceptual truths about our moral concepts to explain our reliability about morality. This is true regardless of whether someone is a non-naturalist like Cuneo and Shafer-Landau, or a naturalist like Copp, and regardless of whether she accepts attitude-independence or not. Of course, *some* metamoral views will have a better time here than others in using the conceptual strategy. For example, consider views that involve a *particular* kind of attitude-dependence where the fundamental moral facts depend on facts about *which concepts* we use. So there is room for differentiation here between metamoral views, with respect to how useful a version of the conceptual strategy will be to defending them. Still, the issues I have raised for the conceptual strategy are ones that will affect the discussion of *whichever* kind of metamoral view one chooses to focus on.[107]

[107]To underscore one reason this last point matters, recall that, as I noted earlier in this paper, one way of reading Copp, 'Darwinian Skepticism About Moral Realism' (n 13) is as using a version of the conceptual strategy to defend a form of *naturalistic* moral realism in response to Street's Darwinian Dilemma argument. If I am right, then endorsing naturalism as such does little to help here.

8.2. Reliability challenges about other domains

Second, consider reliability challenges we might run about other domains. As Enoch, Schechter, and others have emphasised, the reliability challenges that I have focused on in this paper (one about morality and the other about authoritative normativity) closely parallel the challenge to mathematical Platonism developed by Harty Field, drawing on work from Paul Benacerraf.[108] And math isn't the only domain where philosophers do (or should) think about reliability challenges. So what then can we say about how my argument in this paper matters for discussion about using the conceptual strategy for responding to reliability challenges in these other domains? The core answer is that it helps bring out questions we should ask about the pursuit of such a strategy in these other domains, and gives us a sense of the potential issues that will arise for such a strategy.

Consider, first, other domains where we engage in normative or evaluative thought and talk – roughly, thought and talk involving such normative concepts as OUGHT, SHOULD, OBLIGATION, etc. or such evaluative concepts as BETTER and WORSE. Normative epistemology is one place we engage in such thought and talk: e.g., in making claims about what agents should believe. Normative political philosophy is another: e.g., in making claims about how social/political institutions should be set up or run. And so too is *law* – where we make claims about what agents legally should do, or what their legal obligations are. There are many ways in which these domains differ from each other, as well as many ways in which one's best theories about these domains do. For example, one might well be a non-naturalistic realist about authoritative normativity but endorse a form of positivism about *legal* facts on which they are ultimately explained by facts about our social practices, and not by authoritative normative facts. If so, given plausible assumptions about what the relevant facts about our social practices are, it's likely not going to be as hard to explain our reliability about the legal facts as it will be to explain our reliability about authoritative normative facts. These kinds of differences might well make the conceptual strategy more or less attractive in these other domains.

My discussion in this paper brings out some questions one should ask in assessing how promising the conceptual strategy is to help explain our reliability in judgment in each respective domain. Here are three crucial ones:

(1) What kind of normative import do we give to the different normative systems, or the reasons, values, etc. that are tied to them? For example: what sort of connection to authoritative normativity do we think the law has? What about etiquette? Or the norms we appeal to when doing epistemology?

(2) How extension-revealing vs. how authority-revealing do we think the relevant concepts (e.g., JUSTICE, EPISTEMIC RATIONALITY, or LEGAL OBLIGATION) are that we might use in running the conceptual strategy?

(3) What kind of range of nearby concepts are there for each concept, and how easily might we be confusing them with the one we target?

[108]See Field (n 1), drawing on Benacerraf (n 1).

In light of my arguments in this paper, we can see that how we answer each of these questions will help give us a sense of whether a version of the conceptual strategy will be useful for answering the reliability challenge in these different areas.

My discussion also helps give us some sense of how promising a version of the conceptual strategy might be for non-normative areas. But the import is less clear here. This is in part because the idea of a concept being *authority-revealing* is tied to specifically normative concepts, and because (in a connected vein) these domains aren't ones where worries about normative authority arise in the same way as they do for ethics (or for epistemology, political philosophy, etc.). Thus, my discussion has fewer direct implications for these non-normative areas: e.g., for using a version of the conceptual strategy to try to explain our reliability in math or logic.[109]

An interesting kind of reliability challenge, which my discussion ties into in important ways, concerns our reliability about a specific normative topic that I briefly brought up earlier in this paper: namely, the issue of which concepts we should use. This topic – which I take to be at the heart of the area of inquiry Alexis Burgess and I call *conceptual ethics* – interacts with the conceptual strategy in an interesting way.[110] Suppose we think (as many of us presumably do) that our normative concepts are generally reliable at helping us learn about the normative facts that (put intuitively) really and truly matter the most. Or, more importantly, even if we don't think all of our actual normative concepts reliably help us in this way, suppose we think that we are generally reliable in our normative judgments about conceptual ethics. That is: suppose we think there is a correlation between the normative facts about which concepts we should use and our beliefs about that topic. What explains the reliability of *those* normative beliefs?[111] Indeed, one might think that explaining *that* reliability should in fact be thought of as a core part of fully responding to the reliability challenge concerning our beliefs about authoritative normative facts in ethics (or elsewhere). In other words: one might think to adequately explain our reliability about authoritative normative facts, one might need to explain our reliability about the normative facts about which *concepts* we should use in ethics.[112]

One could try to run a version of the conceptual strategy with respect to this reliability challenge about conceptual ethics. If so, one will need to appeal to some conceptual truths at this juncture. The worry is that it's not clear *which* conceptual truths those should be, or how helpful they will be for explaining our reliability. But one option briefly worth mentioning is this. Suppose we are wondering what explains why using the concept AUTHORITATIVE VALUE helps us reliably get at (put intuitively) the *normatively important* normative facts. Perhaps we should use conceptual truths *not* about AUTHORITATIVE VALUE but rather ones about the concept we use to judge how important the normative facts in the extension of that concept are. Thus, perhaps we should use the second-

[109]This is assuming that these areas aren't fundamentally normative. This assumption obviously might be challenged, especially in the case of logic.

[110]My terminology of 'conceptual ethics' draws from Burgess and Plunkett, 'Conceptual Ethics I' (n 42) and Burgess and Plunkett, 'Conceptual Ethics II' (n 42).

[111]This challenge ties into the second kind of worry that Schechter raises for the conceptual strategy Schechter (n 1) 454, about maintaining a commitment to what he thinks of as a "non-plenitudinous" form of non-naturalistic realism. I don't think the challenge I explain below is exactly the same as the challenge that motivates Schechter's worry there. But it is in much the same vein.

[112]See David Plunkett, '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' (2016) 3 Ergo 27 for connected discussion.

order concept *of* a concept being authoritatively normative.[113] If so, we then face the following issue for the conceptual strategy: how extension-revealing is *that* concept?[114] Perhaps that second-order concept involves enough constraints such that ordinary competent users of that concept (given their ordinary knowledge about reality) can know *which* concepts are in fact authoritatively normative, *even if* those concepts are in turn ones that are not very extension-revealing at all with respect to the authoritative normative facts (e.g., about what we authoritatively should do).[115] If that is right, then it gives some cause for optimism about the prospects of using a version of the conceptual strategy to respond to this reliability challenge about conceptual ethics.

It might, however, be more difficult to explain the relevant kind of reliability in conceptual ethics than this suggests. This is because we might want to explain the reliability of our normative beliefs in conceptual ethics about *all* of our normative concepts, *including* our beliefs about when we should use the second-order concept of a concept being an authoritatively normative one. It's not at all clear how to explain *that* reliability. More to the point in this context: it's not at all clear how conceptual truths about *our* concepts could play a definitive role here, without collapsing into something akin to a seemingly trivial form of self-vindication, or without begging the question in an objectionable way.[116]

There is obviously a lot more to say on this kind of reliability challenge, as it arises for our beliefs in conceptual ethics.[117] I think that this challenge has not yet been sufficiently wrestled with, either by those working in conceptual ethics or those working on reliability challenges in ethics. If I am right, this challenge (which I hope will receive more discussion in the coming years) interacts in interesting ways with my core argument in this paper about the conceptual strategy.

8.3. Other epistemological arguments in ethics

Finally, let's turn to epistemological arguments about authoritative normative facts in ethics, or moral facts, *other* than the reliability challenge. Importantly, there are a wide variety of different ways in which people have reconstructed the core of Street's Darwinian Dilemma argument, as well as simply a wide variety of arguments that are in the vicinity of her argument, or of the reliability challenge.[118] Might the conceptual strategy work for

[113]See McPherson, 'Authoritatively Normative Concepts' (n 9), Tristram McPherson, 'Ardent Realism without Referential Normativity' (2019) Inquiry, and Tristram McPherson and David Plunkett, 'Evaluation Turned on Itself: The Vindicatory Circularity Challenge to the Conceptual Ethics of Normativity' in Russ Shafer-Landau (ed), *Oxford Studies in Metaethics*, Vol. 16 (Oxford University Press Forthcoming) for connected discussion.

[114]For connected discussion, see McPherson, 'Ardent Realism without Referential Normativity' (n 110) and McPherson and Plunkett, 'Evaluation Turned on Itself' (n 110).

[115]See McPherson, 'Ardent Realism without Referential Normativity' (n 110) for closely connected discussion.

[116]This ties into an important issue I discussed earlier: namely, how much one can use one's own normative views in explaining the reliability of those views.

[117]For connected discussion, see McPherson and Plunkett, 'Evaluation Turned on Itself' (n 110).

[118]For some of the many different reconstructions of a Street-inspired argument, or articulations of a closely connected kind of argument, see Berker (n 3), Locke (n 86), Matthew S Bedke, 'Intuitive Non-naturalism Meets Cosmic Coincidence' (2009) 90 Pacific Philosophical Quarterly 188, Clarke-Doane (n 1), Karl Schafer, 'Evolution and Normative Scepticism' (2010) 88 Australasian Journal of Philosophy 471, Skarsaune (n 31), Andreas L Mogensen, 'Do Evolutionary Debunking Arguments Rest on a Mistake About Evolutionary Explanations?' (2016) 173 Philosophical Studies 1799, Wielenberg (n 31), Jeff Behrends, 'Meta-normative Realism, Evolution, and Our Reasons to Survive' (2013) 94 Pacific Philosophical Quarterly 486, Setiya (n 1), Louise Hanson, 'The Real Problem with Evolutionary Debunking Arguments' (2017) 67 Philosophical Quarterly 508, Russ Shafer-Landau, 'Evolutionary Debunking, Moral Realism, and Moral Knowledge' (2012) 7 Journal of Ethics

responding to one of those arguments, even if not to the reliability challenge? To exhaustively answer this question, one would need to take each such argument on a case-by-case basis, which is obviously a task that is beyond the scope of this paper. Nonetheless, I think there is something general that can be said here. This is that many of my core points will carry over to many such attempts.

To illustrate, consider the case of the "irrelevant influences" challenge that Vavova argues is at the core of Street's Darwinian Dilemma argument, and which, perhaps more importantly, she thinks forms the basis for an important kind of evolutionary argument against non-naturalistic moral realism.[119] Modifying the statement of the argument Vavova considers somewhat (both in terms of terminology and detail) to make it fit more smoothly with the rest of my discussion, the basic argument can be put as follows:

P1. *Non-Naturalism*. Non-naturalistic moral realism is true.

P2. *Influence*. Evolutionary forces have influenced our token moral beliefs.

P3. *Off-track*. Evolutionary forces aim at fitness, not at moral truths (as conceived of by non-naturalistic moral realism).

P4. *Gap*. The fitness enhancing beliefs and the moral truths come apart.

C. *Mistaken*. We have good reason to think our moral beliefs are mistaken.[120]

As this outline of the irrelevant influences argument makes clear, the core issue here isn't about explaining the (purported) reliability of our moral beliefs. Rather, the issue is that we have an argument to the effect that our moral beliefs (or at least many of them) are mistaken. This is a conclusion that many of us (including Vavova) want to reject.

As I noted earlier, Vavova expresses optimism that a version of the conceptual strategy can help with resisting the conclusion of this argument.[121] Her basic thought seems to be that, because of conceptual truths about the relevant moral concepts, we can know that we

and Social Philosophy, Benjamin James Fraser, 'Evolutionary Debunking Arguments and the Reliability of Moral Cognition' (2014) 168 Philosophical Studies 457, Justin Morton, 'A New Evolutionary Debunking Argument Against Moral Realism' (2016) 2 Journal of the American Philosophical Association 233, Tomas Bogardus, 'Only All Naturalists Should Worry About Only One Evolutionary Debunking Argument' (2016) 126 Ethics 636, Ramon Das, 'Evolutionary Debunking of Morality: Epistemological or Metaphysical?' (2016) 173 Philosophical Studies 417, Elizabeth Tropman, 'Evolutionary Debunking Arguments: Moral Realism, Constructivism, and Explaining Moral Knowledge' (2014) 17 Philosophical Explorations 126, Matthew Braddock, 'Evolutionary Debunking: Can Moral Realists Explain the Reliability of Our Moral Judgments?' (2016) 29 Philosophical Psychology 844, William J FitzPatrick, 'Debunking Evolutionary Debunking of Ethical Realism' (2015) 172 Philosophical Studies 883, Guy Kahane, 'Evolutionary Debunking Arguments' (2011) 45 Noûs 103, Andrew Moon, 'Debunking Morality: Lessons from the EAAN Literature' (2017) 98 Pacific Philosophical Quarterly 208, Billy Dunaway, 'Epistemological Motivations for Anti-realism' (2018) 175 Philosophical Studies 2763, Michael Klenk, 'Evolution and Moral Disagreement' (2018) 14 Journal of Ethics and Social Philosophy, Silvan Wittwer, 'Evolutionary Debunking, Self-Defeat, and All the Evidence' in Michael Klenk (ed), *Higher-Order Evidence and Moral Epistemology* (Routledge 2020), and Daniel Z Korman and Dustin Locke, 'Against Minimalist Responses to Moral Debunking Arguments' in Russ Shafer-Landau (ed), *Oxford Studies in Metaethics*, Vol. 15 (Oxford University Press Forthcoming).

[119]Vavova, 'Evolutionary Debunking of Moral Realism' (n 14) and Katia Vavova, 'Debunking Evolutionary Debunking' in Russ Shafer-Landau (ed), *Oxford Studies in Metaethics*, vol 9 (2014). Note that Vavova also thinks this kind of argument is a good way of making sense of the "evolutionary debunking" arguments in a range of work beyond Street's, including, for example, the arguments in Bedke (n 115) and Richard Joyce, *The Evolution of Morality* (MIT Press 2006). See Vavova, 'Irrelevant Influences' (n 86) for connected discussion about irrelevant influences arguments in general.

[120]The above reconstruction is a slightly modified version of Vavova's reconstruction in Vavova, 'Evolutionary Debunking of Moral Realism' (n 14) 108.

[121]She does so at the end of ibid. It should be noted, however, that Vavova's preferred way of responding to the irrelevant influences argument is *not* just to endorse a version of the conceptual strategy. See Vavova, 'Debunking Evolutionary Debunking' (n 116) and Vavova, 'Irrelevant Influences' (n 86) for her more developed positive views.

can't be *too* off track in our moral judgments. Consider premise 4 (*Gap*): "the fitness enhancing beliefs and the moral truths come apart". In order for us to have reason to think *this*, we must have some sense of what moral truths *are*. After all, the issue here isn't that the fitness enhancing beliefs come apart from *some* truths. It is that they come apart from the specifically *moral* ones.[122] Once we then understand what it is for truths to be *moral* ones in particular – which (on the strategy we are considering here) we do by understanding our relevant moral concepts – we can see some things that morality must be like, if there is morality at all (or if there is moral wrongness at all, etc.). Then, once we know the relevant conceptual truths, we can discover that, in fact, we aren't that *off track* after all in our moral judgments. As with using the conceptual strategy to respond to the reliability challenge, we can insist that certain (purported) foundational moral truths are conceptual truths (e.g., that, as long as nihilism is false, then killing others for fun is not morally permitted, and that we generally morally should prefer pleasure over pain, etc.). And given the centrality of those truths for discovering other, derivative moral truths – perhaps, for example, in part through the process of making these judgments cohere with each other in a kind of "reflective equilibrium" and through using general-purpose reasoning capacities – we end up in good shape with respect to a whole range of moral judgments. Or at least so the basic thought goes.

The core issues I have raised for the conceptual strategy when used for the reliability challenge also apply here. This is because the *same* features of the relevant moral concepts that were called on to do the heavy-lifting in the case of the reliability challenge are also being called on to do the heavy-lifting here: namely, that the relevant moral concepts are fairly extension-revealing (and are so because of conceptual truths about those moral concepts). Thus, the same line of argument I ran in this paper can be run here, starting with the possibility (which I illustrated by discussing the Gibbardian fitting attitude account of moral concepts) that our relevant moral concepts are not very extension-revealing, to the issues tied to the tradeoffs between the extension-revealing and authority-revealing aspects of normative concepts, and, finally, to the issues with running a version of the conceptual strategy when we focus on authoritative normative facts as such, which are (I have argued) the kinds of facts that ultimately matter the most in ethics. All of these issues carry directly over here as well. This is, I think, illustrative of how this will often go if someone uses something akin to the conceptual strategy to respond to an epistemological argument in the vicinity of Street's Darwinian Dilemma argument.

This is not to say that there won't be potential differences here between different arguments. For example, an interesting feature of the irrelevant influences argument worth noting is this: it might raise a problem with the idea of focusing on authoritatively normative concepts, in the way that I (following Street) have been advocating for doing in the context of the reliability challenge. In the context of the irrelevant influences argument, the core worry is that if we do so, then the arguments under consideration might lose focus in a problematic way. In short, in order to know that we are *off track* in our judgments about X, we need to know something about what is correct about X. That might well be harder to do with authoritative normativity than with morality.[123] Moreover,

[122]As Vavova puts it: "We could begin our response to the debunker by noticing that we need some assumptions about what morality is like to recognize that we might be wrong about morality." Vavova, 'Evolutionary Debunking of Moral Realism' (n 14) 113.

[123]See Vavova, 'Debunking Evolutionary Debunking' (n 116) for connected discussion.

the less we are able to start with something we know is correct about X, the more we might just slide into a form of general skeptical worry about our ability to know anything. And, as Vavova emphasises, irrelevant influences arguments are *not* supposed to collapse into a general form of skepticism.[124]

These issues tie into a range of further issues about how to understand and best develop irrelevant influences arguments in general, and into a host of foundational issues in epistemology that Vavova discusses in addressing this topic.[125] It is beyond the scope of this paper to address all of those issues here. Instead, the point I want to make is more limited. This is that we should be cautious to not overstate the worry here about running an irrelevant influences argument about authoritative normativity, as opposed to one about morality. In short, it would be a mistake to think that the switch to authoritatively normative concepts from moral concepts dooms the irrelevant influences argument against non-naturalistic realism (if such an argument is indeed doomed). We can know we are thinking about a given topic – e.g., what an agent authoritatively ought to do – without knowing the kinds of details about the extensions of the relevant concepts that the proponent of the conceptual strategy has in mind in her discussions, e.g., which acts we authoritatively ought to perform, or what those acts are like (e.g., that they involve fairness etc.). The concept AUTHORITATIVE PRACTICAL OUGHT fixes a topic (as do other authoritatively normative concepts), assuming we grant that concepts fix topics to begin with. But it does not do so in a way that is extension-revealing. Moreover, as I have argued, the same would be true about the relevant moral concepts (e.g., MORAL WRONGNESS) if a broadly Gibbardian fitting attitude account of them is correct. So it is not clear how deep the contrast here goes between authoritatively normative concepts and moral concepts.

Further issues obviously remain when thinking about the irrelevant influences argument on the table, including many that Vavova emphasises in her work. For example, one such issue is this: what kind of information (and how much of it) does one need to think we are *really* off track in our judgments in these domains (and not just that it might be possible that we are)?[126] That issue will matter a lot in the context of thinking about the conceptual strategy as applied here to the irrelevant influences challenge. Indeed, it is, I think, one of the key questions for thinking about how the conceptual strategy interacts with this kind of challenge.

9. Conclusion

In this paper, I have discussed a certain kind of strategy for responding to the reliability challenge for non-naturalistic realism about morality. This strategy, which I have called "the conceptual strategy", appeals to (purported) conceptual truths about relevant moral concepts in order to respond to the challenge. I argued that this strategy faces a number of issues, tied to the tradeoff between the extension-revealing and authority-revealing aspects of normative concepts. These issues might of course be dealt with by future versions of the conceptual strategy. But the issues run deep, in a way that I think has not been sufficiently wrestled with by those sympathetic to the conceptual strategy. Following this discussion, I then

[124]See Vavova, 'Irrelevant Influences' (n 86), as well as Vavova, 'Evolutionary Debunking of Moral Realism' (n 14) and Vavova, 'Debunking Evolutionary Debunking' (n 116).
[125]See especially Vavova, 'Irrelevant Influences' (n 86).
[126]See ibid.

turned to exploring how the conceptual strategy fares with respect to dealing with the reliability challenge concerning what I have called "authoritative normative facts". I argued this is the version of the challenge that ultimately matters the most in ethics (or which, at the very least, matters just as much as the parallel challenge about morality), and that the issues for the strategy run even deeper when applied to this challenge. Following this, I then explored how my core argument matters for a range of other arguments, including different versions of the reliability challenge (targeting different metamoral or metaethical views, or targeting domains other than morality or ethics), as well as other epistemological arguments in ethics (e.g., what Vavova calls "irrelevant influences" arguments).

As this above gloss of my work in this paper underscores, a key premise in my discussion has been that there is a tradeoff between the extension-revealing and authority-revealing aspects of normative concepts. In this paper, I have tried to motivate the idea that there is such a tradeoff, in large part by pointing to different examples. I have also made some arguments for why we would expect there to be this kind of tradeoff, given the different roles different normative concepts play in our lives (and which we should want them to play). I have also discussed some proposed analyses of given concepts (e.g., McPherson's account of AUTHORITATIVE PRACTICAL OUGHT) that, if correct, would go some way toward explaining why those concepts are more or less authority-revealing or extension-revealing. But nothing I have said in this paper approaches a full explanation of *why* there is a general tradeoff between the extension-revealing and authority-revealing aspects of normative concepts. I think that explaining *that* – assuming that I am right that there is this tradeoff, or at least something closely akin to it – is a deep philosophical issue, which matters for a range of work in normative and metanormative inquiry. Tied to this, there is also simply the issue of getting clearer on exactly what this tradeoff consists in, which goes beyond the schematic way I have mapped out for what it is for a concept to be extension-revealing or authority-revealing. This too is something that I think matters in significant ways for a range of work in normative and metanormative inquiry. My hope is that my discussion in this paper can be seen as helping to make the case for why thinking about these issues matters for future philosophical theorising, both about the normative concepts we in fact have, as well as about the ones we should have.

Acknowledgements

Earlier versions of this paper were presented at the Philosophy Mountain Workshop, a workshop on normative epistemology at University of Freiburg, the ConceptLab seminar series at the University of Oslo, a workshop on new waves in the study of normative concepts at the Frankfurt School of Finance and Management, and the Philosophy Desert Workshop. Thanks to everyone who participated in those sessions. Thanks also to Sarah Buss, Terence Cuneo, Billy Dunaway, David Enoch, Rachel Fraser, Allan Gibbard, Coby Gibson, Dustin Locke, Daniel Gun Lim, Errol Lord, Tristram McPherson, Eliot Michaelson, Michael Morck, Kate Nolfi, Peter Railton, Adrian Russian, Raul Saucedo, Joshua Schechter, Dan Singer, Sharon Street, Amie Thomasson, Silvan Wittwer, Katia Vavova, Victoria Xiao, Alice Zhang, the editors of *Jurisprudence* (George Pavlakos and Veronica Rodriguez-Blanco), and to the anonymous referees who reviewed this paper.

Disclosure statement

No potential conflict of interest was reported by the author.