

THE COSTS OF EPISTEMIC REALISM

Abstract: The present paper considers some neglected costs of epistemic realism. It argues that those responding to the realist's main argument against epistemic expressivism—the so-called *perspective objection*—have overestimated the power of that argument, since a central premise of it can actually be used to turn the tables on the realist. More specifically, the premise entails that, unless the realist accepts a far-reaching scepticism, she must do two things. First, she must reject the idea that true belief is a central epistemic goal. Second, she must hold that the diversity of views in discussions about epistemic normativity is a sign of cognitive-behavioural incoherence, if not of widespread irrationality, on the part of epistemologists. Such are the costs of epistemic realism.

Keywords: Epistemic Realism; Epistemic Expressivism; Scepticism; Epistemic Goals; Epistemic Normativity.

1. Introduction

Expressivists explain normativity, not by way of some independent value referred to through normative discourse, but in terms of the norms accepted or sentiments expressed by agents through such discourse. As Simon Blackburn (1993) says in relation to moral discourse, expressivism (or *quasi-realism*, as he calls it) 'avoids the view that when we moralize we respond to, and describe, an independent aspect of reality' (157). Rather, when we assert values, we 'voice' (Blackburn 1998: 50) or 'express' (Gibbard 1990: 8) our state of mind. This makes for what J. Adam Carter and Matthew Chrisman (2012) refer to as 'the core expressivist maneuver' of addressing questions about what values *are* in terms of what it is to *judge* something to be valuable.

Expressivists disagree about what exactly the relevant judgments express. Suggestions include states of norm-acceptance (Gibbard 1990), practical states (Blackburn 1998), 'plan-laden'

normative beliefs (Gibbard 2003), and combinations of beliefs and desires (Ridge 2006). Common to most expressivist accounts, however, is a naturalistically motivated anti-realism about normative facts.¹ The particular kind of facts denied by the epistemic expressivist is *epistemic* facts, the most relevant feature of which is that they are *authoritative*, to borrow a term from Terence Cuneo (2007).² According to Cuneo, epistemic facts are authoritative in that they ‘are, imply, or indicate categorical reasons for agents to behave in certain ways,’ in virtue of which ‘whether I have an epistemic reason, say, to believe a proposition for which I have compelling evidence is not contingent upon whether I care about believing what is true’ (59).

The purpose of the present paper is not to defend epistemic expressivism. In epistemology, as in other theoretical pursuits, theories are accepted on the basis of their costs and benefits. Benefits consist in possibilities for explaining significant phenomena, while costs comprise undesirable implications. Recent discussions of epistemic expressivism have focused on the costs of epistemic expressivism (e.g., Lynch 2009; Cuneo 2007), the benefits of epistemic expressivism (e.g., Ahlstrom-Vij 2013; Field 2009; Chrisman 2007; Gibbard 2003 and 1990), and the benefits of epistemic realism (e.g., Cuneo 2007). What follows focuses on the fourth and heretofore neglected part of the dialectical space: the costs of epistemic realism.

¹ I am setting aside concerns about *creeping minimalism* (Dreier 2004), or the problem of how the sophisticated minimalist machinery of recent forms of expressivism makes the distinction between traditional realism and anti-realism hard to draw in the contemporary debate. Consider for example Blackburn’s fairly liberal attitude towards talk about ethical facts and properties (e.g., in his 1999: 216; and his 1993: 81) and the ontologically non-committal stance of Gibbard on the existence of normative facts (see, e.g., his 2003: x and 18). For present purposes, I will assume two things: First, a clear distinction between realism and anti-realism can be drawn in the epistemological domain. Second, it is to be drawn in terms of the acceptance or denial—or, to accommodate Gibbard’s non-committal stance, *non-acceptance*—of the existence of epistemic facts, as understood by epistemic realists like Cuneo (2007).

² But see footnote 19.

As for the nature of those costs, it will be argued that those responding to the realist's main argument against epistemic expressivism—the so-called *perspective objection*³ (Section 2)—have overlooked the fact that a central premise of that argument can be used to turn the tables on the realist. More specifically, the premise entails that, unless the realist accepts a far-reaching scepticism concerning claims about normativity (Section 3), she must reject the idea that true belief is a central epistemic goal (Sections 4-5), and hold that the diversity of views in discussions about epistemic normativity is a sign of cognitive-behavioural incoherence, if not of widespread irrationality, on the part of epistemologists (Section 6). The purpose of identifying these costs is not to suggest that we should be expressivists, but to contribute to a more complete picture of the costs and benefits of epistemic realism and expressivism, respectively. Only against the background of such a picture can we make an informed choice between the two.

2. The Perspective Objection

The main charge against epistemic expressivism starts out with the claim that it presupposes a distinction between an internal and an external perspective. Cuneo explains:

The internal perspective . . . is supposed to be the perspective that captures what it is like to be an agent engaged in ordinary epistemic thought and discourse; it is the arena in which it appears to an agent that she is giving and assembling epistemic reasons, epistemically evaluating beliefs, uttering epistemic truths, representing epistemic reality, and so forth. The external perspective, by contrast, is supposed to be the perspective of the naturalist philosopher who in Terry Horgan and Mark Timmons's words is engaged in 'metaphysical specula-

³ While the perspective objection is the main objection to *epistemic* expressivism, there are of course other problems for expressivism generally that might also apply to epistemic expressivism, including the Frege-Geach problem (Geach 1965). See Ridge (2006) and Boisvert (2008) for two attempts to come to terms with that problem.

tion' or 'theoretical inquiry', but believes that there are no epistemic reasons or facts (Cuneo 2007: 170).

The 'naturalist philosopher' at issue here is the epistemic expressivist who denies not only that there are any epistemic facts, but—according to Cuneo—also that there are any epistemic reasons. The problem for the expressivist is that the idea of an external perspective makes no sense in epistemology:

. . . it is very difficult to see how there could be a perspective in which a person at once engages in theoretical inquiry and does not believe (or take it for granted) that there are epistemic reasons. After all, anything we could recognizably call 'theoretical inquiry' or 'inquiry from a naturalist perspective' involves viewing ourselves as assembling reasons, epistemically evaluating claims, offering arguments, and so forth (2007: 171).

The idea that the epistemic expressivists must inhabit a philosophically uninhabitable perspective amounts to what Cuneo refers to as *the perspective objection* (see also Lynch 2009). Available responses to this objection include the following. First, Carter and Chrisman (2012) defend a metaphysically *agnostic* form of expressivism, on which the expressivist is not committed to an external perspective, on account of remaining neutral on whether there are epistemic facts. Second, Klemens Kappel (2011) argues for the *anti-realist* response that there are epistemic reasons, and theoretical inquiry, in the absence of epistemic facts. I will not evaluate the merits of these responses here. Instead, I will assume that at least one of them works, and then defend the so-far overlooked point that we can use a central premise of the perspective objection to turn the tables on the epistemic realist. The premise is the following:

- (P) When engaging in a theoretical inquiry into epistemic normativity, we are engaged in the business of assembling and exchanging epistemic reasons.

As Cuneo points out, it would be hard to make sense of what we are doing when inquiring about epistemic normativity, if we are *not* thereby engaged in the business of assembling and exchanging epistemic reasons. Contrary to what Cuneo suggests, however, acknowledging that this is so doesn't present a problem for the expressivist, who (I'll assume) can accommodate (P) along the lines of Carter and Chrisman (2012) or Kappel (2011). The challenge instead lies with the realist. But before explaining why, let us get clearer on the realist's theoretical commitments.

3. Realism and Reason-Responsiveness

The epistemic realist at issue takes the following to be true:

- (ER) There are categorical epistemic reasons in virtue of there being epistemic facts.

According to Cuneo (2007), epistemic reasons are categorical in that they exercise a particular kind of *authority* over us, and thereby 'inescapably govern our conduct' (59). Epistemic reasons are authoritative 'inasmuch as the decisiveness of some such reasons for an agent is not a function of whether she wants to act in an epistemically commendable fashion, or belongs to a certain social group, or has entered into certain agreements with others, and so forth' (59), nor 'contingent upon whether I care about believing what is true' (*ibid.*)

I will follow Cuneo in taking a commitment to (ER) to be *necessary* for being an epistemic realist, whether or not it is also sufficient. The reason is that, if what is to be argued below is on the right track, anyone accepting (ER)—which, if Cuneo is right, includes *all* epistemic realists—is bound to face serious problems when attempting to say something substantive about what ep-

istemic facts there are, and does so simply in virtue of the combination of (ER), (P), and an additional commitment to which we now turn.

To get at this additional commitment, consider that the realist is going to want to commit to epistemologists inquiring into the nature of epistemic normativity *responding* to epistemic facts. This simply means that, when epistemologists engage in the relevant type of inquiry—i.e., when cognitively evaluating beliefs, giving and asking for reasons, and so forth—they are *sensitive* to the categorical reasons there are in virtue of the epistemic facts, and possibly tacitly so. (Might there be *degrees* of responsiveness? We will discuss this question in Section 5.) Why should the epistemic realist commit to such responsiveness? Because to deny that is to hold that any epistemological theorizing in accordance with good reason is accidental at best. And if so, why pay any attention to the conclusions epistemologists reach on the basis of such inquiry? To fail to *track* reason is to *lack* reason, and consequently also knowledge and justification, on the relevant matters. In other words, if epistemologists are *not* responsive to the epistemic facts and reasons, the only available stance *vis-à-vis* claims made on the basis of their inquiry is one on which we hold that we neither know nor are justified in believing the claims in question. That is, in the absence of responsiveness, the only available stance *vis-à-vis* the relevant domain of inquiry is one of *scepticism*, entailing that any claims about epistemic facts or reasons lack justification, and consequently also fail to amount to knowledge.⁴ Assuming that the epistemic realist finds such scepticism unpalatable, she would have to commit to the following:

⁴This might be taken to be too quick. Perhaps we have reason to pay attention to the conclusions of epistemologists, even if epistemological theorising doesn't generate justification or knowledge, because some of the theories involved are—even if only accidentally—on the right track, and we might make progress towards knowledge by engaging carefully with them. But such progress will only be possible if we in so engaging are thereby responsive to the categorical reasons there are in virtue of the epistemic facts, in which case epistemological theorising *is* ultimately responsive to the categorical reasons there are, in accordance with the anti-sceptical commitment under consideration.

- (R) When engaging in epistemological inquiry, we are responsive to the categorical epistemic reasons there are in virtue of the epistemic facts.

While (ER) commits the epistemic realist to the claim that *there are* epistemic facts, it does not commit her to saying anything in particular about *what* epistemic facts there are. However, I will argue that, as soon as the realist attempts to tell us something substantive about what epistemic facts there actually are, we see that the combination of (ER), (P), and (R) has two implications that she is going to find hard to swallow.

In order to show this, let us tease out some more specific ideas on the part of the realist as to what epistemic facts there are. One way to do so is by looking more closely at her conception of epistemic merit terms like ‘justification’, ‘warrant’, and ‘knowledge’. After all, when believing with justification or warrant, or when knowing something, we believe with (good) *reason*. So, if we want to find out what reasons there are, and as such what epistemic facts there are, we can look at what the realist says about the relevant merit terms. For that purpose, let us return to Cuneo.

Appealing to certain ‘platitudes’ about epistemic facts, Cuneo suggests that ‘[e]pistemic merits such as being justified, entitled, warranted, a case of knowledge, and so forth are almost always understood to be alethic merits,’ i.e., as ‘merits the presence of which in an agent’s attitude indicates that the content of his attitude is true, that the content of his attitude is likely to be true, or that he has done a good job of trying to get at the truth’ (58). Cuneo does not provide more details on the issue, so the next two sections will consider the two most promising ways to flesh out the relevant relation between epistemic facts, reasons, and alethic merit, by considering the two most prominent approaches to epistemic normativity. Sections 4 and 5 will spell out the relevant reasons in *consequentialist* terms, i.e., in terms of means to epistemic ends. Since we, following Cuneo, will be concerned with alethic merit in particular, the relevant end is true belief. Sec-

tion 6 will then consider the option of spelling out the relevant reasons in non-consequentialist terms, and in terms of epistemic *duties* in particular.

Before proceeding, however, note that what follows does not assume that a commitment to realism *implies* any commitment on the question of whether the relevant reasons are to be understood in consequentialist, deontological, or other terms. That would be to confuse metaepistemology with normative epistemology. Still, we would want a commitment to realism to be *compatible* with at least one normative story about what the facts and reasons postulated by the realist are like. The following sections consider the two main contenders for the purpose of surveying the realist's options. As we shall see, the problem is that opting for either has implications that she would have to consider unpalatable.

4. True Belief as a Central Epistemic Goal

True belief, Michael Lynch (2009) claims, is an end of inquiry, and thereby also an epistemic end.⁵ Moreover, if true belief is an epistemic *end*, then true belief is a *good*. More specifically, Lynch suggests that it is 'prima facie good that one believes all and only what is true' (77), which is to say that 'there is always something to say for it' and it's as such 'good considered by itself but not necessarily good all things considered' (78). More specifically, it is always *epistemically* good to believe all and only true belief, whether or not it is good all things considered.

⁵ On the idea that true belief is an epistemic end, see Alston (1985: 83-84), Bonjour (1985: 7-8), Moser (1985: 4), and Foley (1993: 19), as well as David (2001) for an overview and further references. On the idea that the domain of inquiry demarcates the epistemic domain, see Alston (2005: 30), Haack (1993: 199), and Bonjour (1985: 83-84).

Taking it to always be epistemically good to believe all and only truths is compatible with taking there to be other epistemic goals, in addition to true belief.⁶ However, Lynch clearly holds the following:

(T) True belief is a central epistemic goal.⁷

Lynch takes true belief to be a *central* epistemic goal in (at least) two respects. First, the reason we pursue other candidate epistemic goods, such as coherence and justification—and, we might add, evidence and reliability—is because we take them to be conducive to the goal of true belief (89). As such, true belief is a central epistemic goal in that it explains a significant part of what we *do* when engaging in inquiry. Second, it is *constitutive* of inquiry that it aims to produce true belief (90-91), and Lynch’s main objection to epistemic expressivism is exactly that it leaves room for conceptions of inquiry that imply otherwise.⁸ As such, the truth norm is also central in that it plays an important role in explaining what inquiry *is*.

Having introduced (T), Lynch shifts from talk about what is *good* to talking about what one *should* believe, and suggests that ‘we shouldn’t just believe what is true; we should also avoid error’ (77). But that seems too fast. What is good to believe and what we should believe are not trivially connected. An analogy might be helpful. Hedonism tells us what is *good*, namely happiness. But hedonism doesn’t tell us what to *do*—unless combined with a normative ethical theory

⁶ On this point, note that the idea that true belief is an epistemic goal is typically embraced also by the value pluralist, who simply denies that true belief is *unique* in this respect. See, e.g., Kvanvig (2005).

⁷ We might worry that (T) implies that all kinds of trivial truths are epistemically valuable. If the reader shares this worry, she may qualify (T) in terms of true belief that pertains to matters that we care about. Such a qualification makes no difference to the arguments to be provided in what follows.

⁸ See Ahlstrom-Vij (2013) for a discussion of Lynch’s argument.

that provides a bridge principle between the right and the good. Classic utilitarianism provides such a principle: do whatever maximizes the good.

At this point, we may return to Cuneo's idea that epistemic merits are alethic or 'truth-connected' merits. As noted above, epistemic merit terms tell us something about reasons on account of how instantiating the relevant merits—e.g., justification, warrant, and knowledge—involves *believing* with (good) reason. Moreover, one way to read Cuneo's take on epistemic merits as alethic merits is in *consequentialist* terms, i.e., as merits constituting merits on account of being conducive to epistemic ends, and to true belief in particular. (We will consider the most plausible non-consequentialist construal in Section 6.) This suggests a particular bridge principle, which we may refer to as the *consequentialist bridge principle*:

(CBP) If one's beliefs are formed by way of a belief-forming process conducive to attaining the epistemic goal(s), then one has reasons for those beliefs.

This is a good time to take stock. We started out by asking what epistemic facts the epistemic realist might take there to be, and then went on to reconstruct some suggestions by Cuneo and Lynch in terms of (T) and (CBP), where the former talks about true belief as a central goal, and the latter ties epistemic goals to epistemic reasons. By combining (T) and (CBP)—or (T+CBP) for short—we get the idea that beliefs produced by processes conducive to true belief are beliefs held with reason. That is an epistemic fact, on this reading of Cuneo and Lynch. This doesn't presuppose that true belief is the *only* epistemic goal. Still, if we follow Lynch in taking true belief to be a *central* goal of inquiry, and also take epistemic merit terms to be understood in consequentialist terms, then (T+CBP) will be at work in the great majority of epistemic cases, and capture one central and important way to believe with reason.

5. The Problem of Non-Convergence

On the reading of the realist proposed in the previous section, to be responsive to the categorical reasons there are in virtue of (T+CBP) is to be responsive to the truth.⁹ More specifically, someone is *perfectly* responsive to (T+CBP) if *all* of her beliefs are formed in ways conducive to believing truly, and she thereby believes a great many true things and all of her beliefs are at least likely to be true. By contrast, someone is *imperfectly* responsive to (T+CBP) if many of her beliefs are formed in ways conducive to believing truly, and she thereby believes many true things and most of her beliefs are at least likely to be true. Note that this is not to deny that you can have sufficient reason to believe falsehoods. It's simply to maintain that, if having (epistemic) reasons is to be understood in terms of being at least imperfectly responsive to whatever categorical reasons there are, and to be responsive thus is to be responsive to the truth, having sufficient reason to believe is incompatible with not at the very least believing many true things, and for many beliefs that are not in fact true to at least be likely to be true.

However, it might be thought that this incompatibility is a mere artefact of (CBP) simplifying too much what a consequentialist has to think about the relationship between reasons and truth.¹⁰ By way of example, consider the following, alternative bridge principles:

(CBP_N) If one's beliefs are formed by way of a belief-forming process conducive to attaining the epistemic goal(s) *under normal circumstances*, then one has reasons for those beliefs.

(CBP_I) If one's beliefs are formed by way of a belief-forming process conducive to attaining the epistemic goal(s) *given appropriate inputs*, then one has reasons for those beliefs.

⁹ For the sake of brevity, I will occasionally talk about being responsive to epistemic facts, instead of being responsive to the epistemic reasons there are in virtue of the epistemic facts that there are.

¹⁰ Many thanks to an anonymous reviewer for raising this matter.

If the relevant bridge principle is weakened along either of these lines, someone being (imperfectly) responsive to (T) might be compatible with that person *not* forming her beliefs in ways conducive to believing truly. But note that a consequentialist looking to invoke that possibility to suggest that epistemologists are responsive to reasons yet not forming beliefs in ways conducive to believing truly will have to say more than this. In particular, she will also have to say—on (CBP_N)—that epistemologists are forming beliefs under *abnormal circumstances*, or—on (CBP_I)—that their belief-forming processes are operating on *inappropriate inputs*. That’s not to say that anyone is necessarily *blameworthy*, and, for what it’s worth, something like blameworthiness would seem to be what the notions of ‘reason’ in (CBP_N) and (CBP_I) would be modelling. It’s also not to say that the only appropriate attitude towards the conclusions of epistemologists would be one of scepticism, as they may still have (good) reason to believe what they do, so long we accept that (CBP_N) and (CBP_I) capture genuine notions of ‘reason.’ But just as in a case where such scepticism is warranted, it would still be the case that we may rightly ask why we should pay any attention to what epistemologists are saying on the relevant matters. Because if we hold on to (T), it would still be the case that something has gone *wrong*, in that the epistemologists involved have failed (whether they’re blameworthy or not) to get at the truth. Indeed, they have failed *systematically* to get at the truth, not just in particular instances. Had they not, their beliefs would have been formed in ways conducive to believing truly, in which case the implications for the present argument would be no different than on (CBP).

That’s to say that, whether or not any consequentialist would prefer something along the lines of (CBP_N) or (CBP_I) over (CBP) on independent grounds, there are good reasons not to want to take on the additional commitment of also saying that epistemologists in particular are responsive to the type of reasons mentioned in one of these weakened principles, but *not* to those mentioned in (CBP). Because to say that is to say that epistemologists are getting the relevant facts systematically wrong. I will assume that the consequentialist is not going to want to say

that. And if she is not going to want to say that, then whether she takes (CBP) to be the correct bridge principle, she will accept that, as far as the reason-responsiveness of epistemologists is concerned, it's such that she at the very least believes many true things, and for many beliefs that are not in fact true, they are at least be likely to be true. But that's of course to say that, in the case of epistemological inquiry, things are exactly as they would be, were (CBP) the correct bridge principle.

With that in mind, let's pull together all of the commitments we have now uncovered: The realist believes there to be epistemic facts, giving rise to categorical epistemic reasons, as per (ER), and that engaging in a theoretical inquiry into epistemic normativity is to engage in the business of assembling and exchanging epistemic reasons, as per (P). Moreover, to be engaged thus is to be responsive to the categorical epistemic reasons there are in virtue of such facts, as per (R), and being responsive thus is to be responsive to the truth in the manner just outlined in relation to (T+CBP). If we assume at least imperfect responsiveness, we can thereby expect those inquiring in this manner to believe many true things about epistemic normativity, and for most of their theories about such normativity to be (at least) likely to be true.

But if so, we should expect to see a *convergence* among theories of epistemic normativity. After all, if epistemic reasons are categorical, *everyone* theorizing about epistemic normativity is bound by the *same* reasons, simply in virtue of engaging in a form of inquiry. Moreover, if (T+CBP) is an epistemic fact, then to be responsive to the relevant reasons is to be responsive to the *truth*.¹¹ The problem for the epistemic realist is that convergence is not what we are seeing. What we are seeing is a great theoretical diversity of mutually incompatible positions, ranging from epistemo-

¹¹ This is also why the problem developed in this section is not a problem about epistemic reasons *as such* but specifically about the kind of *categorical* epistemic and moreover alethic reasons postulated by the epistemic realist, if having such reasons is a matter of truth-conduciveness. If reasons are either not categorical or not necessarily truth-conducive, then a set of inquirers may all be responsive to reasons without necessarily converging on the same belief, either because they are not bound by the same reasons, or because the relevant reasons are not truth-conducive.

logically deontological theories spelled out in terms of moral duties (Clifford 1886) to epistemologically deontological theories spelled out in terms of epistemic duties (Feldman 2002); from epistemologically consequentialist theories on which true belief is the only epistemic goal (Ahlstrom-Vij and Dunn 2014; David 2005; Goldman 1999) to consequentialist theories on which true belief is only one epistemic goal among many, alongside understanding (Kvanvig 2005), wisdom (Whitcomb 2007), and knowledge (DePaul 1993); and on to pragmatist theories on which there are no epistemic goals or norms at all (Stich 1990).¹²

Since arguments from non-convergence against realist views are nothing new, let us be clear about what the above line of reasoning does *not* show. Consider J. L. Mackie's (1977) well-known 'argument from relativity' against moral realism. According to Mackie, 'the actual variations in the moral codes are more readily explained by the hypothesis that they reflect ways of life than by the hypothesis that they express perceptions, most of them seriously inadequate and badly distorted, of objective values' (37). Unlike Mackie's, the present argument is *not* that epistemic anti-realism provides a better *explanation* than does realism for the absence of convergence. That is, the argument does not take the following form: *there is diversity, hence (by abductive inference) anti-realism is true*.¹³ It is simply that the combination of (ER), (P), (R), and (T+CBP) generates the prediction that we should see far more convergence among theories about epistemic normativity

¹² This claim about theoretical diversity is an empirical claim, of course. There has not been a lot of empirical investigation into the distribution of views of philosophers. The only exception is Bourget and Chalmers (2014), who find that, while philosophers do not have very accurate views about the views of other philosophers, there is certainly a great diversity of views in philosophy, including in epistemology. The one outlier in epistemology is non-skepticism, which 81.6% of respondents indicated they either 'accept' or 'lean toward'.

¹³ Consequently, the present argument does not over-generalize in the manner that a straightforward argument from diversity to anti-realism might do. After all, diversity crops up in a great many places—not the least within philosophy—and any straightforward argument from diversity to anti-realism runs the risk of counting against more kinds of facts than most of us would be comfortable rejecting.

than we are in fact seeing. That puts pressure on the realist to relinquish one of these claims, but—unlike in the case of Mackie’s argument—not necessarily to accept anti-realism.¹⁴

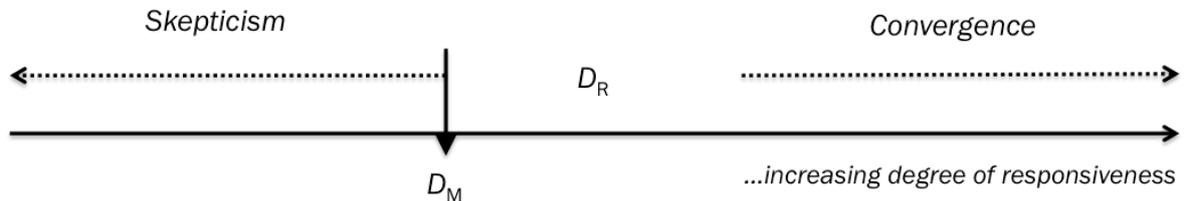
Moreover, there is a familiar rejoinder to Mackie’s argument: variations in moral *codes*—i.e., patterns of moral judgments about particular cases—do not imply variations in fundamental moral *principles*, since even people who agree on fundamental principles may arrive at different codes on account of applying shared principles under different circumstances.¹⁵ That’s to say that Mackie’s argument fails to show that there is non-convergence where it matters, i.e., at the level of fundamental, moral principles. But it is clear that an analogous rejoinder does *not* work in response to the above problem of non-convergence, which highlights another respect in which the present argument differs from Mackie’s. In the case of theories about epistemic normativity, where we see a lack of convergence is exactly at the level of fundamental epistemic principles. Whether or not there is any diversity in epistemologists’ patterns of judgments about *particular cases*, it is clear that they disagree over the *principles* supposed to account for those patterns.

In light of this, the epistemic realist might instead follow some moral realists in appealing to the indeterminacy of the relevant facts (Shafer-Landau 1994), or to the relevant theories being underdetermined by the data (Tersman 2006). However, any attempt to accommodate non-convergence in this manner is *ipso facto* an attempt to postulate an ‘epistemic roadblock’ between the relevant facts and us. It is open to the epistemic realist to account for an absence of theoretical convergence by identifying such roadblocks. But doing so will raise the question whether the account is still compatible with our being sufficiently responsive to the relevant facts for (R) to hold and scepticism concerning claims about epistemic normativity be kept at bay.

¹⁴ This is also what sets the problem of non-convergence for the epistemic realist that concerns us here apart from the variety of arguments from moral disagreement that have been provided since (and before) Mackie—all of which terminate in the conclusion that realism is false. See, e.g., Enoch’s (2009) comprehensive overview.

¹⁵ See, e.g., Rachels (1999: 23).

Indeed, this question gets to a general problem for any epistemic realist who wants to accept (T+CBP), while steering clear of the scepticism that looms if she rejects (R). The problem is best illustrated by considering a potential rejoinder to the above argument from non-convergence. The objection is that we so far have talked about responsiveness as an all-or-nothing affair, despite the fact that there might very well be *degrees* of responsiveness. Moreover, if we leave room for degrees of responsiveness, we might be able to identify a *particular* degree of responsiveness that is high enough to avoid aforementioned type of scepticism, yet low enough not to imply convergence—or so the objection goes. In visual terms, we get the following:

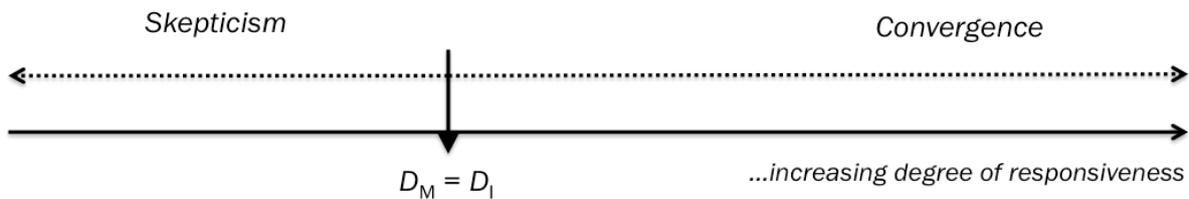


What the epistemic realist needs is for there to be some degree, D_R , that is high enough to avoid scepticism but low enough not to imply convergence. As for avoiding scepticism, this means that, if there is such a degree as D_R , it needs to fall on the right side of the point designated ' D_M ' in the above graph, pin-pointing the *minimal* degree of responsiveness needed to steer clear of scepticism. But D_R also needs to fall to the left of whatever point at which we can expect to see convergence. Is there such a point as D_R ? I will argue that there is not.

When setting up the argument, we need to remember that we have talked about being imperfectly responsive to (T+CBP) in terms of believing many true things, and for most of the things that one believes to be at least likely to be true. Let us refer to this *imperfect* degree of responsiveness as D_I . Given that degree of responsiveness, we might believe several things falsely, but for most of the things that we believe, we will have the probabilities on our side. As such, scepticism concerning claims about normativity does not seem an immediate threat. Any lower degree of responsiveness, however, would be one on which we do *not* believe many true things, and it is *not* the case that most of the things that we believe are at least likely to be true. At that point, scepti-

cism *vis-à-vis* the relevant domain would seem called for. What this means is that, in the above graph, D_1 is identical to D_M , i.e., the point designating the minimal degree of responsiveness needed to steer clear of aforementioned type of scepticism.

However, we also saw above that no more than responsiveness to degree D_1 is needed to generate the problem of non-convergence. If those theorizing about epistemic normativity believe many true things about epistemic normativity, and most of their theories about epistemic normativity are at least likely to be true, then we should expect far greater convergence upon a single theory of epistemic normativity than we are, in fact, seeing. In other words, D_1 is not only identical to D_M , but also constitutes a degree of responsiveness on which we can already expect convergence. This means that there is no point *in between* D_M and the point at which we can expect to see convergence, and that we may revise the above graph as follows:



As illustrated by the fact that D_R has dropped out of the picture, there being no point in between D_M and D_1 means that there is no degree of responsiveness high enough to avoid scepticism concerning claims about normativity, yet low enough not to imply convergence—contrary to the rejoinder under consideration.

In response to this, the realist may tinker, not with the degree of responsiveness, but with the proportion of theorists satisfying the relevant degree of responsiveness.¹⁶ More specifically, she may object that the above argument seems to be assuming that *every* theorist is responsive to degree D_1 , and suggest that we may explain non-convergence with reference to how only a proportion of theorists are responsive to degree D_1 . But how big of a proportion? If a *sizable* propor-

¹⁶ This was suggested to me by Matthew Chrisman.

tion, then we should expect convergence. If *less* than a sizable proportion, then the realist runs into the following dialectical problem:

There being less than a sizable proportion of people responsive to the relevant degree amounts to there being some particular subset of elite theorists about normativity who, unlike all the others, are likely to be right about what epistemic facts there are. But it would not be enough for the realist to argue that there *exists* such a subset, whoever happens to be in it. Her reason to postulate the existence of such a subset would be that, unless there is such a subset, then she has a hard time explaining non-convergence without falling prey to scepticism concerning claims about epistemic normativity. That reason would in no way impress her anti-realist interlocutors. So, the realist would need to argue for the existence of some *particular* subset, populated by some particular brand of theorists that have gotten the epistemic facts right. But to argue for *that* claim is simply to get involved in the object-level theorizing about normativity with respect to which we find great theoretical diversity—and no clear signs of a resolution.

One final way that the realist might try to get around the above challenges is as follows: Those theorizing about epistemic normativity have *knowledge-how* about how to conduct inquiry in an epistemically appropriate manner and, as such, steer clear of scepticism. But they do not have *knowledge-that* about the epistemic facts that they are, thereby, responsive to, which is why they are not converging on a single theory.¹⁷ This response fails, however. If those theorizing about epistemic normativity have know-how about how to conduct inquiry in an epistemically responsible manner, and a central (indeed, according to Lynch, a *constitutive* part) of so doing is to conduct inquiry in a manner that is responsive to (I+CBP), their beliefs about what they are inquiring into should at least be likely to be true. But in that case, we should (again) expect to see far more convergence than we are in fact seeing.

¹⁷ Thanks to Joëlle Proust for suggesting this response to me.

Consequently, the argument from non-convergence puts pressure on the epistemic realist to surrender either (ER), (P), (R), or (T+CBP). Assuming that the epistemic realist wants to remain a realist, she cannot surrender (ER). Assuming that she does not want to deny that theorizing about epistemic normativity involves assembling and exchanging epistemic reasons, she cannot deny (P). Assuming that she does not want to embrace scepticism about normativity claims, she cannot surrender (R)—nor can she reduce the degree of responsiveness at work in (R), or tinker with the proportion of theorists responsive to that degree, if what has just been argued is correct. Which brings us to the first implication of epistemic realism: (T+CBP) needs to be rejected. This should worry at least those who, like Lynch, take the truth-goal to be a *constitutive* fact about inquiry, and moreover want to criticize epistemic expressivism on the ground that it fails to respect that (alleged) fact.¹⁸

Where does this leave the realist? What got her in trouble in relation to the issue of non-convergence was the universally binding nature of categorical reasons, together with those reasons being truth-conducive. This suggests that the realist may escape the problem of non-convergence by rejecting the consequentialist framework, and specifically severing the connection between reasons and truth. In the next section, it will be suggested that, even if the epistem-

¹⁸ It should also worry those who take truth to be the norm of belief (e.g., Shah 2003), at least if we follow Lynch (2009) in taking such a ‘truth norm’ to apply also to inquiry, which in turn demarcates the domain of the epistemic (see footnote 5). To see why, notice first that defenders of the truth norm typically take the relevant correctness of believing in accordance with the truth norm to be *prescriptive*: one *should* (in an epistemic sense) believe that *p* if and only if it is true that *p*. Moreover, assume—as seems reasonable—that epistemic reasons and the epistemic ‘should’, or ‘should_E’ for short, are connected as follows: one should_E believe with epistemic reason, and should_E not believe without epistemic reason. In that case, to be responsive to reasons is to be responsive to the truth, and we run into the problem of non-convergence.

ic realist were willing to go down this epistemological route, it would not rid her theory of implausible implications.¹⁹

6. Widespread Irrationality or Cognitive-behavioural Incoherence

The epistemic realist takes there to be epistemic facts. If what was argued in the previous section is correct, taking (T+CBP) to be among those facts gives rise to a problem of non-convergence. But on one normative tradition, epistemic reasons have less to do with attaining goals like true belief—as the epistemological consequentialist would have it—and more with conducting inquiry in a manner that satisfies one’s epistemic duties (Clifford 1886; Feldman 2002; Berker 2013). If the epistemic realist were to side with this tradition, she can happily reject (T+CBP) and instead take the following to be an epistemic fact:

(D) To believe with reason is to believe in a manner that satisfies one’s epistemic duty.

If the defender of (D) is to avoid the problem of non-convergence, satisfying your epistemic duty cannot imply either that the relevant beliefs are true or even that they are *likely* to be true.

¹⁹ It was noted at the outset that the present paper is not an attempt to defend epistemic expressivism. For that reason, I have not investigated whether the epistemic expressivist avoids the problem of non-convergence. That said, it is worth noting that, since what gets the realist into trouble with non-convergence is partly the categorical nature of the reasons generated by epistemic facts, we should worry about the manner in which some expressivists (e.g., Blackburn and Gibbard) take their theories to leave room for cross-personal authority. To the extent that the relevant kind of authority mirrors the categorical force of reasons, in line with the general ambition of some expressivists (perhaps more so in Blackburn than in Gibbard) to mimic as many of the realist’s theoretical features as possible, it will run in to the same problems as the epistemic realist, if we import their suggestions into the domain of epistemic normativity.

However, if we understand what it is to satisfy one's epistemic duties in a manner that does not imply even a likelihood of truth, then we avoid the problem.

Questions have been raised about whether we can make sense of talk about duties in the epistemic realm (Alston 1988; Plantinga 1988). While epistemic deontologists, naturally, have argued that we can (e.g., Feldman 2000; Steup 2000), there is no need to settle the issue here. If we can't make sense of talk about epistemic duties, so much the worse for the epistemic realist, as far as her available options are concerned—and so much the better for her opponents.

So, let us assume that (D) is an epistemic fact. An epistemic non-deontologist (e.g., a consequentialist) will believe that there are no epistemic duties and as such deny (D). In believing that, she will also have to believe that, for any complete set of (putative) epistemic duties, what's in that set does *not* capture what it is to believe with reason. (If it did, there would be epistemic duties after all, which is what the non-deontologist is denying.²⁰) Let us refer to the complete set of epistemic duties that there in fact are as ED. It follows from what we said a moment ago that the non-deontologist would then have to believe the following about that set:

- (i) what's in ED does *not* capture what it is to believe with reason.

However, if epistemic norms are categorical, non-deontologists are nevertheless *committed* to reasoning in accordance with the reasons that there are simply in virtue of engaging in inquiry. Again, if epistemic reasons are categorical, everyone theorizing about epistemic normativity, and as such engaging in inquiry, is bound by the same reasons. Moreover, if (R) holds, deontologists and non-deontologists alike will be *responsive* to those reasons. On (D), that means that they will

²⁰ This is true in epistemology, at least, as epistemic consequentialists tend to reject the idea that there are epistemic duties. Not so in moral philosophy, where the consequentialist might be perfectly happy with taking there to be moral duties, so long as they can be spelled out in consequentialist terms.

be disposed to exhibit belief-forming patterns that are in accordance with ED. Since non-accidental, their disposition to that effect will be explainable. But in the case of someone believing (i)—someone who’s denying that what’s in ED captures what it is to believe with reason—what will that explanation be? It would have to be that she *on some level* takes it to be the case that the reasons she’s responding to are the right reasons. But on *what* level? One possibility is that she believes the following:

- (ii) what’s in ED *does* capture what it is to believe with reason.

But in that case, she is *inconsistent*. And notice that, while she will be thinking of what’s in ED *as* epistemic duties in relation to (i), for her to be inconsistent she does not need to do the same in connection with (ii). In relation to the latter, she might think of the set members of ED under any number of descriptions, including as epistemic duties, or more generally as norms of belief-formation, or (even more generically) as the ways to form belief. But so long as she is denying, by (i), that what’s in ED does not capture what it is to believe with reason, while, by (ii), also believing that it does, she’s inconsistent. And if consistency is a necessary condition on rationality, as it is typically taken to be, it follows that the non-deontologist is *irrational* with respect to the very set of commitments in virtue of which she is a non-deontologist, i.e., her denial of (D) and belief in (i).

But maybe this very verdict of irrationality should lead us to question whether the non-deontologist truly believes (ii)—particularly if we can find some *other*, perhaps sub-doxastic level on which she takes that what’s in ED captures what it is to believe with reason, with reference to which we can then explain her responsiveness. To this end, consider what Tamar Gendler (2008) has referred to as *alief*. Gendler introduces the notion of alief to account for certain instances of belief-behaviour discord, such as in cases where someone sincerely believes that the transparent

glass walkway projecting over the Grand Canyon will carry her, yet trembles while walking on it. The doxastically disavowed content guiding behaviour is *alieved* rather than *believed*.

With that distinction in mind, perhaps we can say that the non-deontologist believes (i) but simply alieves (ii). If we say that, we can explain why she is disposed to exhibit belief-forming patterns that are in accord with ED—after all, there’s a (sub-doxastic) level on which she takes what’s in that set to capture what it is to believe with reason—without having to attribute any irrationality on her part. At the same time, it doesn’t follow that all is well with such a person. Indeed, Gendler (2008) herself suggests that ‘if I believe that *P* and alieve that not-*P*, something is amiss’ (651). ‘I am violating certain norms of cognitive-behavioural coherence’ since ‘the well-functioning aliever is one whose aliefs and beliefs largely coincide’ (*ibid.*). So, while the move to alief may rule out inconsistency, it does not rule out what seems an objectionable form of cognitive-behavioural incoherence.

This suggests that rejecting (T+CBP) is not enough; even the mere combination of (ER), (P), and (R) is problematic. After all, the above argument in terms of (D) generalizes to every controversy regarding epistemic facts. Given the wide diversity of theories of epistemic normativity, and the fact that not all of them can be right if epistemic realism is true, the existence of epistemic facts implies that a great many philosophers—namely, all of those who fail to account for the ‘right’ epistemic reasons in their epistemological theories, *whatever those reasons as a matter of epistemic fact happen to be*—are wrong. And since everyone still is responsive to the ‘right’ reasons (as per R), the realist is committed to holding that a great many of those engaging in the relevant discussions not only are wrong but also are, if not irrational, then suffering from cognitive-behavioural incoherence.²¹

²¹ Say, for example, that the realist opts for a virtue epistemological understanding of epistemic normativity. The first thing she would need to do is *not* understand such virtue in consequentialist terms (see, e.g., Greco 2010), as this would have her face the problem of non-convergence of Section 5. However, even assuming a non-

7. Conclusion

Where can the epistemic realist go from here? She cannot reject (or even water down) (R), unless she wants to accept a far-reaching scepticism concerning claims about epistemic normativity. Nor can she deny (P), since it would be hard to make sense of what we are doing when engaging in a theoretical inquiry into epistemic normativity, if we are *not* thereby engaged in the business of assembling and exchanging epistemic reasons. But she also cannot reject (ER)—at least not if she wants to remain an epistemic realist. However, there are significant costs to holding on to (ER), (R), and (P). That package of commitments forces the epistemic realist to reject the idea that true belief is a central goal of inquiry and maintain that the diversity of views in the discussions about epistemic normativity is not merely a sign of many philosophers being mistaken about epistemic normativity, but also of widespread irrationality or cognitive-behavioural incoherence.

What has been argued in this paper does not establish that there is no possible way to flesh out the nature of the categorical reasons postulated by the epistemic realist that will enable her to avoid the problems presented above. But the fact that the two main options available—the consequentialist and the deontological—come with great costs for the realist does not bode well. Of course, the costs of epistemic expressivism might be even greater. But as noted at the outset, the purpose of this paper is not to defend epistemic expressivism, but to add to previous discussions about epistemic realism and expressivism by focusing on some substantive but heretofore ne-

consequentialist notion of virtue (e.g., along the lines of Baehr 2011), she would still face the problem that, if to believe with reason is to believe in a manner that has one manifest epistemic virtue, then any non-virtue epistemologist would be characterized by a state of mind analogous to the non-deontologist in the working example above, and as such either have to be considered irrational or to be suffering from cognitive-behavioral incoherence, on account of failing to accept the relevant virtue theory.

glected costs of the former. Only against a more complete picture of the relative costs and benefits of epistemic realism and expressivism, respectively, can we make an informed choice about what theory to opt for between the two.

References

- Ahlstrom-Vij, K. (2013). Moderate Epistemic Expressivism. *Philosophical Studies* 163 (2): 337-357.
- Ahlstrom-Vij, K. and J. Dunn (2014). A Defence of Epistemic Consequentialism. *Philosophical Quarterly* 64 (257): 541-551.
- Alston, W. (1985). Concepts of Epistemic Justification. *The Monist* 68 (1): 57-89.
- (1988). The Deontological Conception of Epistemic Justification. *Philosophical Perspectives* 2: 257-299.
- (2005). *Beyond "Justification": Dimensions of Epistemic Evaluation*. Ithaca: Cornell University Press.
- Baehr, J. (2011). *The Inquiring Mind: On Intellectual Virtues and Virtue Epistemology*. Oxford: Oxford University Press.
- Berker, S. (2013). The Rejection of Epistemic Consequentialism. *Philosophical Issues* 23: 363-387.
- Blackburn, S. (1993). *Essays in Quasi-Realism*. Oxford: Oxford University Press.
- (1998). *Ruling Passions: A Theory of Practical Reasoning*. Oxford: Oxford University Press.
- (1999). Is Objective Moral Justification Possible on a Quasi-realist Foundation? *Inquiry* 42: 213-228.
- Boisvert, D. (2008). Expressive-Assertivism. *Pacific Philosophical Quarterly* 89: 169-203.
- BonJour, L. (1985). *The Structure of Empirical Knowledge*. Cambridge: Harvard University Press.
- Bourget, D. & Chalmers, D. (2014). What do Philosophers Believe? *Philosophical Studies* 170: 465-500.
- Carter, J. A., & Chrisman, M. (2012). Is Epistemic Expressivism Incompatible with Inquiry? *Philosophical Studies* 159 (3): 323-339.
- Chrisman, M. (2007). From Epistemic Contextualism to Epistemic Expressivism. *Philosophical Studies* 135: 225-254.
- Clifford, W. K. (1886). The Ethics of Belief. In L. Stephen & F. Pollock (Eds.), *Lectures and Essays by the Late William Kingdon Clifford* (2nd ed., pp. 339–363). New York: Macmillan and Co.
- Cuneo, T. (2007). *The Normative Web: An Argument for Moral Realism*. Oxford: Oxford University Press.
- David, M. (2001). Truth as the Epistemic Goal. In M. Steup (Ed.), *Knowledge, Truth, and Duty* (pp. 151–169). New York: Oxford University Press.
- (2005). Truth as the Primary Epistemic Goal: A Working Hypothesis. In M. Steup & E. Sosa (Eds.), *Contemporary Debates in Epistemology* (pp. 296–312). Cambridge: Blackwell.

- DePaul, M. (1993). *Balance and Refinement: Beyond Coherentism in Moral Inquiry*. New York: Routledge.
- Dreier, J. (2004). Meta-ethics and the Problem of Creeping Minimalism. *Philosophical Perspectives* 18: 23-44.
- Feldman, R. (2000). The Ethics of Belief. *Philosophy and Phenomenological Research* 60(3): 667-695.
- (2002). Epistemological Duties. In P. Moser (Ed.), *The Oxford Handbook of Epistemology* (pp. 361–384). Oxford: Oxford University Press.
- Field, H. (2009). Epistemology Without Metaphysics. *Philosophical Studies* 143: 249-290.
- Foley, R. (1993). *Working Without a Net*. New York: Oxford University Press.
- Geach, P. (1965). Assertion. *Philosophical Review* 74: 449-465.
- Gendler, S. G. (2008). Alief and Belief. *The Journal of Philosophy* 105(10): 634-663.
- Gibbard, A. (1990). *Wise Choices, Apt Feelings*. Cambridge, MA: Harvard University Press.
- (2003). *Thinking How to Live*. Cambridge, MA: Harvard University Press.
- Goldman, A. (1992). Foundations of Social Epistemics. In his *Liaisons: Philosophy Meets the Cognitive and Social Sciences* (pp. 179–207). Cambridge, MA and London, England: The MIT Press.
- (1999). *Knowledge in a Social World*. Oxford: Oxford University Press.
- Greco, J. (2010). *Achieving Knowledge*. Cambridge: Cambridge University Press.
- Haack, S. (1993). *Evidence and Inquiry: Towards Reconstruction in Epistemology*. Oxford: Basil Blackwell.
- Kappel, K. (2011). Is Epistemic Expressivism Dialectically Incoherent? *Dialectica* 65(1): 49-69.
- Kvanvig, J. (2005). Truth is Not the Primary Epistemic Goal. In M. Steup & E. Sosa (Eds.), *Contemporary Debates in Epistemology* (pp. 285–296). Cambridge: Blackwell.
- Lynch, M. (2009). Truth, Value and Epistemic Expressivism. *Philosophy and Phenomenological Research* 79(1): 76-97.
- Mackie, J. L. (1977). *Ethics: Inventing Right and Wrong*. Harmondsworth: Penguin
- Moser, P. K. (1985). *Empirical Justification*. Dordrecht, Netherlands: Reidel.
- Plantinga, A. (1988). Chisholmian Internalism. In D. F. Austin (Ed.), *Philosophical Analysis: A Defense by Example* (pp. 127-151). Boston: Kluwer Academic Publishers.
- Rachels, J. (1999). *The Elements of Moral Philosophy* (3rd Ed.). New York: McGraw-Hill.
- Ridge, M. (2006). Ecumenical Expressivism: Finessing Frege. *Ethics* 116: 302-336.
- Shafer-Landau, R. (1994). Ethical Disagreement, Ethical Objectivism and Moral Indeterminacy. *Philosophy and Phenomenological Research* 54: 331-344.
- Shah, N. (2003). How Truth Governs Belief. *Philosophical Review* 112: 447-483.
- Steup, M. (2000). Doxastic Voluntarism and Epistemic Deontology. *Acta Analytica* 15(24): 25-56.

Forthcoming in C. McHugh, J. Way, and D. Whiting (eds), Metaepistemology, Oxford University Press.

Stich, S. (1990). *The Fragmentation of Reason: Preface to a Pragmatic Theory of Cognitive Evaluation*. Cambridge, MA: The MIT Press.

Tersman, F. (2006). *Moral Disagreement*. Cambridge: Cambridge University Press

Whitcomb, D. (2007). *Intellectual Goods: An Epistemic Value Theory*. Doctoral Dissertation, Rutgers University.