UNIVERSITY OF LEEDS

This is a repository copy of *Expressivism and the Reliability Challenge*.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/152189/>

Version: Accepted Version

Article:

Golub, C orcid.org/0000-0003-3934-5842 (2017) Expressivism and the Reliability Challenge. *Ethical Theory and Moral Practice*, 20 (4). pp. 797-811. ISSN 1386-2820

<https://doi.org/10.1007/s10677-017-9794-1>

© Springer Science+Business Media Dordrecht 2017. This is an author produced version of a paper published in *Ethical Theory and Moral Practice*. Uploaded in accordance with the publisher's self-archiving policy.

Reuse

Items deposited in White Rose Research Online are protected by copyright, with all rights reserved unless indicated otherwise. They may be downloaded and/or printed for private study, or other acts as permitted by national copyright laws. The publisher or other rights holders may allow further reproduction and re-use of the full text version. This is indicated by the licence information on the White Rose Research Online record for the item.

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.

eprints@whiterose.ac.uk
<https://eprints.whiterose.ac.uk/>

Expressivism and the Reliability Challenge

Camil Golub

The final publication is available at Springer via DOI: [10.1007/s10677-017-9794-1](https://doi.org/10.1007/s10677-017-9794-1). Please cite the published version.

Suppose there are objective normative facts and our beliefs about these facts are by-and-large true.¹

How did this come to happen? This is the reliability challenge to normative realism. It is significantly different from other epistemological challenges in that it assumes the truth of realism, and asks realists to *explain* the correlation between our beliefs and the facts, rather than to argue *that* we are reliable. In answering this demand, realists may rely on any normative or metaphysical claims that are part of their view.² The dialectical force of the challenge comes precisely from this generous setup: if even after being granted their own conception of the subject-matter of normative discourse, realists are unable to provide a good explanation of our reliability, or at least to convince us that an explanation is in principle available, then they are in a problematic epistemological position. Or so the argument goes.³

Importantly, the reliability challenge does not rely on any specific causal story meant to undermine realist commitments. This is not to say that facts about the origins of our normative beliefs may not be relevant in pressing this challenge. Evolutionary considerations, in particular, can be used to highlight one kind of explanation of our reliability that is available in other regions of thought, but not in the normative domain: namely, an evolutionary vindication of our reliability. Take the following passage from Jamie Dreier's (2012) paper on the reliability challenge:

-
- 1 I focus here on practical normativity. For brevity, I will use *normative* and *normativity* to mean *practically normative* and *practical normativity*, respectively, throughout the paper.
 - 2 The reliability challenge differs in these respects from the evolutionary debunking arguments due to Joyce (2001, 2006) and Street (2006), which do not allow realists to take for granted the truth of our beliefs about objective normative facts.
 - 3 Enoch (2011), Dreier (2012) and Joshua Schechter ("Does Expressivism Have an Epistemological Advantage over Realism?", unpublished manuscript) offer similar statements of the reliability challenge, inspired by Field's (1989) challenge to mathematical Platonism. How the failure to explain our reliability would undermine the realism is still a disputed question: it might provide a defeater for the justification we previously had for our realist beliefs, it might undermine the knowledge status of said beliefs, or it might just be that the absence of an explanation of our reliability is an important theoretical cost for realism. I will not attempt to settle this issue here.

“When we ask why human beings like us have largely accurate beliefs about how hot things are, we do not want a catalogue of the beliefs we have plus an explanation, piecemeal, of why this one and that one are true. Instead, we want an explanation for how it comes to be that so much of what we believe about temperatures is correct. Plainly, the satisfying explanation goes by way of our perceptual sensitivity to temperatures. And the accuracy of our perceptual sensitivity to temperatures is, presumably, due to natural selection. Our ancestors had to be pretty good at detecting the temperatures of objects in their environment, and organisms with inaccurate detectors fared worse and left fewer descendants. But there is no parallel explanation of why we reliably form normative beliefs.” (p. 271)

Dreier is not assuming that only an evolutionary explanation of reliability is acceptable, nor is he claiming that the actual history of normative thought undermines its epistemic status. Rather, his point is that realists should give us *some* explanation of our reliable access to normative facts that would be as unified and illuminating as the evolutionary explanations available in other domains of thought.

Anyone who claims that our normative beliefs are by-and-large true faces the demand to explain our reliability. But normative subjectivists will have an easier time answering it: it is no surprise that our beliefs line up with the normative facts if the latter depend on the former. Even certain realists might have a good response to this challenge: for instance, naturalist realists for whom normative facts play a causal role in the best explanation of our judgments about them. The reliability challenge is particularly pressing for *non-naturalist* normative realists, who typically hold that normative facts are both causally inefficacious and constitutively independent of our attitudes. This conjunction of claims might seem to rule out *in principle* the possibility of a good explanation of our reliability.

The question I want to explore is the following: does expressivism about normative discourse—the view that normative claims express desire-like mental states—provide any advantage in

the face of this challenge?⁴ Until recently, this question would have sounded odd. Expressivism was long thought to be immune to epistemological challenges. In particular, the reliability challenge does not seem to arise for a traditional expressivist view on which normative claims do not express truth-apt beliefs: on such a view, there seems to be no reliability *thesis* to begin with, let alone any question about how our reliability could be explained. But contemporary expressivists like Simon Blackburn (1993, 1998) and Allan Gibbard (2003) have adopted realist-sounding commitments—e.g., there are objective normative facts, and many of our beliefs about such facts are true—by relying on a minimalist account of truth, fact, etc. Thus, they have become vulnerable to the challenge of explaining the correlation between our beliefs and the facts. And given that these expressivists reject a naturalistic epistemology for normative thought, it is not clear whether they have a better answer to this challenge than non-naturalist realists.⁵

I will argue that expressivism does have something to offer in the face of the reliability challenge, and that these resources have not been properly explored yet.

Expressivists typically rely on two strategies when facing epistemological concerns. The first is to treat any question about truth and objectivity as internal to normative discourse, and defend our beliefs with first-order normative arguments. We can reassure ourselves of our reliability, they argue, by considering the normative verdicts we are most confident in: for instance, that nobody can have a reason to torture children for fun.⁶ However, this would not be an adequate response to the reliability challenge, which *assumes* that our beliefs are aligned with the facts, and asks about *why* this is so.⁷

4 This might be an advantage that expressivism has *over* non-naturalist realism, or one that it can provide *to* non-naturalist realism, depending on how the two views are related. More on this below.

5 Gibbard (2003) endorses certain claims associated with *naturalist* realism—e.g., that normative concepts are realized by natural properties, and that certain causal explanations involving normative facts are legitimate—but argues that there can be no naturalistic vindication of our normative judgments, remaining thus closer to non-naturalist realism. See also his claim in Gibbard (2011) that quasi-realism mimics the kind of realism that rejects a causal model for normative knowledge. It is an interesting question how exactly Gibbard's epistemological non-naturalism can be reconciled with his naturalism about properties, but I will assume that there is no deep instability in his views on these matters.

6 This is, for instance, how Gibbard (2011) and Blackburn ("Sharon Street on The Independent Normative Truth as Such", unpublished manuscript) respond to Street's (2011) evolutionary challenge to quasi-realism.

7 Schechter (manuscript) makes a similar point about the irrelevance of a normative defense of our reliability in this

The second strategy is to argue that the desire-like nature of normative judgments explains why we should be content with a defense of our reliability from a normative standpoint. This idea does seem relevant for the reliability challenge. Perhaps if we pay attention to the nature of normative thought we will see that the demand to explain our reliable access to normative facts is misguided. Jamie Dreier (2012) has recently offered a defense of Blackburn's and Gibbard's “quasi-realism” along these lines.

However, I will argue, the desire-like mental states posited by expressivists as the content of normative claims do not evade the reliability challenge. In a minimalist framework, these mental states are also truth-apt beliefs, so the demand to explain our reliability is intelligible and pressing.

In order to defuse this challenge, expressivists need to isolate a deeper sense in which normative thought is not representational. My proposal is that they rely on a *negative functional thesis* that has long been present in Blackburn's and Gibbard's views, but hasn't been put to work in epistemological contexts: the claim that normative thought does not have the function of tracking normative facts, or any other kind of facts, understood as a naturalistic thesis about why we use normative concepts. This thesis supports an argument to the effect that we should not demand an explanation of our access to objective normative facts that would mirror the explanations available in regions of thought that have a tracking function. We should be content instead with explanations of our reliability that take for granted the truth-conduciveness of certain psychological traits and belief-forming methods.

In developing my arguments, I will remain neutral on the issue of whether “quasi-realist” expressivism is an anti-realist view, or has converged with familiar versions of realism. I will only assume that expressivism is compatible with *objectivism*, i.e. with claiming that normative facts are counterfactually and constitutively independent of our attitudes.⁸ But at the end I will suggest that,

context. Moreover, the strategy of defending our reliability with normative arguments is not a specifically *expressivist* resource to begin with. Realists with little sympathy for expressivism can adopt a similar stance.

⁸ This leaves open that genuine realism might consist in more than mere objectivism. The question of how to draw a divide between quasi-realism and realism has become known as the *problem of creeping minimalism*, due to Dreier (2004), who proposes that the divide between the two views should be understood in explanatory terms. In Golub (forthcoming), I argue that there need be no explanatory conflict between quasi-realism and anything recognizable as a

whatever other reasons non-naturalist realists might have to reject expressivism, the negative functional thesis should not be one of them. Indeed, the appeal to this thesis in the face of the reliability challenge dovetails with claims that non-naturalists have made about the illegitimacy of applying naturalistic epistemological standards to our normative beliefs.

1. How the reliability challenge applies to quasi-realism

The expressivist project promised to allow us to engage in normative discourse in good faith while avoiding the epistemological issues faced by normative realism. If all we are doing in making normative claims is expressing certain conative attitudes—e.g., states of approval and disapproval—rather than beliefs that purport to describe the normative realm, then there seems to be no issue about how we can detect normative facts: there is no sense in which normative judgments aim at representing such facts, so it is misguided to ask how we can explain their reliability.⁹

However, recent expressivists have moved away from this traditional model. By adopting minimalist notions of truth, fact, etc.—notions that would apply to any superficially assertoric discourse in virtue of its syntax and logical discipline—expressivists such as Blackburn and Gibbard have come to embrace many commitments associated with realism. For instance, they will say that it is *true* that genocide is wrong, or that the claim that genocide is wrong *describes a fact*, taking these to be nothing more than endorsements of the first-order verdict that genocide is wrong. Moreover, they have given an expressivist account of objectivity claims. Our actual attitudes toward certain actions or states of affairs remain the same even when we consider scenarios in which our attitudes are different, they point out. For instance, we disapprove of genocide even when thinking of a world in which we

general notion of realism. See Blackburn (1993), Rosen (1998), and Fine (2001) for earlier discussions of similar issues.

9 Schechter (manuscript) argues that the possibility of our attitudes being mistaken, which even old-school expressivists can allow for, is enough to give rise to a reliability challenge: how is it that we came to have by-and-large correct attitudes? I bracket this issue here, and focus on how the challenge applies to minimalist expressivism.

ourselves do not disapprove of it. On the expressivist picture, it is such unconditional attitudes that we express when we say, e.g., that it is an objective fact that genocide is wrong.¹⁰

If this project of appropriating realist commitments in a minimalist framework is successful, then quasi-realists have earned the right to make the following claims: (i) there are objective normative facts, and (ii) our beliefs about such facts are by-and-large true. Therefore, quasi-realists are subject to the demand that they explain the correlation between our normative judgments and the normative facts.

Does anything in expressivism help in responding to this challenge? Contrary to what philosophers like Terence Cuneo (2008), Sharon Street (2011), and Joshua Schechter (MS) have argued, I believe that expressivism does provide such epistemological advantages.¹¹ But I also believe that these benefits have not been properly identified yet. That is, I agree with the critics of quasi-realism when they argue that minimalism about truth and facthood does not blunt the reliability challenge, and more importantly, that expressivists cannot dodge this challenge simply by appealing to the desire-like nature of normative thoughts. I will defend these points in what follows. Then I will develop my own proposal regarding the epistemological benefits of expressivism.

First, a few words on the irrelevance of minimalism. Quasi-realists cannot dissolve the reliability challenge by appealing to their deflationary understanding of truth, factuality, etc. As many have noted,¹² the same epistemological issues arise no matter whether we adopt an inflationary conception of truth and factuality or a minimalist one, on which these notions are mere devices for generalization and abbreviation. In particular, the correlation between our beliefs and the facts calls out for an explanation even if, in calling something a normative fact, we are only rehearsing a first-order normative verdict. Quasi-realists claim that, in many cases where we endorse such verdicts, we get it

10 See Blackburn (1993), Essay 8; Blackburn (1998), Ch. 9; or Gibbard (2003), Ch. 9. As already mentioned, I will remain neutral on whether this combination of expressivism and objectivism—which I call *quasi-realism*, following established usage—can still be considered an anti-realist metaethical view, or has converged with realism.

11 Cuneo and Schechter focus on the reliability challenge. Street extends her evolutionary debunking argument to quasi-realism, but some of the issues she raises carry over to the present discussion.

12 Field (1989) and Fine (2001), among others.

right. This is what they are asked to explain.¹³

To see even more clearly why minimalism is epistemologically idle, consider the fact that standard versions of minimalism are meant to apply in any domain of discourse, and deny the intelligibility of more robust notions of truth and fact. In other words, for a thoroughgoing minimalism, claims about the existence of truths and facts can *only* be understood as endorsements of first-order commitments, in the normative domain and elsewhere. So, if minimalism helped dissolve the reliability challenge, we would never face the demand to explain our reliable access to objective facts. But this cannot be true. Not only is the reliability challenge intelligible with respect to, say, our beliefs about shapes or temperatures, but it can be successfully answered in these cases.

If something in quasi-realism helps address the reliability challenge, this cannot be its metaphysical minimalism. It must be what expressivism reveals about the nature of normative thought.

2. Desire-like mental states and the reliability challenge

In dealing with epistemological issues, quasi-realists typically appeal to the desire-like character of normative judgments. It is because of this psychological fact, they argue, that a defense of our commitments with normative arguments is a good enough response to any epistemological challenge. Gibbard (2011), for instance, suggests that to demand more than a normative defense of our commitments is to fail to do justice to the expressive nature of normative claims. Similarly, in *Thinking How to Live* (2003) he argues that we should not be troubled by the lack of a vindication of our reliability of the sort available in naturalistic regions of thought, because normative judgments express

¹³ Bedke (2013) argues that quasi-realism avoids the epistemological troubles of non-naturalist realism because it does not incur any robust metaphysical commitments. But I do not see how quasi-realists can make this argument. From a thoroughly minimalist standpoint, the quasi-realists' commitments to truth and objectivity are as robust as they can intelligibly be, and give rise to the question of how it is that our normative beliefs are by-and-large true. (Moreover, even if minimalism provided some epistemological comfort in the face of this challenge, this would not have much to do with expressivism. Non-expressivist realists like Dworkin (1996) and Scanlon (2014) offer similar deflationary accounts of metaethical commitments, so the alleged benefits of minimalism would also be available to them.)

planning states, and coherent planning requires that we trust our own capacity to plan correctly.¹⁴

Dreier (2012) offers an argument along these lines focused on the reliability challenge: quasi-realism is immune to this challenge because it combines an attitudinal semantics for normative claims with minimalism about metaphysical and semantic notions. Here is a reconstruction of the argument:

1. If the conative attitudes expressed by normative claims are not subject to the reliability challenge, then merely adding on expressive devices for conveying such attitudes cannot give rise to a reliability challenge either.
2. The conative attitudes invoked in the expressivist account of normative discourse are not subject to the reliability challenge.
3. Therefore, quasi-realism is not subject to the reliability challenge.

Dreier makes a compelling case for (1). He asks us to imagine several steps in the evolution of a normative language—starting from a purely descriptive stage where speakers have conative attitudes but no linguistic devices for expressing them, and then adding expressive devices until we get to a normative language much like ours—and he challenges the opponent of quasi-realism to isolate the stage where the reliability challenge first arises. Given that none of the steps in the development of this language involves taking on any substantial commitments beyond those contained in the initial attitudes, he argues, there can be no reliability challenge at the fully expressive level that was not already there at the starting level. The epistemological standing of a set of attitudes cannot be undermined by the mere addition of expressive devices to our vocabulary.

As for (2), Dreier states that there can be no explanatory challenge to the ground-level attitudes posited by expressivists as the content of normative claims. The only challenges that can be raised for them, he says, are normative ones—e.g., a demand to justify a particular attitude, or the whole of our

¹⁴ See Gibbard (2003), p. 261. Blackburn (1998) makes similar points about the epistemological upshot of expressivism. Note, however, that what is meant to do epistemological work here is not, strictly speaking, the expressivist *semantics* or any other thesis about language, but a *psychological* thesis about normative thought.

evaluative standpoint—to which we can respond by defending our attitudes with normative arguments.

I agree with Dreier that the minimalist expansion into traditional realist territory cannot bring any *new* challenges for expressivists: if a reliability challenge arises at the fully expressive level, it must also apply to the conative attitudes expressed by normative claims. But this only shows that the quasi-realists' success in mimicking realism made visible a problem that they had faced all along: they did not *inherit* any epistemological difficulties from realists, but had their own reliability challenge to begin with, simply in virtue of having certain attitudes. If we accept the very minimalism that the quasi-realist project relies on, these attitudes *are* beliefs, even before being expressed publicly through declarative sentences, and adopting them amounts to thinking them *true*—and this is enough for the reliability challenge to arise, i.e. for (2) to be false. Quasi-realists still owe us an explanation of why we would have come to form normative beliefs that are by-and-large true.

Insisting on the desire-like character of normative judgments does not dissolve this explanatory demand, nor does it make it easier to answer. The fact that the truth-apt beliefs expressed by our normative claims have a motivational dimension is irrelevant for the question of how it is that these beliefs are by-and-large true.¹⁵

Here is what quasi-realists need to provide, if they are to weaken the reliability challenge. They should claim that normative thought—despite being truth-apt or representational in the sense captured by minimalism—is *not* representational at a deeper explanatory level. This negative thesis will not prevent the reliability challenge from arising. But it can be used to argue that we should not demand a particular kind of explanation of our reliability in the normative domain, akin to the explanations available in regions of thought that are representational in the relevant sense, and should be content instead with the explanations that quasi-realists can provide.

In order to make this argument, expressivists must look beyond their positive picture of

¹⁵ Thanks to Dan Waxman for helpful discussion on this issue. Cf. the arguments in Street (2010), pp. 377-9, and Schechter (manuscript) for similar theses.

normative thought. The thesis that normative judgments are desire-like does not entail a contrast in representationality between normative and ordinary descriptive thought—at least not in a minimalist framework, which undermines the traditional belief-desire dichotomy and other cognate distinctions. For all that this thesis tells us, normative judgments might still be representational in any relevant sense, the only thing separating them from ordinary descriptive judgments being their additional motivational valence. If expressivism is to be useful in the face of the reliability challenge, it must offer an independent *negative* thesis about normative thought. In the next section, I isolate such a negative thesis in Blackburn's and Gibbard's expressivism, and articulate its epistemological value.

Dreier himself acknowledges that more needs to be said about the epistemological advantages of quasi-realism.¹⁶ Although he takes himself to have shown *that* quasi-realism is immune to the reliability challenge, he looks for an explanation of *why* this is so, and tentatively suggests that his account of the divide between quasi-realism and realism, proposed in Dreier (2004), might offer an answer: perhaps expressivists are not vulnerable to the reliability challenge because they do without normative facts in explaining what it is to have a normative thought.

Now, we have seen that quasi-realism is not *immune* to the reliability challenge. Moreover, it is disputable whether Dreier's (2004) proposal delivers a stable divide between quasi-realism and realism.¹⁷ Nonetheless, I believe Dreier's suggestion is on the right track. We can find epistemological benefits in expressivism if we look at the explanatory roles that expressivists *deny* to normative facts.

3. The negative functional thesis and its upshot

Expressivism is more than a semantic program, at least in the versions developed by Blackburn and Gibbard. The attitudinal semantics is only one piece in a broader project of trying to understand the function of normative discourse in naturalistic terms—that is, to understand why creatures like us,

¹⁶ See Dreier (2012), pp. 285-6.

¹⁷ See fn. 8.

living in a natural world seemingly devoid of normative facts, would have a use for normative concepts. Answering this question is the fundamental task that Blackburn and Gibbard set for themselves. The answer they provide is meant to be therapeutic, by revealing that there is nothing defective about our normative practice, and we can keep engaging in it in good faith: normative discourse does not have the function of tracking any kind of facts, but it has a different, attitude-expressive function, which it serves well.¹⁸ While expressivists and their allies have focused so far on the positive part of this functional picture in epistemological contexts, it is the negative functional thesis of expressivism that helps in the face of the reliability challenge.

The thesis that normative concepts do not have the function of tracking normative facts, understood as a naturalistic claim about what explains our normative thought and practice, has been at the center of Gibbard's and Blackburn's projects. Gibbard (1990) argued at length for the claim that normative thought is not *naturally representational* in the way that ordinary descriptive thought is: there is no domain of facts such that something is a normative thought only if it naturally represents a fact from that domain.¹⁹ Blackburn too has insisted throughout his work that the best explanation of our normative beliefs and linguistic practice does not involve a causal role for normative facts—in other words, that normative thought and language do not exist because they track certain aspects of reality.²⁰ Moreover, Gibbard and Blackburn articulate the contrast between their views and normative realism precisely by appealing to such negative claims about explanation: expressivists do without normative facts in their explanations of normative thought and language, they say, while realists explain our normative beliefs and linguistic practice by appeal to normative facts.²¹

18 More precisely, normative terms have the function of conveying and stabilizing our action-guiding attitudes, for the purposes of solving coordination problems in social contexts. See, e.g., the first few chapters of Gibbard (1990).

19 Gibbard (1990), Ch. 6. Gibbard used this idea of natural representation, defined in terms of evolutionary function, to articulate an inflationary notion of fact, and argue that normative thought is non-factual. But the functional claims of expressivism can be adopted by those who accept minimalism about factuality, as shown by Gibbard's own evolution.

20 See, e.g., Blackburn (1984), p. 169, or Blackburn (1993), p. 7.

21 Blackburn (1984) attributes to realists the thesis that values “are themselves part of the genesis of our beliefs. It would be because values, etc. are distributed in some way around the world, and because we are capable of reacting to them (...) that we moralize as we do.” (pp. 181-2) Similarly, Blackburn (1993) defines realism as follows: “[T]he existence of

Putting aside whether we can isolate a conflict between expressivism and realism if we focus on issues about explanation and function,²² what matters for our current purposes is that the following is a core commitment of Blackburn's and Gibbard's expressivism:

Negative Functional Thesis Normative concepts are not naturally representational, i.e. do not have the function of tracking normative facts, or any other kind of facts.

Given their focus on historical explanations of normative thought, particularly on evolutionary explanations, Blackburn and Gibbard arguably understand this as a claim about *etiological* function: *x* is among the functions of a mental or linguistic item, in this sense, if being or doing *x* plays a key role in the best causal explanation of the emergence and proliferation of items of that kind. *Tracking* stands for the kind of co-variation between mental or linguistic items and features of the world that would explain the emergence and proliferation of items of that kind as representational devices. The Negative Functional Thesis thus understood denies that the mental mechanisms that produce normative concepts and beliefs were selected for producing mental items of that kind because they tracked something.

However, the Negative Functional Thesis can be given a broader reading, on which its fate is not tied to the success of etiological accounts of representational content. The philosophical search for a naturalistic theory of representational content is still in its early stages, and while etiological theories have been prominent in recent decades,²³ the jury is still out on whether they can successfully address the challenges that they face.²⁴ One might be optimistic that a suitably refined etiological account will overcome these obstacles, but I do not believe the Negative Functional Thesis should be hostage to this debate. It is more fruitful to think of expressivism as denying the natural representationality of

facts explains the way in which our knowledge expands and progresses: here an explanatory role seems to carry with it an ontological commitment which (...) is surely problematic to the quasi-realist." (p. 18)

22 Many realists might not assign a tracking function to normative concepts either: for instance, non-naturalist realists who believe that normative facts are not causally efficacious. I discuss this issue in Golub (forthcoming).

23 See Dretske (1981, 1988) and Millikan (1984), among others.

24 See Fodor (1990) for a discussion of challenges to etiological theories of content, e.g. the seeming possibility of mental items that have content but no history, or how to distinguish content-determining from non-content-determining causes.

normative thought, in the sense captured by the best theory of naturalized representational content, whatever that theory turns out to be.²⁵

This Negative Functional Thesis can be used to weaken the reliability challenge. It supports an argument to the effect that we should not demand a certain kind of explanation of our reliable access to normative facts, because this would saddle normative thought with a function that it does not have.

First, note that the Negative Functional Thesis does not only distinguish between normative thought and regions of thought that are naturally representational and succeed in fulfilling their tracking function, e.g. our beliefs about shapes or temperatures. It also sets normative thought apart from regions that have a tracking function yet do not track anything. For instance, astrological facts play no causal role in an explanation of why people use astrological concepts, because there are no such facts. Nevertheless, astrological concepts arguably have a tracking function: it is plausible that they are the outputs of mental modules the general function of which is to track facts, such as a module for detecting agency or for identifying the causes of natural phenomena. In contrast, the best explanation of why we use normative concepts, according to expressivism, does not depict said concepts as having the function of tracking any kind of facts.²⁶ This contrast between naturally representational regions of thought that fail to track anything and non-tracking regions of thought will be relevant to my argument.

Consider now the explanations of our reliability that quasi-realists *can* offer. Not only can they defend our reliability with normative arguments, but they can offer substantive explanations of why we ended up with largely true beliefs. For instance, they may hold that we are likely to form true normative beliefs when we pay attention to the relevant non-normative facts, deliberate carefully and our

25 Sinclair (2006) proposes a similar understanding of the negative functional thesis of expressivism, on which the relevant notion of function might involve evolutionary explanations or explanations of current practices. The idea that this functional thesis is not a mere add-on to expressivism but should be seen as one of its central tenets—indeed, as the best way to preserve its negative insights in a minimalist framework—is defended in O’Leary-Hawthorne and Price (1996).

26 To be clear, the Negative Functional Thesis is compatible with holding that our normative claims *represent* normative facts, in the sense captured by minimalism. On a minimalist account, to say that an utterance of “Genocide is wrong” represents the fact that genocide is wrong is not to say anything about why we use this linguistic item.

reasoning is not obstructed by pernicious biases. They may point to the kind of upbringing that fosters good judgment on normative matters. Or they may even appeal to the evaluative tendencies that are part of our evolutionary heritage, such as the tendency to disvalue pain or to value reciprocity: it is no surprise that we ended up largely reliable in our normative judgments, they might say, given that these tendencies bring us closer to the truth, as a matter of normative fact.²⁷

These explanations are unifying and illuminating. They do not consist in simply explaining why we have the beliefs that we do and then independently defending each belief. They explain the correlation between our beliefs and the facts by appealing to more general connections between our belief-forming methods and normative truths.

Nonetheless, many philosophers will not be content with such explanations. Quasi-realists, the objection goes, still leave unexplained certain connections between our psychology and the facts: we have yet to be told why certain features of our belief-forming processes are truth-conducive, for instance why our inherited evaluative tendencies bring us in line with basic normative truths. In contrast, in other regions of thought, we can provide a more robust explanation of our reliability, which does not treat any such connection between our beliefs and their subject matter as a brute fact.

Quasi-realists should acknowledge this explanatory asymmetry between normative thought and other regions of thought, but argue that we should not be troubled by it: this asymmetry is precisely what we ought to expect, given the non-tracking function of normative thought, and it is misguided to demand a more robust explanation of our reliable access to objective normative facts, as such expectations are only appropriate for regions of thought that have a tracking function.

Let us consider again how much the kind of explanation of our reliability suggested above would accomplish, and what might be the source of any lingering dissatisfaction with such explanations. Take the explanation that appeals to evaluative tendencies that are arguably part of our

²⁷ Such explanations are equally available to non-expressivist realists. See, in particular, Schafer (2010), Enoch (2011), Vavova (2015), and Clarke-Doane (2015) for explanations that appeal to our evolutionary heritage.

evolutionary heritage, such as disvaluing pain and valuing reciprocity. Call it the *evolutionary explanation*.²⁸ Why would it not be enough to explain our reliability by relying on the premise that valuing reciprocity and disvaluing pain are truth-conducive attitudes?

Again, the complaint cannot be that this explanation does not account for the vast and striking correlation between our beliefs and the facts. The evolutionary explanation provides a unifying account of why that broad correlation holds, by appealing to a much less striking alignment between features of our psychology, e.g. our valuing of reciprocity and disvaluing of pain, and certain normative truths, e.g. reciprocity is good and pain is bad.

Nor can the dissatisfaction come from a sense that quasi-realists are begging the question. Remember the setup of the reliability challenge: quasi-realists are allowed to appeal to normative theses that are part of their view. Their task is not to convince us *that* we have true beliefs, but to explain our reliability. What their opponents must claim is that, even if we accept the relevant connections between our evaluative tendencies and the normative truths from a normative standpoint, these connections are still *striking*, or call out for an explanation.

Someone might suggest that the problem is the *modal insecurity* of these connections between our psychology and the facts: the facts could have easily been different, or we could have easily had different evaluative tendencies, so for all that the evolutionary explanation tells us, it is only a matter of *luck* that we ended up reliable in our judgments. That is, it might be thought that any good explanation of our reliability should establish that our beliefs are *sensitive* to the facts and *safe*—i.e., that they would still be true in nearby worlds where the facts were different, or where the beliefs themselves were different—and the evolutionary explanation fails to achieve this, given that it does not posit either a causal or a constitutive link between our inherited evaluative tendencies and the facts.²⁹

28 I am not suggesting that this explanation amounts to an evolutionary vindication of our normative beliefs of the sort available for, say, our beliefs about shapes or temperatures.

29 Several contributors to the debate on the reliability challenge hold that a good explanation of reliability should establish modal security thus understood. See Clarke-Doane (2015), Barkhausen (2016), and Warren (forthcoming).

Now, it is not obvious that establishing the modal security of our beliefs is a necessary condition for a good explanation of reliability.³⁰ The fact that we have by-and-large true beliefs may well be highly contingent, but that by itself does not entail that it cannot be *explained*. After all, we can explain many unlikely events in the natural world.³¹ But even if this were a legitimate *desideratum*, the evolutionary explanation seems to deliver as much modal security as we can reasonably expect. First, sensitivity is secured by appealing to the metaphysical necessity of normative facts: if reciprocity is indeed valuable, then there is no nearby world in which reciprocity is *not* valuable and yet we still value it.³² As for the safety condition, we can rely on the very fact that our beliefs were influenced by evolutionary factors to argue that we could not have easily had very different beliefs: the relevant nearby worlds are arguably worlds in which we would have been on the same evolutionary pathway, and hence would have still valued reciprocity and disvalued pain.³³

What else could explain the dissatisfaction with the evolutionary explanation, if it is not a concern about how striking, implausible, or lucky it is that we are reliable? The only remaining option, I believe, is a lingering sense that quasi-realists have not yet provided an explanation that reveals how our valuing of reciprocity and disvaluing of pain are *connected* with facts about what is valuable.³⁴

30 Notably, Dreier (2012) and Schechter (manuscript) do not rely on this modal reading of the reliability challenge. Field (1989), in his reliability challenge to mathematical Platonism, only appeals to sensitivity as a relevant modal constraint.

31 To take a classic example from the philosophy of science: we can explain the occurrence of paresis by appealing to the fact that the patient had latent untreated syphilis, even though the chance that someone with latent untreated syphilis will develop paresis is about one fourth.

32 To be clear, this is not a specific virtue of the evolutionary explanation. The metaphysical necessity of normative facts ensures that the sensitivity constraint would be met by *any* explanation of our reliability.

33 Clarke-Doane (2015) makes both of these arguments. As he points out, the set of worlds relevant for assessing modal security cannot be too wide, lest we end up counting most of our ordinary beliefs as insecure.

34 An anonymous reviewer suggests that quasi-realists might have more to say in response: if we endorse Gibbard's view that normative concepts are realized by natural properties and normative facts can play a role in causal explanations (see fn. 5), then the evolutionary explanation need not bottom out in brute connections between our psychology and the facts. We can claim that our evaluative tendencies are, at least in part, the result of interactions with normative facts—for instance, that the natural property that constitutes goodness causally influenced our tendency to value reciprocity. I cannot explore this suggestion in depth, but I suspect that neither Gibbard nor his opponents would think of these causal explanations as fully addressing any lingering concerns driving the reliability challenge. If normative thought does not have the function of tracking any facts, then the best explanation of why we have certain evaluative tendencies—one that would apply in a wide range of evolutionary pathways—does not involve normative facts. Even if our evaluative tendencies were causally influenced by normative facts, this does not give us sufficient insight into why we ended up on a truth-conducive path, given that other evaluative tendencies would have been equally compatible with the coordinating function of normative thought. Only a tracking account of normative thought could meet this explanatory demand.

Quasi-realists, however, can offer a debunking diagnosis of this epistemological anxiety: it comes from being in the grip of a tracking epistemological model, which is illegitimately imported to the normative domain. In other regions of thought, e.g. our beliefs about temperatures, we can describe how our mental states and cognitive mechanisms are causally connected with external objective facts, and as a result we can say that we have certain beliefs *because* those beliefs are true. The absence of a similar story about our access to normative facts may seem troubling in comparison. We should not be troubled by this explanatory asymmetry, however: once we realize that normative thought does not have a tracking function, we should not demand that our beliefs about objective normative facts be vindicated in the same way that our beliefs about temperatures can.

To be clear, this is not an argument that the reliability challenge is inappropriate in the normative domain. The demand to explain why our normative beliefs are by-and-large true is intelligible, and quasi-realists can meet it successfully, e.g. through the evolutionary explanation. What they cannot offer is an explanation of why our belief-forming methods are truth-conducive that would resemble tracking accounts available in other regions of thought. My thesis is that we should not demand such an explanation of our reliability, given that normative thought does not have a tracking function.

Someone might object that the demand for a more robust explanation of our reliability is not based on ascribing a tracking function to normative thought, and does not assume that a tracking account is the only way to uphold the epistemological credentials of a set of beliefs. After all, a normative subjectivist or constructivist need not offer a tracking account of our reliability. She can appeal to *constitutive* connections between our attitudes and the facts: if normative facts depend on our attitudes, then no matter what evaluative tendencies we were endowed with, we would be on the right track in our normative judgments, as long as we were reasonably cognizant of our own psychology. This explanation reveals *connections* between our psychology and the facts where quasi-realists leave

unsettling gaps, it might be argued, and this should count against quasi-realism.

However, why should we accept that the mere positing of constitutive connections between our psychology and facts is a theoretical virtue, independently of any legitimate epistemological concerns? Again, quasi-realists can offer illuminating, non-circular, and modally robust explanations of our reliability. The sense that their explanations are incomplete comes, I believe, from the undue influence of tracking accounts of reliability on our explanatory expectations. Subjectivism only fills an illusory explanatory gap. This is what the Negative Functional Thesis helps us see.

A different objection might come from the opposite direction. The explanations available to quasi-realists are indeed unified, illuminating, and modally robust, and no further explanatory demand is legitimate, someone might argue, but this has nothing to do with the Negative Functional Thesis.

However, this cannot be the end of the story. In other domains of thought we *are* justified in asking for a deeper explanation of our reliability, namely a tracking explanation. Take again astrological thought. Imagine an explanation of reliability in this domain modeled on the explanations available to normative quasi-realists: it would appeal to some psychological traits that underlie the formation of astrological beliefs, and to their stability in nearby-worlds.³⁵ This explanation would also take for granted that the relevant traits correspond to basic astrological facts, without claiming that the former are causally explained by the latter. Now, such an explanation would not be good enough. Only a depiction of the causal connections between astrological facts and our cognitive faculties would assuage our epistemological worries about astrology.

Why is the normative case different? The Negative Functional Thesis provides an answer. It allows us to distinguish between naturally representational systems that systematically misfire and non-tracking regions of thought. We have thus a principled reason to treat the demand for a tracking explanation of our reliability as legitimate for astrological beliefs, but not for normative beliefs.

³⁵ Arguably, astrological beliefs are in fact more modally insecure than normative beliefs, but we can easily imagine a scenario in which they were just as stable, without being caused by astrological facts.

Perhaps this appeal to the Negative Functional Thesis in the face of the reliability challenge will not make any new converts to the objectivist cause. But for those of us with objectivist leanings, this thesis provides a new and illuminating way of reassuring ourselves that we should not be troubled by the explanatory asymmetry between normative and ordinary descriptive thought.

This quasi-realist response to the reliability challenge resembles what some non-naturalist realists have said in the face of similar challenges: that our reliability should be judged and accounted for by the autonomous standards of the normative domain, and not by importing epistemological models appropriate for naturalistic thought. For instance, Ronald Dworkin (1996) argues that the lack of an “external” explanation of our reliability should do nothing to undermine our fundamental moral commitments. Such an explanatory constraint on the justification of our commitments is appropriate for beliefs about the physical world, given their content—“they are beliefs about objects and events that can interact causally with the human nervous system” (p. 119)—but is not invited in any way by the content of moral beliefs: morality is an autonomous domain that makes no causal claims, he argues, so its epistemology should not be constrained by a model imported from natural science. Similarly, T.S. Scanlon (2014) argues that the epistemology of the normative domain only needs to be *internally appropriate* and should not be expected to resemble the naturalistic epistemological model for discovering truths about the physical world, given that when we engage in normative thought we do not aim at “being in touch with the facts in a causal sense” (p. 85).³⁶

The Negative Functional Thesis of expressivism deepens and provides an empirical grounding for these realist arguments. It fleshes out in naturalistic terms the sense in which normative thought does not aim at the same kind of understanding as naturalistic regions of thought, and thus the reason why it should not be held by the same epistemological standards. Rather than talking about the nature of normative facts themselves, as realists usually do, the quasi-realist response to the reliability

³⁶ See also Nagel (1986), p. 141.

challenge focuses on the way we represent such facts, or rather on the way we do *not* represent them. Far from being in tension with normative objectivism, such a naturalistic inquiry into the function of normative thought can dissipate the epistemological anxiety caused by the lack of a certain kind of explanation of our access to objective normative facts. It makes good on the realist ideas about the autonomy of the normative, and the narrow scientism presupposed by the demand for an “external” explanation of our reliability. Thus, we find here a point of convergence between quasi-realism and non-naturalist realism, and a way in which realists can find epistemological benefits in expressivism.

4. Conclusion

I have argued that an important piece of the expressivist project—namely, the thesis that normative thought does not have the function of tracking any kind of facts—can be used to weaken the reliability challenge to quasi-realism. Importantly, these epistemological benefits do not hinge on the success of the semantic project associated with expressivism. Even if expressivists ultimately fail to offer a compositional psychologicistic semantics for normative discourse, their Negative Functional Thesis could still hold and be useful in dissipating epistemological worries about our objectivist commitments.

Acknowledgments

I am grateful to Brian Ballard, Michelle Dyke, Tom Nagel, Daniel Skibra, David Velleman, several anonymous referees, the editors of this special issue, and especially Hartry Field, Josh Schechter, and Sharon Street, for their helpful comments on previous versions of this paper. Many thanks also to audiences at New York University, Bilkent University, and the 2016 Conference on Objectivity in Ethics at the University of Utrecht for their useful feedback. Finally, I have greatly benefited from discussions with Max Barkhausen, Harjit Bhogal, Paul Boghossian, Jared Warren, and Mike Zhao, and especially from many conversations with Dan Waxman.

References

- Barkhausen M (2016) Reductionist moral realism and the contingency of moral evolution. *Ethics* 126(3):662-689
- Bedke M (2013) Ethics makes strange bedfellows: intuitions and quasi-realism. In Haug M (ed) *Philosophical methodology: the armchair or the laboratory?* Routledge
- Blackburn S (1984) *Spreading the word*. Clarendon Press
- Blackburn S (1993) *Essays in quasi-realism*. Oxford University Press
- Blackburn S (1998) *Ruling passions*. Oxford University Press
- Clarke-Doane J (2015) Justification and explanation in mathematics and morality. In: Shafer-Landau R (ed) *Oxford studies in metaethics*, vol 10. Oxford University Press
- Cuneo T (2008) Moral realism, quasi-realism, and skepticism. In: Greco J (ed) *The Oxford handbook of skepticism*. Oxford University Press
- Dreier J (2004) Meta-ethics and the problem of creeping minimalism. *Philos Perspectives* 18:23-44
- Dreier J (2012) Quasi-realism and the problem of unexplained coincidence. *Analytic Philos* 53:269-287
- Dretske F (1981) *Knowledge and the flow of information*. MIT Press
- Dretske F (1988) *Explaining behavior: reasons in a world of causes*. MIT Press
- Dworkin R (1996) Objectivity and truth: you'd better believe it. *Philos & Public Affairs* 25:87-139.
- Enoch D (2011) *Taking morality seriously: a defense of robust realism*. Oxford University Press
- Field F (1989) *Realism, mathematics and modality*. Basil Blackwell
- Fine K (2001) The question of realism. *Philosophers' Imprint* 1:1-30.
- Fodor JA (1990) *A theory of content and other essays*. MIT Press
- Gibbard A (1990) *Wise choices, apt feelings: a theory of normative judgment*. Harvard University Press

- Gibbard A (2003) *Thinking how to live*. Harvard University Press
- Gibbard A (2011) How much realism? Evolved thinkers and normative concepts. In: Shafer-Landau R (ed) *Oxford studies in metaethics*, vol 6. Oxford University Press, pp 33-51
- Golub C (forthcoming) Expressivism and realist explanations. *Philos Stud*
- Joyce R (2001) *The myth of morality*. Cambridge University Press
- Joyce R (2006) *The evolution of morality*. MIT Press
- Millikan RG (1984) *Language, thought, and other biological categories*. MIT Press
- Nagel T (1986) *The view from nowhere*. Oxford University Press
- O'Leary-Hawthorne J, Price H (1996) How to stand up for non-cognitivists. *Australas J Philos* 74:275-292
- Rosen G (1998) Blackburn's Essays in Quasi-realism. *Noûs* 32(3):386-405
- Scanlon TM (2014) *Being realistic about reasons*. Oxford University Press
- Schafer K (2010) Evolution and normative scepticism. *Australas J Philos* 88:471-488
- Sinclair N (2006) The moral belief problem. *Ratio* 19:249-260
- Street S (2006) A Darwinian dilemma for realist theories of value. *Philos Stud* 127:109-166
- Street S (2010) What is constructivism in ethics and metaethics? *Philos Compass* 5:363-384
- Street S (2011) Mind-independence without the mystery: why quasi-realists can't have it both ways. In: Shafer-Landau R (ed) *Oxford Studies in Metaethics*, vol 6. Oxford University Press
- Vavova K (2015) Evolutionary debunking of moral realism. *Philos Compass* 10:104-116
- Warren J (forthcoming) Epistemology versus non-causal realism. *Synthese*