

# Two Accounts of Moral Objectivity: from Attitude-Independence to Standpoint-Invariance

Jeroen Hopster<sup>1</sup>

Accepted: 3 March 2017 / Published online: 31 March 2017

© The Author(s) 2017. This article is published with open access at Springerlink.com

**Abstract** How should we understand the notion of moral objectivity? Metaethical positions that vindicate morality’s objective appearance are often associated with moral realism. On a realist construal, moral objectivity is understood in terms of mind-, stance-, or attitude-independence. But realism is not the only game in town for moral objectivists. On an antirealist construal, morality’s objective features are understood *in virtue of* our attitudes. In this paper I aim to develop this antirealist construal of moral objectivity in further detail, and to make its metaphysical commitments explicit. I do so by building on Sharon Street’s version of “Humean Constructivism”. Instead of the realist notion of *attitude-independence*, the antirealist account of moral objectivity that I articulate centres on the notion of *standpoint-invariance*. While constructivists have been criticized for compromising on the issue of moral objectivity, I make a preliminary case for the thesis that, armed with the notion of standpoint-invariance, constructivists have resources to vindicate an account of objectivity with just the right strength, given the commitments of ordinary moral thought and practice. In support of this thesis I highlight recent experimental findings about folk moral objectivism. Empirical observations about the nature of moral discourse have traditionally been taken to give *prima facie* support to moral realism. I argue, by contrast, that from what we can tell from our current experimental understanding, antirealists can capture the commitments of ordinary discourse at least as well as realists can.

**Keywords** Objectivity · Humean constructivism · Moral realism · Ideally coherent eccentrics · Folk objectivism · Moral relativism · Moral metaphysics

## 1 Introduction

There is a broad consensus among philosophers that morality has certain objective features, and that it is a metaethical *desideratum* to account for this. For example, many judgements in

---

✉ Jeroen Hopster  
j.k.g.hopster@uu.nl

<sup>1</sup> Utrecht University, Utrecht, Netherlands

ordinary moral discourse purport to be objective; many metaethicists, in turn, regard this as a datum that their theories should accommodate. *How* to accommodate morality's objective aspirations, however, is an issue of ongoing debate, fuelled at least in part by controversy over how the notion of 'moral objectivity' should be understood.<sup>1</sup> Objectivity has various connotations; its precise meaning may differ across contexts and domains. Nichols (2014, p. 734) contends that "[g]iving a precise characterization of objectivity is itself a major philosophical endeavour." Timmons (2010, p. 544) even characterizes the issue of moral objectivity as "the central topic of metaethics."

On one prominent account, moral objectivity is understood in terms of mind-independence, stance-independence, or – my term of choice in this article – *attitude-independence*. Call this a realist understanding of moral objectivity (e.g. Shafer-Landau 2003).<sup>2</sup> On a realist understanding, to say that a moral judgement purports to be objective is to say that this judgement purports to have an attitude-independent truth-value (Finlay 2007, p. 822). Moral judgements are not like judgements of taste, but more akin to statements of fact (Enoch 2014). There are facts about what is right and wrong, good and bad, and these facts are objective in the sense of being attitude-independent. Consider the claim that the Earth revolves around the sun: this claim is true, it states a fact, and this truth or fact is fully independent of what any agent thinks or feels about it. Similarly, moral realists maintain that moral truths or facts are fully independent of the attitudes of any agent.

A recurring criticism against a realist understanding of objectivity, which has been pushed both in the domains of ethics and mathematics (Clarke-Doane 2014), is that it raises epistemic quandaries. For example, Sharon Street (2006, 2016a) has forcefully argued that if moral truths or facts are fully attitude-independent, then we have no guarantee that our moral judgements reliably track these truths or facts. Other critics, such as John Mackie (1977), argue that a realist view is not only epistemically, but also metaphysically problematic.<sup>3</sup> However, whatever difficulties moral realism might face, there is a general consensus that at least with concern to the aim of capturing the moral appearances, it does a good job. For instance, the realist has resources to argue that in cases of moral disagreement, at least one of the disagreeing parties must be mistaken, since his or her judgement does not correspond with the attitude-independent moral facts or truths. *Prima facie*, this seems to accord with people's intuitions. Consider an imaginary agent who values an action that many people would regard as decidedly immoral – say an agent who values torturing innocent human beings for fun. Presumably, many people would want to claim that this agent is morally wicked, and objectively so. Torturing innocent people for fun is morally wrong, no matter what this agent thinks or feels about it, or so our intuition tells us. Since

<sup>1</sup> Another potential source of controversy is the vagueness of 'morality', thus characterized. I give a more precise characterization of morality's objective features in **sections 2 and 3**.

<sup>2</sup> Although moral realism is a metaethical term of art, philosophers tend to characterize it in slightly different ways. As a result, the variety of metaethical claims labelled 'realist' is rather diverse (see Finlay 2007 for an overview). For a helpful recent characterization of moral realism, see Leibowitz and Sinclair 2016, especially pp. 4–5, 60–61, 126–127. In footnote 5 I explain how I draw the distinction between realism and antirealism in this article. Note that what I call a realist view of moral objectivity, according to which moral objectivity is understood in terms of attitude-independence, sometimes goes by the name *robust realism*.

<sup>3</sup> Error theorists maintain that a substantial subset of moral judgements (e.g. all affirmative moral judgements) are systematically false or systematically unjustified, since their truth-makers – objective properties that are at the same time prescriptive – are too queer to be instantiated (cf. Mackie 1977). In making this claim, the error theorist assumes that ordinary discourse is committed to a realist notion of objectivity – i.e. objectivity understood in terms of attitude-independence.

realists take objective moral truths to be attitude-independent, they can procure this intuition in a straightforward manner.

Many metaethicists agree that realist theories do a good job at accommodating morality's 'objective pretensions', and that in this respect they have a leg up on rivaling theories (e.g. Finlay 2007, p. 844). Some philosophers implicitly treat a realist understanding as the default understanding of moral objectivity (e.g. Mackie 1977; Enoch 2014), or even equate 'moral realism' and 'moral objectivism' (Kahane 2011), thus implying that all moral objectivists are moral realists.<sup>4</sup> This is an uncharitable move, however, since there has been an extensive, predominantly Kantian tradition of providing alternative accounts of moral objectivity, which do not centre on the notion of attitude-independence. Call these *antirealist* accounts of moral objectivity. What unites antirealist accounts of moral objectivity, as I shall understand them, is that they understand morality's objective features *in virtue* of our attitudes, or at least partly in virtue thereof.<sup>5</sup> For example, an antirealist might argue that moral judgements are objective if they result from a procedure that elicits dialogical understanding (Habermas 1995). Or she might maintain that moral judgements are objective if they are independent of the whims of anyone's *particular* attitudes: "Objectivity (...) is not so much a 'view from nowhere,' but a 'view of no one in particular'" (Sen 1993, p. 129).

My primary aim in this paper is to develop an antirealist view of moral objectivity in further detail, and to make its metaphysical commitments explicit. Extant antirealist accounts are underdeveloped in this respect. For example, to define objectivity as a "view of no one in particular" is insufficiently informative: the definition fails to provide a positive answer to the question *in virtue of what* moral judgements are objective. Building on the metaethical position that Sharon Street (2008b, 2009, 2010, 2012) has developed under the header "Humean Constructivism" (henceforth: HC), I seek to give antirealist theories of objectivity a more solid theoretical foundation. I expand on Street's constructivist account by framing moral objectivity in terms of standpoint-invariance: *moral judgements that purport to be objective, purport to have the property of being standpoint-invariant*. To distinguish between realist and antirealist accounts of moral objectivity, then, comes down to distinguishing between an account of objectivity in terms of attitude-independence, and an account of objectivity in terms of standpoint-invariance.

Street herself has expressed some doubts whether Humean constructivists can vindicate a sufficiently strong form of objectivity.<sup>6</sup> Indeed, there seems to be a consensus among many metaethicists that realists are better than antirealists at accommodating morality's 'objective

<sup>4</sup> Kahane 2011; see esp. p. 121, fn. 1 & p. 123, fn. 39.

<sup>5</sup> Different authors carve the lines between realist and antirealist positions in metaethics in different ways (Finlay 2007). One way of framing the distinction is as follows: realists vindicate some form of objectivism, whereas antirealists do not. Thus understood, the account of objectivity in terms of standpoint-invariance that I develop in this paper counts as realist. However, on the taxonomy that I follow in this paper, I use the label 'realism' to capture all views that adhere to a *specific* account of moral objectivity – objectivity in terms of attitude-independence – and the label 'antirealism' to capture all views that defend an account of moral objectivity *in virtue of our attitudes*. Hence, for the purposes of the present discussion, realists and antirealists stand united in defending an account of moral objectivity, but differ with respect to the *metaphysical framework* in terms of which they defend this account. While this is not the only way of drawing the realism-antirealism distinction, it is a common one (cf. Street 2016a, fn. 8).

<sup>6</sup> Street has touched upon the topic of objectivity at various places (e.g. 2010 p. 380; 2012; p. 41; 2016a, pp. 42–43), expressing slightly different views. Her general position, as I understand it, is that whereas so-called "Kantian" versions of constructivism vindicate a strong form of moral objectivity, the "Humean" version of constructivism that she defends *compromises* on moral objectivity (Street 2010, 2012; presentation at the workshop *Objectivity in Ethics* 2016).

pretensions'. The argument in support of this consensus goes roughly as follows. Realists advance a strong notion of objectivity, and this strong notion appears to capture the commitments of our ordinary moral discourse and thinking. If we assume that a metaethical account of moral objectivity scores plausibility points if it accords with ordinary discourse and thinking about moral objectivity, this finding counts in favour of realism (I discuss this assumption in **section 2**). Antirealist accounts, by contrast, lack the strength of their realist counterpart. *Prima facie*, antirealist objectivity merely seems to boil down to intersubjective agreement. The realist's strong notion of attitude-independence captures the prescriptive force and inescapability that seem characteristic of objective judgements in the moral domain; the antirealist seems to capture neither. Hence, a realist account of objectivity accords better with the intuitions of participants in moral discourse than the weaker account of the antirealist.

Contrary to this claim, I shall advance several reasons for thinking that the antirealist is at least as well positioned as the realist in accounting for our folk intuitions and in capturing our ordinary discourse.<sup>7</sup> First, building on recent experimental findings, I suggest that ordinary discourse and thinking are not committed to the realist's strong sense of objectivity, but rather to a graded form objectivism. Second, elaborating on Street's Humean Constructivism, I argue that the antirealist's account of objectivity is much stronger than appears at first glance, and that the antirealist may well be able to procure the intuition that moral judgements have an inescapable authority – at least, to the extent that experimental research indicates that it *should* procure this intuition. Hence, instead of compromising on the issue of moral objectivity, I argue that the antirealist is well-equipped to *vindicate* an account of objectivity with just the right strength, apt for the moral domain.

The paper proceeds as follows. In **section 2** I begin by clarifying how findings about ordinary moral discourse might be relevant for metaethical theorizing. In **section 3** I give a brief overview of recent experimental research on 'folk moral objectivism,' which allows us to refine our *explanandum* – i.e. morality's objective features. In **section 4** I outline the metaethical position that Street has developed under the header Humean Constructivism. In **section 5** I use Street's position as the basis to articulate an antirealist account of moral objectivity: objectivity as standpoint-invariance. I then go on to address three objections against HC that have been raised in the literature – the 'relativity objection' (**section 6**), the 'alignment objection' (**section 7**), and the 'revision objection' (**section 8**) – and outline the resources that the constructivist has to counter these objections. In the concluding **section 9** I contrast my antirealist account of moral objectivity with its realist counterpart, and make a preliminary case for the thesis that, even if a realist understanding of objectivity may be appropriate in other domains, in the moral domain an antirealist account of objectivity has a better fit.

## 2 The Metaethical Relevance of Ordinary Discourse

Many metaethicists regard it as a *desideratum* to account for morality's objective appearance. In this section and the next one, I outline various findings about ordinary moral thought and talk that have been cited in support of morality's apparent objectivism. But first we should get

<sup>7</sup> Naturally, there may be other respects in which realists and antirealists could be better or worse positioned; which position is the most plausible all things considered will have to be settled on the basis of comprehensive grounds. Traditionally, it has been held that whereas moral realists are better at capturing the moral appearances, antirealists are better at reconciling morality with our understanding of the mind and world (Finlay 2007, p. 844). In this article, I seek to cast doubt on the former claim; the latter claim is beyond the purview of my discussion.

clear about the metaethical relevance of these findings. Why think that the commitments of ordinary moral discourse carry any metaethical weight?<sup>8</sup>

It can be tricky to take people's ordinary thought and talk about moral objectivity at face value. After all, folk beliefs might be mistaken, and folk discourse might be deficient. Instead of mirroring people's ordinary beliefs about moral objectivity, our best metaethical theory might force us to revise these beliefs; instead of vindicating the commitments of ordinary discourse, our best metaethical theory might tell us to debunk them. Philosophers may hold views about the metaphysics of moral objectivity that are incompatible with laymen's intuitions, and it may well be that philosophers are correct in doing so, and that the laymen are mistaken.

How moral agents themselves conceive of their own discourse and what is the best philosophical understanding of that discourse, then, are two distinct issues. But even though these issues are logically distinct, they are not unrelated. Ordinary moral thought and talk are not beyond the purview of metaethical analysis. On the contrary: these are among the very phenomena that metaethicists seek to shed theoretical light upon. Whether or not ordinary discourse and laymen's beliefs about moral objectivism are ultimately vindicated, it is the metaethicist's task to accommodate them at the level of theory. As Blackburn (1984, p. 180) notes, metaethicists aspire to explain "why our discourse has the shape it does." Facts and findings about ordinary discourse constitute metaethical data; in turn, metaethicists seek to provide an adequate theoretical framework to account for them. In the final analysis, the plausibility of a metaethicist's framework will depend on how well this framework fits with our best theoretical understanding of the moral domain, all things considered – and findings about actual moral discourse can influence how many plausibility points a metaethical theory scores in this respect. Hence, while ordinary thought and discourse are not infallible, this does not preclude their metaethical relevance.

Indeed, metaethicists have long regarded it as a theoretical *desideratum* to accommodate the features of ordinary discourse. Consider some of the findings that have traditionally been cited in support of morality's apparent objectivism. These include the observation that moral judgements are often stated in the declarative mode, just like statements of fact (Bagnoli 2015); the observation that just like in theoretical deliberation, in practical deliberation we tend to assume that there are correct and mistaken answers about what we have reason to do (Blackburn 1984; Shafer-Landau 2003); and the observation that in terms of phenomenology, moral judgements seem to be more objective than judgements of taste, and almost or equally objective as factual judgements (Enoch 2014). These observations all take input from colloquial moral intuitions, moral practices, and moral language. They are the findings that give rise to the metaethical *explanandum* of accounting for morality's objective appearance.

Lastly, note that even though the commitments of ordinary discourse need not be taken at face value, it takes effort to discard them. All other things being equal, a metaethicist who claims that large parts of ordinary discourse are in need of revision faces a larger burden of proof than a metaethicist who vindicates moral discourse. If there is a good fit between some metaethical theory and some moral practice, then, unless and until arguments to the contrary have been given, this fit serves to corroborate the theory and to justify the practice. For example, if many people regard moral judgements as objective, and a metaethicist can vindicate their objectivity, then there is a good fit between practice and theory. By contrast, if many people regard moral judgements as objective, but a metaethical theory does not justify

<sup>8</sup> I thank the anonymous reviewers for pressing me to address this issue explicitly.

morality's purported objectivity, then this theory faces the more daunting task of accounting for the fact that many people's intuitions about moral objectivity are mistaken.

### 3 Refining the Explanandum

In the previous section I began to highlight some of the findings that have traditionally been cited in support of the claim that metaethicists have to account for morality's objective appearance. Only over the last decade or so, however, moral psychologists and experimental philosophers have begun to scrutinize the objectivist commitments of ordinary moral speakers in a rigorous, experimental fashion. This has resulted in a fairly extensive and ever-growing body of studies, with subjects that have been tested across different cultures and age-groups (e.g. Nichols 2004; Goodwin and Darley 2008; Sarkissian et al. 2011; Goodwin and Darley 2012; Wright et al. 2013; Beebe 2014; Khoo and Knobe, *Moral Disagreement and Moral Semantics*. Noûs, (forthcoming). While these studies by and large support the traditional consensus that moral discourse purports to be objective, they provide a more fine-grained understanding of the sense in which it does. In this section, I highlight some of the findings that the new wave of experiments has brought to light, and show how they refine the view that ordinary discourse is objectivist. Rather than being straightforwardly objectivist, I argue, the new experiments suggest that moral discourse is committed to a *graded* sense of objectivism.

Experimenters typically test the objectivist commitments of ordinary respondents by first presenting them with a range of personal, conventional, moral and scientific judgements. In some studies, respondents are asked to classify these judgements in one of these categories themselves. For example, a judgement predominantly classified as moral is that "Consciously discriminating against someone on the basis of race is unacceptable" (Wright et al. 2013). After being presented with a list of such statements, respondents are typically asked whether in cases of disagreement, they think that one of the disagreeing parties is mistaken. If the respondents think that it is possible for two disagreeing parties both to be correct about an issue, this is regarded as a denial of objectivity. If the respondents think that at least one of the disagreeing parties must be mistaken, this is regarded as an ascription of objectivity.

There is much more to say about the experimental setup of these studies, and their potential pitfalls. Unfortunately, a full discussion is beyond the scope of this paper. Let it suffice to say here that, while I do not think that these experiments are unproblematic in all respects, they are more rigorous than observations about the commitments of ordinary moral talk and practice solely derived from the armchair. For example, metaethicists typically cite people's intuition that in cases of moral disagreement, one of the disagreeing parties is mistaken, as an indicator of their objectivist commitments (e.g. Enoch 2014). These studies scrutinize the same intuition, but test it in experimental fashion. Given their rigour, these findings seem to constitute the best data about folk moral objectivism currently available for metaethicists.

While the experiments corroborate many prior ideas about the objectivist commitments of ordinary discourse, such as the idea that moral judgements purport to be more objective than judgements of taste, they also indicate that there are complexities about folk objectivism that metaethicists have heretofore overlooked. I highlight three of them. First, the *degree of objectivity* attributed to moral judgements varies substantially depending on the moral issue in question. For example, in Wright et al.'s (2013) study a majority of subjects self-classified the (un)acceptability of first trimester abortion as a *moral* issue, but it was classified as an *objective* issue only in 15% of the cases. The (un)acceptability of robbing a bank to pay for

holiday, on the other hand, which was also self-classified as a moral issue, was regarded as objective in 60% to 85% of the cases. This diversity in attributions of moral objectivity has also been found by other studies, suggesting that we might be mistaken to think that all moral judgements purport to be equally objective. Instead, experimental research suggests that in ordinary discourse, judgements on some moral issues are regarded as more objective than judgements on other moral issues (Goodwin and Darley 2012).

A second complication concerns the *scope* of appeals to objectivity. Philosophers have traditionally held that moral claims purport to be inescapable, in the sense that they are binding on all agents. Again, while empirical research supports this intuition at a general level, more fine-grained analyses suggest a more diverse picture. Asked about the scope of application of moral judgements, experimentally tested subjects typically respond as *graded* objectivists: while they have absolutist intuitions when moral judgements concern individuals from their own culture, they express increasingly relativist intuitions when these judgements concern individuals from increasingly different cultures or ways of life, such as aliens from outer space (Sarkissian et al. 2011). People tend to be cultural absolutists, but galactic relativists.

Building on the latter study, Khoo and Knobe, *Moral Disagreement and Moral Semantics*. Noûs, (forthcoming) have recently highlighted a third complication. They show that there are cases of moral conflict in which people are inclined to say both that the two speakers disagree and that it is not the case that at least one of them must be saying something incorrect. To reflect the moral semantics of ordinary discourse, they propose, we should distinguish between two types of disagreement: a type of disagreement in which the disagreeing parties are thought to occupy positions that exclude each other, and a type of disagreement where the positions of the disagreeing parties are non-exclusionary. In at least some cases of moral disagreement, the study suggests, ordinary respondents regard the positions of disagreeing parties as non-exclusionary.

To sum up, as the current state of research on folk objectivism indicates, moral discourse is committed to a *graded* sense of objectivity, in at least three ways. First, some moral judgements are held to be very objective, others less so. Second, even for those types of judgements that are held to be very objective, people do not think that their moral authority applies across the board. Third, some moral disagreements may be non-exclusionary, in the sense that none of the disagreeing parties are mistaken. The task for metaethicists who want to shed light on ordinary moral discourse, then, is to provide an *explanans* for this commitment to *graded objectivism*.

#### 4 Humean Constructivism

In the following sections I develop an antirealist account of objectivity that is particularly well suited to make sense of this commitment to *graded objectivism*. I do so by building on Sharon Street's version of Humean Constructivism, which I outline in this section. Following Street, I shall present metaethical constructivism as a theory about normative practical reasons. In the next section, I return to reasons that are specific to the moral domain.

Normative practical reasons are reasons to pursue, value, or favour something. For example, to make the normative judgement "*X* is a reason for *A* to *Y*" can be understood as "for agent *A*, *X* counts in favour of *Y*." Note that the concept of a 'normative reason' is distinct from the concept of a 'psychological reason': *X* can be a normative reason for *A* to *Y*, whether or not

A is psychologically inclined to  $Y$ .<sup>9</sup> But while ‘normative reasons’ and ‘psychological reasons’ should not be equated, the constructivist does emphasize that normative reasons are always reasons *for someone*. The judgement “ $X$  is a reason to  $Y$  full stop” is underdetermined: to know whether this judgement holds true we need to know *whose* reason it is. For Johnny, to hear a song by Prince is a reason to turn up the volume; for Jackie, it is not. For Louis, to hear that a dish contains bluefin tuna is a reason not to order it; for Laura, it is.

What is it for an agent to judge *correctly* that  $X$  is a normative reason to  $Y$ ? In other words: what are the truth-makers of normative judgements? On a straightforward subjectivist account, if  $A$  *thinks* that  $X$  is a reason for her to  $Y$ , her thinking makes it so. Metaethical constructivism is *not* straightforwardly subjectivist. According to the constructivist, whether  $X$  provides agent  $A$  with a reason to  $Y$ , depends on whether this judgement withstands scrutiny from agent  $A$ ’s full set of normative judgements, in combination with the non-normative facts. Call the full set of normative judgements that an agent makes the agent’s *evaluative standpoint*.<sup>10</sup> The constructivist maintains that the truth of any normative judgment is relative to the evaluative standpoint of the agent who makes the judgement, combined with the non-normative facts. Hence, it’s not Johnny’s *actual mental state*, but rather Johnny’s *evaluative standpoint* that sets the standards of correctness for his normative judgement.

What is it for an agent to *mistakenly* judge that  $X$  provides her with a reason to  $Y$ ? To commit a normative mistake is to make a judgement that *does not* withstand scrutiny from an agent’s evaluative standpoint, even if the agent, in making this judgement, *thinks* it does. There are several ways in which normative mistakes can come about. For example, an agent might be ignorant or misinformed about the relevant non-normative facts, she might not have given her judgement sufficient reflection and fail to see the full consequences of her existing commitments, or be incoherent in any other sort of way (cf. Street 2016b). Let’s consider these mistakes in some more detail, taking Laura – who thinks that the fact that a dish contains bluefin tuna ( $X$ ) is a reason to order it ( $Y$ ) – as our example.

#### A) *Factual misinformation*

First, let’s assume, plausibly, that it follows from Laura’s standpoint that she has reason to be informed about the relevant non-normative facts. That is to say, as a self-governed epistemic norm, Laura cares about being correctly informed by the facts that are relevant to her normative judgement: if her normative judgement at  $t_1$  rests on poor or false information, then at  $t_2$  she might retroactively realize that her judgement was mistaken by her own lights. For example, at  $t_1$  Laura doesn’t know that bluefin tuna is an endangered species ( $Z$ ). But once she becomes aware of this, she readily submits that she has reason not to order it. In other words, once at  $t_2$  Laura becomes aware that  $Z$ , she realizes that for her,  $X$  in combination with  $Z$  entails *not*  $Y$ . Since she judged herself to have reason to be informed about the relevant non-normative facts when she made her original judgement at  $t_1$ , she now realizes that, by her own lights, her judgement at  $t_1$  was mistaken.

<sup>9</sup> In what follows, when speaking of ‘reasons’ I refer to reasons in this non-psychological, normative sense. Following the current fashion in metaethics, the concept of a ‘normative reason’ may be regarded as a primitive concept of normative analysis (cf. Scanlon 1998, p. 19).

<sup>10</sup> Adopting Street’s terminology, we may alternatively call this an agent’s ‘practical point of view’, an agent’s ‘practical standpoint’, or her ‘full set of evaluative attitudes.’ Here and in what follows, I stick with the notion of an agent’s ‘evaluative standpoint.’



## B) *Incoherence*

Second, even if Laura is aware of the relevant non-normative facts, she might still err in her normative judgement. For example, although Laura is aware that bluefin tuna is an endangered species, and although she endorses the judgement that one shouldn't order dishes that contain endangered species, she fails to take this into explicit consideration when placing her order. Hence, while Laura judges that *X* and *Z* and thinks that she has normative reason to *Y*, she isn't alert to the fact that *Z* gives her instrumental reason to *not Y*. But once she thinks harder about her judgement, she realizes that *not Y*, instead of *Y*, is what actually follows from her evaluative standpoint, when stripped from inconsistencies. In making her initial judgement, then, Laura was mistaken.

Given the complexity involved in making normative judgements, and the myriad non-normative facts that might be relevant to a judgement, there are numerous ways in which an agent might err by her own lights (cf. Flanagan et al. 2014). Indeed, the Humean constructivist maintains that human agents make such mistakes all the time. Human beings are not ideally coherent; we rarely make normative judgements by weighing all the relevant options, reflecting on all the relevant norms, and gathering all the relevant non-normative facts. Instead we apply shortcuts and rules of thumb, we make decisions under conditions of less than full information, and as a result we regularly make mistakes by the lights of our own evaluative standpoint.

To see just how broad this range of potential normative mistakes is, consider mistakes that an agent could make due to failures of imagination. Suppose that it follows from an agent's standpoint that she has reason to be imaginative about the normative judgements that she could make. That is to say, the agent values being the sort of person who does not necessarily act upon one of the reasons that she is already familiar with. Instead, she values to use her imagination to probe whether she might not act upon reasons which she has not heretofore considered. Take the case of Laura. Perhaps Laura has not spent much time reflecting on moral issues; until  $t_2$  she has simply taken it for granted that *X* is a reason for her to *Y* (Laura's parents regularly ordered tuna dishes). But with hindsight, she blames herself for not considering the possibility that *X* might be taken as a reason to *not Y* – for, she is upfront to admit, she *could* have reasonably considered this, and since it followed from her evaluative standpoint that she had reason to be imaginative, she *should* have done so. Thus, even if up until  $t_2$  Laura has not even contemplated the possibility that *X* might be taken as a reason to *not Y*, it might still be the case that she is mistaken to judge at  $t_1$  that *X* is a reason for her to *Y*, since it already followed from her evaluative standpoint at  $t_1$  that she had reason to be imaginative about the normative judgements that she *could* make, including the judgement that *X* is a reason to *not Y*.

What follows from an agent's evaluative standpoint may be unveiled by answering the following question: what judgement would this agent make if she were ideally coherent, imaginative, and informed about the relevant non-normative facts? If the judgement she would make under these conditions coincides with the judgement that she actually makes, then her actual judgement is correct. If being more coherent or informed would lead her to make a different judgement, then her actual judgement is mistaken. Given the many ways in which such mistakes can come about, the Humean constructivist submits that, at least among the human agents in our world, ideal coherence is likely to be a rare trait.

## 5 Objectivity as Standpoint-Invariance

With our brief exposition of HC in place, let's scrutinize which facets of normative objectivity the constructivist can procure. Start with the weakest sense of objectivity: intersubjective agreement. On the constructivist's view, even though normative agreement is not guaranteed *in principle*, it is quite likely that, at least in our world, there will be an extensive overlap between the normative judgements that withstand scrutiny from the standpoints of different agents. After all, all human beings share the vast majority of their evolutionary history; as a result, we have many biologically grounded needs and preferences in common (cf. Street 2006). Cross-culturally endorsed norms typically have an adaptive rationale: their universality can often be explained as a consequence of the fact that they are the kinds of norms that tended to promote the survival and reproductive success of human ancestors. Moreover, human agents also share several social practices and institutions, as well as a reservoir of factual knowledge – and increasingly so, in our globalizing world. As a result, there is a substantial overlap between the normative judgements that different agents endorse.<sup>11</sup>

However, normative judgements in the *moral* domain typically make a stronger claim to objectivity. Apart from being intersubjectively valid, moral judgements typically purport to have an inescapable authority: they express norms that are “thought to give us reasons to act in conformity with them, even if we have no desires or interests that are served by those norms” (Handfield 2016, p. 62). Or, using Joyce's (2006) term, moral judgements seem to have ‘practical clout’. Can the constructivist procure the idea that moral judgements are strongly objective, in the sense of having this inescapable authority?

Prima facie, this seems to be a difficult task. According to HC, the standards of correctness of moral judgements are standpoint-relative: their correctness depends on whether they withstand scrutiny from the standpoints of individual agents. Therefore, it may appear that HC cannot procure the intuition that moral judgements can have normative force irrespective of what an agent wants. Agent *A* only has reason to *X* if *X* follows from her own standpoint: the ultimate source of a normative judgement's prescriptive authority, the constructivist maintains, comes from within, not from some external source.

However, recall that what follows from an agent's evaluative standpoint need not coincide with what the agent *thinks* that follows from her standpoint. The constructivist maintains that there are objective standards that govern whether a normative judgement is correct or mistaken, irrespective of whether an agent thinks of her judgement as being correct or mistaken. Not the agent's actual mental states, but the agent's *evaluative standpoint* legislates whether or not her normative judgement is correct. As a result, an agent's actual normative

<sup>11</sup> To appreciate this, it is helpful to think of the vast number of possible normative judgements any agent *could* make (cf. Street 2006, 2008a, 2016a). A priori there are innumerable judgements that might follow from an evaluative standpoint – “from the judgement that infanticide is laudable, to the judgement that plants are more valuable than human beings, to the judgement that the fact that something is purple is a reason to scream at it” (Street 2006, p. 133). Relative to this vast set of possible judgements, the variation in the normative judgements that human beings *actually* make is quite limited. For example, there may well be universal agreement on the judgement that “seeing a fire is a reason not place your hand in it” and that “smelling that some substance is covered with faeces is a reason not to eat it” – normative judgements that we may find so obvious that ordinarily we do not bother to spell them out.

Naturally, as one reviewer pointed out, just as there is pervasive agreement on some normative issues, there is also pervasive disagreement on some normative issues. The constructivist does not deny this; indeed, she acknowledges that some disagreements may rest on fundamental differences in evaluative standpoint (see this section, **section 6** and the end of **section 9**).

judgements may legitimately be criticized by other agents: agent *B* may be in a position to argue that it follows from agent's *A* standpoint that she has reason to *X*, even if agent *A* has not yet realized this.

Hence, an agent's evaluative standpoint need not be fully transparent. Moreover, an agent's standpoint is neither immutable: other agents can try to change it, for example by updating her non-normative beliefs, or by inviting her to take a different perspective and fuelling her imagination. Note that the incentive to change other people's attitudes is typical for *moral* judgements, but much less so for other evaluative judgements, such as judgements of taste, convention, or aesthetics. While there may be intersubjective agreement on the truth that ice-cream is tasty, agents will typically show little motivation to *push* for social convergence on this issue. For *moral* judgements, on the other hand, there tends to be a much greater incentive to influence people's attitudes. As Gibbard (1990, p. 171) notes, in making a claim to moral objectivity "I do more than simply evince the norms that I accept; what I do is to exert a conversational pressure. In effect I demand acceptance of what I am saying." Moral judgements concern issues where we *want* other agents to conform to our attitudes.

HC can accommodate this datum: it can procure the idea that at least in terms of appearance, objective moral judgements typically have practical clout. If an agent defends a moral truth that she regards as objective, she typically thinks that this truth *should* withstand scrutiny from the evaluative standpoints of other moral agents, irrespective of whether they *think* it does. Hence, the agent is committed to achieving social convergence on this point (cf. Mitchell-Yellin 2015, p. 5). Consider Louis, who is a preservationist. Louis judges that preserving the bluefin tuna as a biological species is valuable. Moreover, he thinks that others have reason to value this as well: not only Laura, but *no* human agent has reason to order a dish that contains bluefin tuna, Louis maintains. In other words, he regards the judgement that human agents have reason not to order tuna as inescapable: it is binding on all human agents. Other agents, such as Laura, may of course disagree: she does not think that this judgement withstands scrutiny from her own evaluative standpoint. But for Louis, whether Laura *thinks* that this judgement withstands scrutiny from her own standpoint is quite beside the point: what concerns him is whether it *does*. And since he maintains that the judgement is objective – he thinks that it *should* follow from Laura's standpoint that she has reason not to order tuna – the challenge for Louis is to update Laura's attitudes and non-normative beliefs in such a way, that she comes to realize that, by her own lights, she has reason not to order tuna.

Crucially, on this account of objectivity, the truth or falsity of the moral judgement that *X* is not *independent* of the standpoint of moral agents. Instead, the idea is that however these standpoints may differ, the truth of *X* will still follow from it. In other words, the truth of *X* *invariably* follows from different evaluative standpoints. This invariance is the distinctive feature of an antirealist account of objectivity, and sets it apart from its realist counterpart. On the antirealist account, if a moral judgement purports to be objective, it purports to withstand scrutiny from a diverse set of evaluative standpoints. A fully objective moral judgement purports to be fully standpoint-invariant: no matter how the standpoint of a moral agent varies, the truth of *X* should follow from it. But the constructivist can accommodate the finding, outlined in **section 3**, that not all moral judgements aspire to make such a strong claim to objectivity. Instead, some moral judgements are held to be more objective than others, and their inescapability is a matter of degree. The antirealist accounts for this datum by maintaining that the *degree* of objectivity of a moral judgement is a function of the diversity of evaluative standpoints under which its truth or falsity holds fixed. The more diverse these standpoints are, the more objective the judgement. In this way, the antirealist makes room for a *graded*

understanding of moral objectivity, which, we have seen, is precisely the sense of objectivity that a metaethical account *should* procure, as suggested by experimental findings.

Lastly, note that on an antirealist account, although some moral judgements may aspire to be objective in a very strong sense, this aspiration is not unbounded. Conceivably, there will be agents for whom the truth of *X* *does not* follow from their standpoint, even though many other agents regard *X* as a fully objective moral truth. The constructivist concedes that this is a genuine theoretical possibility – and in the most extreme of such cases, she will maintain, it is most sensible to conclude that fully coherent agents who endorse very bizarre normative judgements, should not be regarded as *moral* agents. This last point brings us to thought-experiments that involve “Ideally Coherent Eccentrics” (ICEs), which are often invoked to push the ‘relativity objection’ against antirealist objectivity.

## 6 The Relativity Objection

One objection that has been raised against HC is that the view is too relativist. Borrowing Gibbard’s (1999) example, consider Caligula, who judges that he has reason to torture other people for fun. Let’s stipulate that Caligula is ideally coherent, imaginative and informed about the non-normative facts, and that his judgement to torture people for fun withstands reflective scrutiny from his evaluative standpoint.<sup>12</sup> The Humean constructivist is committed to the claim that Caligula is *correct* in making this judgement: after all, to make a normative judgement that withstands scrutiny from one’s own evaluative standpoint is to make a correct normative judgement. But this seems to run counter to common-sense. As Bratman (2012, p. 84) notes it is “plausible that our ordinary thinking about reasons for action supports Gibbard’s denial that Caligula has a normative reason to torture even if Caligula’s standpoint wholeheartedly supports torture and does not involve relevant and false non-normative belief.” To formulate the worry in a different way: our ordinary thinking seems to support the idea that some normative judgements – typically *moral* judgements – are objectively binding, in the sense that that *all* agents should be responsive to them. HC, on the other hand, makes the truth of moral judgements escapable. In doing so, the constructivist fails to capture morality’s binding authority.

Note two things about the relativity objection. First, the objection is fostered by observations about ordinary moral thinking: since our colloquial intuitions support the idea that eccentric creatures like Caligula are making a moral mistake, a plausible metaethical theory should procure this idea too. Second, note that the supposed objectivist commitment of our ordinary thinking is quite strong. Naturally, the constructivist could maintain that other agents have reason to prevent Caligula from torturing other agents; after all, it might well follow from *their* evaluative standpoint that Caligula is making a moral mistake. But according to the relativity objection, this would not be enough. Allegedly, our ordinary thinking is not only committed to the idea that it would be *right* if people tried to stop Caligula, but also that Caligula is *mistaken* in his judgement. In other words: allegedly, our ordinary thinking about morality supports the realist’s strong notion of objectivity, rather than the antirealist’s weaker notion.

Which resources does the constructivist have to dispel this objection? First, she might question whether our colloquial intuitions have any relevance for the plausibility of thought-

<sup>12</sup> Alternatively, we might say that Caligula has *most* reason to torture other people for fun, or that he has reason to torture others for fun *all things considered*. From a constructivist viewpoint, however, these additions are not strictly necessary; they are implied by the stipulation that this judgement withstands scrutiny from Caligula’s evaluative standpoint.

experiments that involve ideally coherent creatures with eccentric moral preferences. This is not because ordinary discourse is irrelevant for metaethics – in **section 2** I have provided several considerations for thinking otherwise – but because the thought-experiments involve creatures that are radically different from the agents that people are familiar with. As a result, it is doubtful that our intuitions about such thought-experiments have epistemic value. Arguably, we should suspend our intuitions about what Caligula has most reason to do.

Second, the constructivist may grant that our ordinary thinking about ICEs *is* metaethically relevant, but question whether our ordinary thinking does in fact support the realist's strong notion of objectivity. If the constructivist takes this second route, the experimental research on folk objectivism should be taken into account. After all, these experiments provide the most rigorous findings about what people's ordinary thinking about morality is committed to. Specific scenarios involving ICEs have not yet been experimentally tested, but the extant research cited in **section 3** provides some preliminary indication that the folk notion of moral objectivity is weaker than that of the realist. For example, as Khoo and Knobe, *Moral Disagreement and Moral Semantics*. Noûs, (forthcoming) show, ordinary discourse supports the idea that there can be cases of moral disagreement, in which none of the disagreeing parties is mistaken. The Caligula scenario may well be such a case: perhaps we genuinely disagree with Caligula about what he has moral reason to do, but do not think that Caligula is making a mistake in upholding his eccentric judgement.

Moreover, we should be aware that comparisons with the historical Caligula are likely to be distorting where the thought-experiment is concerned. After all, it is not at all obvious that the historical Caligula was ideally coherent. In fact, if the constructivist position is along the right lines, there is reason to think that the historical Caligula was probably not ideally coherent. While the historical Caligula did have a reputation of killing in a whim, there is no indication that he was highly imaginative, factually informed and very reflective. Given the many coherence failures that human agents are liable to make, we may, therefore, hypothesize that when the historical Caligula judged that he had reason to torture others for fun, he was *mistaken* by his own lights.

Ideal coherence is a very strenuous condition. In **section 4** we discussed some of the many failures that less-than-ideal human agents can make by the lights of their own evaluative standpoint. By stipulation, ICEs commit none of these mistakes. Hence, by stipulation our ideally coherent Caligula is fully aware of the fact that torture causes pain; he is fully aware that his judgement will lead others to condemn and penalize him; he is imaginative and able to think of many other things that he could do for fun. In spite of all of this, he still judges that he has most reason to torture others for fun. We should be alert to the fact that this makes Caligula *a very exotic kind of creature*, who inhabits some possible world which may be quite distant from the world that humans inhabit. He may not share our evolutionary history, he may not be familiar with our sociocultural practices, and he may not share our factual knowledge – for if he did, it seems implausible to suppose that his queer judgement would withstand scrutiny from his own standpoint. As Street (2009, p. 281) notes, “an accurately imagined ICE will look more like an interesting visitor from another planet than a human being.” If this is correct, then the experiments by Sarkissian et al. (2011) provide further reason to think that our ordinary thinking *does not* support the intuition that Caligula is objectively mistaken. After all, their studies suggest that where aliens are concerned, ordinary respondents do not think that objective moral judgements retain their binding authority. Or, in terms of the constructivist position outlined in **section 5**: on our ordinary thinking about morality, creatures like Caligula may not be regarded as *moral* agents.

To sum up, there are grounds for doubting that our colloquial notion of moral objectivity accords particularly well with the strong notion of the realist. Therefore, unless and until

further arguments are given, the objection that the antirealist's notion of objectivity is too relativist is under-supported. The antirealist might even make a preliminary case that the graded relativism typical of ordinary moral discourse is *in tension* with the realist's strong objectivism. However, these are still the early days of experimental research on folk moral objectivism, and we should be careful not to overstate the implications of this research. What I hope to have achieved, at this stage of the dialectic, is at least to have cast reasonable doubt on the claim, traditionally upheld by many metaethicists, that realists have a leg up on antirealist theories in accounting for the objectivist features of ordinary discourse.

## 7 The Alignment Objection

Another objection that has been pressed against HC concerns its purported inability to make sense of the idea that agents can criticize the normative judgements of other agents (Berker 2014). Consider once more the case of Laura and Louis. Laura judges that she has reason to order tuna, and the Humean constructivist contends that whether this judgement holds true depends on whether it withstands scrutiny from Laura's evaluative standpoint. But now consider Louis, who judges that Laura has reason *not* to order tuna. We should, of course, be able to make sense of the idea that Louis can be correct about this judgement. But it may seem that HC cannot make sense of this idea: whether or not Laura has reason to order tuna depends on whether this judgement withstands scrutiny from Laura's – not Louis's – evaluative standpoint. Presumably, their evaluative standpoints do not fully align. Therefore, it seems that Louis is in no position to judge what Laura has reason to do.

Call this the alignment objection.<sup>13</sup> According to the alignment objection, the fact that we have imperfect knowledge about the evaluative standpoints of other agents, precludes us from making correct judgements about the reasons that those agents have. However, in **sections 4** and **5** we have already seen that the constructivist has resources to dispel this objection. Recall that the constructivist readily admits that agents can make mistakes *by their own lights*. After all, there is a potential mismatch between the reasons that an agent *takes* herself to have, and the reasons that she *actually has*, as follows from her evaluative standpoint. As a result, the reasons that an agent *takes* herself can legitimately be criticized, if they do not withstand scrutiny from the agent's own evaluative standpoint – and other agents may be well positioned to press such criticism.

What the alignment objection correctly brings out is that such criticism is fallible. Whether it *does* follow from Laura's evaluative standpoint that she has reason not to order tuna, may not be fully transparent to Louis. The Humean constructivist does not deny this. There is likely to be some subjective uncertainty about what does in fact follow from the standpoint of other agents – and unless an agent is ideally coherent, there will also be some subjective uncertainty about what follows from her *own* evaluative standpoint. Indeed, it is this very uncertainty that allows the constructivist to procure the intuition that we can be *mistaken* about our normative judgements, both with concern to our own reasons, as well as the reasons of others.

But the consequences of this uncertainty are not as dire as the alignment objection suggests. Even if evaluative standpoints are not fully transparent, they are neither fully opaque. Agent A

<sup>13</sup> This objection is distinct from what Bratman (2012, p. 81, 85, 95) calls “the problem of alignment.” By this, Bratman means the challenge for HC to provide an account of the agent's evaluative standpoint that aligns with insights from the philosophy of action.

is familiar with her own attitudes, and she will generally have a clear – if fallible – intuition about what *does* follow from her own evaluative standpoint. For other agents, agent A's evaluative standpoint may be slightly more opaque, but again, not fully so. In particular, agents who are acquainted with agent A are likely to be familiar with many of the values that constitute her standpoint, and even agents who are not acquainted with agent A, but who do share their evolutionary history with her (as well as, perhaps, much of their cultural upbringing and factual knowledge), may be sufficiently well positioned to assess what follows from her evaluative standpoint. Hence, even if different standpoints do not fully align, they may still be sufficiently similar to warrant intersubjective criticism.

## 8 The Revision Objection

A third objection that has been raised against against HC is that the view cannot account for types of hindsight that involve radical changes of value (Jaeger 2015). Consider a former member of the Ku Klux Klan, whose attitudes have radically transformed over time: at t1 he judged that he had reason to discriminate black people, but at t2 he despises his former racist attitudes. An adequate theoretical account should be able to preserve the idea that the Klan member, looking back at his former self, can judge that he was mistaken about his reasons at t1. But HC cannot preserve this idea, or so it seems. The values that the former Klan member endorses at t2 differ radically from the values he endorsed at t1 – but the evaluative standpoint of the Klan member at t1, rather than t2, sets the standards of what he has reason to do at t1. If, by the lights of his own evaluative standpoint at t1, the Klan member had no reason to discard his racist attitudes, then his later self at t2 cannot criticize him for not having done so. But this is unintuitive, or so it seems: supposedly, we would want to procure the idea that one's past self can be mistaken about his reasons, even if the past self could not have recognized these reasons.

The objector is right to point out that according to HC, the standards of correctness for an evaluative judgement are relative to the further set of values that the agent endorses at the time of making this judgement. But the constructivist has resources to argue that the implications of this view are less implausible than the objector suggests. Instead, she might counter, the objector presents a scenario that is implausible in its own right. To see this, note that on a constructivist reading, if the Klan member at t1 had no reason *at all* to detest his own racist attitudes, this implies that even if the Klan member had been ideally coherent, imaginative and informed, he would still have stuck with his racist values. In other words, the example stipulates that there really is *no ground whatsoever* for criticizing the values of the Klan member by the lights of his own values at t1. But if this is the case, it seems problematic to suppose that at t2, the former Klan member thinks that he has reason to criticize his former self. After all, if the agent's later self is aware that his former self could not have come up with *any* reason not to *X*, then it seems peculiar to suppose that his later self would nonetheless criticize his former self for not having come up with a reason not to *X*. The constructivist may argue that we had best conclude that the agent's future self is confused, or that the stipulated scenario is simply incoherent.

Alternatively, the scenario may be such that the Klan member at t2 can *legitimately* criticize his former self: there *really was* a reason for the Klan member to detest his own racism, even though he wasn't aware of this reason. Recall, once more, that HC distinguishes between the reasons that an agent *takes* herself to have, and the reasons she *actually* has, as follows from her evaluative standpoint. Because of the potential mismatch between them, it may well be the

case that it *did* in fact follow from the Klan member's standpoint at t1 that he had reason to detest his own racist attitudes, but that this reason did not become apparent to the agent up until t2. Thus understood, HC has no problem to accommodate the scenario: the constructivist can procure the idea that the past self was mistaken about his reasons, provided that this past self plausibly *could have* recognized his mistake.

## 9 Conclusion

In this paper I have highlighted two accounts accounts of moral objectivity: a realist and an antirealist account. On a realist account, morality's objective features are understood in terms of their *attitude-independence*. Just like the judgement that the Earth revolves around the sun, but unlike the judgement that ice-cream is tasty, the moral realist maintains that the truth or falsity of a moral judgement is independent of what people think or feel about it.

By contrast, on an antirealist account, morality's objective features are understood in terms of their *standpoint-invariance*. Invariance differs from independence. According to the antirealist, the objectivity of moral judgements does *depend* on people's attitudes, but is constituted by the fact that these judgements withstand scrutiny from a diverse set of evaluative standpoints. As a result, moral judgements are not objective in the same sense as facts about planetary orbits: their truth or falsity is not similarly attitude-independent. But moral judgements also differ from judgements of taste, convention and aesthetics, in the sense that moral judgements typically purport to have an inescapable authority, which these other classes of evaluative judgement do not. When an agent judges a moral truth to be objective, the constructivist maintains, she is committed to achieving social convergence about this truth. It is this commitment that gives moral judgements their prescriptive force, which seems to be absent from evaluative judgements in non-moral domains. Hence, if the constructivist account is correct, we are mistaken to think that moral judgements *either* have to be factual judgements, *or* judgements of taste. Instead, at least where their objectivist aspiration is concerned, moral judgements belong to a different category altogether. This idea fits well with research about folk objectivism: while moral judgements are generally regarded as much more objective than judgements of taste, most people regard them as slightly less objective than factual judgements (Goodwin and Darley 2008; Wright et al. 2013).

It might, perhaps, be objected that the antirealist fails to vindicate the 'real' sense of objectivity. By definition, the objector might maintain, objectivity means attitude-independence; anyone who "settles for less" has changed the subject. We should keep in mind, however, that objectivity is a rich notion, which is appealed to in different senses and across a wide variety of discourses (Nozick 2001; Burge 2010). It is not obvious that one account should have to be vindicated at the cost of all others. Instead, it seems to be more fruitful to assess which account of objectivity is most suitable for which specific domain. While there may certainly be domains in which an understanding of objectivity in terms of attitude-independence is most apt, I have presented some considerations for thinking that in the moral domain, an understanding of objectivity in terms of standpoint-invariance is more suitable.

It has been beyond the scope of this paper to provide a full comparison of both realist and antirealist accounts. It may, of course, turn out that the realist can accommodate the experimental findings I have presented in this paper equally well as the antirealist can. However, if my argument has been successful, I have at least been able to show that, contrary to the current fashion in metaethics, the aptness of a realist understanding of moral objectivity need not be



taken for granted. Instead, I have articulated an alternative, antirealist account, and have presented some grounds for thinking that this account is better suited to capture the commitments of ordinary moral thought and practice.

On the account that I have presented, the objectivity of moral judgements may be quite strong, but it is not unbounded. ICEs who are correct in their eccentric moral judgements constitute a genuine possibility, at least in theory. Moral conflicts can have a tragic dimension, the antirealist submits, in the sense that there is no ultimate ground to decide which of the conflicting parties is right. In practice such cases of deep conflict may be rare – ideal coherence, we have noted, is an achievement that might hardly ever be achieved by human agents in our actual world – but they constitute a theoretical possibility nonetheless. Those with realist inclinations may consider this a weakness of the account. Those who think that a vindication of moral objectivity need not preclude some degree of moral relativism, on the other hand, might consider it one of its deeper appeals. The tragic dimension of moral discourse is a dimension that the antirealist admits to exist – and I submit that in a different sense of ‘realism’, she is all the more realistic in doing so.

**Acknowledgements** I thank two anonymous referees for their helpful comments. Earlier versions of this article were presented at the conferences ‘Objectivity in Ethics’ (Utrecht University, 2016) and ‘Realism in Ethics’ (University of Sheffield, 2016).

**Open Access** This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

## References

- Bagnoli C (2015) Moral objectivity: a Kantian illusion? *J Value Inq* 49:31–45
- Beebe JR (2014) How different kinds of disagreement impact folk metaethical judgments. In: Wright JC, Sarkissian H (eds) *Advances in experimental moral psychology: affect, character, and commitments*. Continuum, London, pp 167–187
- Berker S (2014) Does evolutionary psychology show that normativity is mind-dependent? In: D’Arms J, Jacobson D (eds) *Moral psychology and human agency: essays on the science of ethics*. Oxford UP, New York, pp 215–252
- Blackburn S (1984) *Spreading the word: groundings in the philosophy of language*. Oxford UP, New York
- Bratman M (2012) Constructivism, agency, and the problem of alignment. In: Lenman J, Shemmer Y (eds) *Constructivism in practical philosophy*. Oxford University, Oxford, pp 81–98
- Burge T (2010) *Origins of objectivity*. Clarendon Press, Oxford
- Clarke-Doane J (2014) Moral epistemology: the mathematics analogy. *Noûs* 48:238–255
- Enoch, D. 2014. Why I am an Objectivist about Ethics (And Why You Are, Too). In: R. Shafer-Landau ed., *The Ethical Life* (3rd ed.). Oxford: Oxford UP
- Finlay S (2007) Four faces of moral realism. *Philos Compass* 6:820–849
- Flanagan O, Ancell A, Martin S, Steenbergen G (2014) Empiricism and normative ethics: what do the biology and the psychology of morality have to do with ethics? *Behaviour* 151:209–228
- Gibbard A (1990) *Apt choices, wise feelings: a theory of normative judgment*. Clarendon Press, Oxford
- Gibbard A (1999) Morality as consistency in living: Korsgaard’s Kantian lectures. *Ethics* 110:140–164
- Goodwin GP, Darley JM (2008) The psychology of meta-ethics: exploring objectivism. *Cognition* 106: 1339–1366
- Goodwin GP, Darley JM (2012) Why are some moral beliefs perceived to be more objective than others? *J Exp Soc Psychol* 48:250–256
- Habermas J (1995) Reconciliation through the public use of reason: remarks on John Rawls’s political liberalism. *J Philos* 92:109–131

- Handfield T (2016) Genealogical explanations of chance and morals. In: Leibowitz UD, Sinclair N (eds) *Explanation in ethics and mathematics: debunking and dispensability*. Oxford UP, Oxford, pp 58–82
- Jaeger G (2015) The wisdom of hindsight and the limits of Humean constructivism. *Philos Explor* 18:45–57
- Joyce R (2006) *The evolution of morality*. MIT Press, Cambridge, MA
- Kahane G (2011) Evolutionary debunking arguments. *Noûs* 45:103–125
- Leibowitz UD, Sinclair N (eds) (2016) *Explanation in ethics and mathematics: debunking and dispensability*. Oxford UP, Oxford
- Mackie JL (1977) *Ethics: inventing right and wrong*. Penguin, Harmondsworth
- Mitchell-Yellin B (2015) Aligning with the good. *J Ethics Soc Philos*:1–8
- Nichols S (2004) After objectivity: an empirical study of moral judgment. *Philos Psychol* 17:5–28
- Nichols S (2014) Process debunking and ethics. *Ethics* 124:727–749
- Nozick R (2001) *Invariances: the structure of the objective world*. Belknap Press of Harvard UP, Harvard
- Sarkissian H, Parks J, Tien D, Wright JC, Knobe J (2011) Folk moral relativism. *Mind Lang* 26:428–505
- Scanlon TM (1998) *What We owe to each other*. Harvard UP, Cambridge
- Sen AK (1993) Positional objectivity. *Philos Public Aff* 22:126–145
- Shafer-Landau R (2003) *Moral realism: a Defence*. Oxford UP, New York
- Street S (2006) A Darwinian dilemma for realist theories of value. *Philos Stud* 127:109–166
- Street S (2008a) Reply to Copp: naturalism, normativity, and the varieties of realism worth worrying about. *Philos Issues* 18:207–228
- Street S (2008b) Constructivism about reasons. *Oxford Stud Metaethics* 3:207–245
- Street S (2009) In defense of future Tuesday indifference. Ideally coherent eccentrics and the contingency of what matters. *Philos Issues* 19:273–298
- Street S (2010) What is constructivism in ethics and metaethics? *Philos Compass* 5:363–384
- Street S (2012) Coming to terms with contingency: Humean constructivism about practical reason. In: Lenman J, Shemmer Y (eds) *Constructivism in practical philosophy*. Oxford University Press, Oxford, pp 40–59
- Street S (2016a) Objectivity and truth: You'd better rethink it. In: Shafer-Landau R (ed) *Oxford studies in metaethics*, vol 11. Clarendon Press. Cited page numbers refer to the penultimate draft. Available at <https://nyu.academia.edu/SharonStreet>
- Street S (2016b) Nothing “really” matters, but That’s not what matters. In: Singer P (ed) *Does anything really matter: Parfit on objectivity*. Oxford UP, Oxford
- Timmons M (2010) Objectivity in Moral Discourse. In: Barber A, Stainton R (eds.) *Concise Encyclopedia of Philosophy of Language and Linguistics*, 539–545.
- Wright JC, Grandjean PT, McWhite CB (2013) The meta-ethical grounding of our moral beliefs: evidence for metaethical pluralism. *Philos Psychol* 26:336–361