

Quasi-realism

Terence Cuneo

Over forty years ago, Simon Blackburn introduced the philosophical world to a figure he called the “quasi-realist.” According to Blackburn, the quasi-realist is someone who “starting from an antirealist position finds himself progressively able to mimic the thoughts and practices supposedly definitive of realism” (Blackburn 1993: 4; see Blackburn 1994, 1998). In the intervening years, other philosophers – most notably, Allan Gibbard – have joined Blackburn in developing the quasi-realist research program in metaethics, defending it from objections and arguing that it exhibits considerable promise (Gibbard 2003, 2012). These efforts at developing and defending quasi-realism have garnered considerable attention from sympathizers and critics alike. Still, the view has remained elusive: it proves extraordinarily difficult to formulate an accurate and informative statement of what quasi-realism is.

It is no accident that quasi-realism defies ready characterization. A close look at what its proponents say when characterizing the view reveals two ambiguities that explain not just why the position is so difficult to characterize but also to assess. I dedicate the first part of this chapter to identifying these ambiguities and their import. Having located these ambiguities, I close by articulating a challenge that the view must address, which is that it offers no satisfactory account of how we can make moral mistakes.

PRELIMINARIES

Let us begin with some preliminary matters. In the passage quoted above from Blackburn, we’re told that the quasi-realist starts from an antirealist position and then progressively mimics the thoughts and practices (supposedly) definitive of realism. When Blackburn says that the quasi-realist starts metaethical theorizing from an antirealist position, he has a particular type of moral antirealist position in mind. The position is a close relative of A. J. Ayer’s expressivist view according to which there are no moral facts or truths, and

all moral thought and discourse are, or express, states of approbation or disapprobation toward non-moral reality (Blackburn 1993: 187). Similarly, when Blackburn says that the quasi-realist progressively mimics the “thoughts and practices supposedly definitive of realism,” he also has a particular version of moral realism in mind. The version of realism in question is something close to G. E. Moore’s nonnaturalism according to which moral thought and discourse represent moral reality, which is itself not part of the natural order investigated by the sciences (Blackburn 1998: chapter 3; Gibbard 2003). In short, the quasi-realist begins metaethical theorizing from an expressivist starting point but ends up affirming claims that sound a lot like nonnaturalist moral realism.

To make sense of this strategy, we need to understand what distinguishes expressivist positions such as Ayer’s from realist ones such as Moore’s. I propose to do so as follows. Suppose we were to reflect on everyday moral thought and practice in our capacity as ethical theorists. Such reflection reveals that:

- (i) Moral judgments have the marks of a descriptive belief: they are classificatory, truth-evaluable, apt candidates for knowledge, and apt for inference.
- (ii) Not any response to a moral question will do; we can make moral mistakes. Moreover, some answers to moral questions are better than others.
- (iii) Moral judgments have the marks of a practical attitude: they are often directive and motivationally efficacious.

Let us call these the *core metaethical data* (or the “core data,” for short). Much of what transpires in metaethical theorizing consists in different theories attempting to accommodate and explain these (and other) data – or to explain why they do not require accommodation.

Realism typically attempts to accommodate and explain the full range of data stated above. It does so by appeal to the following two theses:

Moral Representation: moral thought and discourse represent moral facts. That is, they have moral representational content.

And:

Moral Facts: there are moral facts.

According to realism, its first thesis fully accommodates and explains data such as (i) and partially accommodates and explains data such as (iii). It is because moral thought and discourse represent moral reality that moral judgments have the marks of descriptive beliefs and can guide our behavior. Realism adds that it’s the combination of its two theses that accommodates and explains data such as (ii). In at least some cases, we make moral mistakes when we inaccurately represent the moral facts. The point to underscore for our purposes is that realism does not simply embrace Moral Representation and Moral Facts. It also commits itself to:

Explanation: when attempting to accommodate and explain the core moral data, appeal to Moral Representation and Moral Facts.

For ease of reference, call these the *realist theses*.

Expressivists have devoted most of their energies attempting to accommodate and explain (i). Their attempted accommodation and explanation is distinctive in two

respects. First, unlike realism, it does not appeal to Moral Representation. Instead, it appeals to:

Attitude: moral thought and discourse are, or express, attitudinal states. These are states of commendation and condemnation that lack moral representational content.

Second, in its most sophisticated quasi-realist guises, expressivism has often *embraced* the first two realist theses stated above. Nonetheless, it has steadfastly maintained that it is not a version of realism, since it rejects Explanation in favor of:

No Explanation: when attempting to accommodate and explain the core moral data, do not appeal to Moral Representation or Moral Facts. Rather, appeal to claims such as Attitude.

Blackburn is explicit about the negative component of this commitment when he writes that expressivism is “visibly antirealist, for the explanations offered make no irreducible or essential appeal to the existence of moral ‘properties’ or ‘facts.’” (Blackburn 1993: 175; cf. Gibbard 2003: 183; Dreier 2004).

I have said that expressivism has devoted most of its energies to accommodating and explaining (i). No Explanation, however, implies that expressivism endeavors to accommodate and explain not just (i) but the full range of core moral data. The textual evidence appears to favor attributing this claim to expressivism. For example, when Gibbard writes that “[a]lmost all of what descriptivists insist on can be embraced and explained by an expressivist” and that what “quasi-realism mimics is not tempered realism as a whole but tempered realism in all but one respect,” he indicates that expressivism has wide ambitions (Gibbard 2003: 20; 2011: 45). The position attempts to “mimic” what realism attempts to accommodate and explain while also embracing No Explanation. Moreover, a look at how expressivists have responded to the charge that it fails to accommodate and explain items that are plausibly viewed as belonging to the core data reveals a distinctive pattern. For example, when Nicholas Sturgeon charges that expressivism cannot capture the explanatory role of moral reality, Blackburn replies that it can (Sturgeon 1991; Blackburn 1993: chapter 11). When Andy Egan contends that expressivism cannot explain certain kinds of moral mistakes, Blackburn maintains that the charge is untrue (Egan 2007; Blackburn 2009). And when Sharon Street charges that expressivism cannot vindicate the trustworthiness of moral judgments, Gibbard replies that this is not so (Street 2011; Gibbard 2011). In this respect, expressivism is unlike positions such as subjectivism and error theory, which do not attempt to accommodate and explain the full range of core moral data, and like realism, which does attempt to accommodate and explain the full range of core data. At any rate, in what follows I’ll refer to the two theses stated above as the *expressivist theses*.

WHAT IS QUASI-REALISM?

Having addressed these preliminary matters, we’re now ready to address the question of how best to understand the quasi-realist project. To help us address this question, let’s have before us a series of passages in which quasi-realists articulate their view. Earlier we heard Blackburn note that he:

invented the figure of the quasi-realist, or someone who “starting from an anti-realist position finds himself progressively able to mimic the thoughts and practices supposedly definitive of realism.”

After describing quasi-realism as an approach that “mimics Moore’s and Ewing’s non-naturalistic moral realism,” Gibbard qualifies what he says, writing that quasi-realism endeavors to thoroughly mimic the “tempered” nonnaturalism of Derek Parfit, T. M. Scanlon, and others:

What quasi-realism mimics is not tempered realism as a whole but tempered realism in all but one respect We [quasi-realists] can’t mimic the claim that understanding normative properties and relations as objective matters of fact is basic to explaining how moral judgments of wrongness work.

(Gibbard 2011: 44)

In the terminology that I’ve used, Gibbard maintains in this passage that quasi-realism mimics tempered realism but not in every respect, since quasi-realism rejects Explanation, which implies that moral reality explains elements of the core moral data.

There are some important differences between these passages from Blackburn and Gibbard that I’ll discuss in a moment. For now, I want to focus on what they have in common: each passage states that quasi-realism ends up saying things that sound very much like realism because the quasi-realist engages in mimicry (Blackburn 1993: 4, 15; Gibbard 2011: 45, 47; 2003: 9, 43, 81, 112, 180–81, 220). What is it for a metaethical approach to engage in mimicry? And what would it be for such mimicry to be successful? Quasi-realists say less about these matters than one might anticipate, given how prominent the activity of mimicry seems to be to their project. And some of what they say is perplexing (compare Blackburn 2005b: 323 and 1993: 16, 52). But even without offering an informative answer to these questions, we can say some informative things about *how* quasi-realists endeavor to mimic what Blackburn calls “the apparently ‘realist’ appearance of ordinary moral thought” (Blackburn 1993: 151). In the first place, mimicking the realist appearances consists in formulating and defending an elaborate “logic of attitudes,” which endeavors to establish that attitudinal states can behave more or less exactly like ordinary beliefs (Hale 1993; Schroeder 2008). If this project were successful, it would help quasi-realists to accommodate and explain data such as:

- (i) Moral judgments have the marks of a descriptive belief: they are classificatory, truth-evaluable, apt candidates for knowledge, and apt for inference.

Second, in order to mimic the realist appearances, quasi-realism embraces deflationist or “minimalist” views regarding representation, truth, and facthood. This approach is manifest in passages such as the following in which Gibbard writes:

In one sense there clearly are “facts” of what a person ought to do, and in a sense of the word “true” there is a truth of the matter. That’s a minimalist sense, in which “It’s *true* that pain is to be avoided” just amounts to saying that pain is to be avoided – and likewise for “It’s a fact that.”

(Gibbard 2003: x; cf. Gibbard 2012: 49; cf. Blackburn 1998: 79)

For reasons that will emerge in the final section, it is a delicate matter to state exactly what the theoretical benefits of propounding deflationary views regarding representation, truth, and facthood are supposed to be. The passages just quoted seem to suggest that “going deflationary” allows quasi-realists to rather closely mimic the sorts of things that nonnaturalist realists say when stating their views. And that, in turn, may help quasi-realists to affirm the legitimacy of certain areas of ordinary moral thought and discourse in which agents claim not only that torture is wrong but also that it is *true* (or that it is a fact) that it is wrong.

So far, we’ve canvassed two strategies that quasi-realism employs to mimic “realist-sounding” moral thought and discourse: developing a logic of attitudes and embracing deflationism with regard to representation, truth, and facthood. This helps us to understand the means by which quasi-realism engages in mimicry. But we don’t yet have a clear understanding of what the central aim of the quasi-realist project is, and what the view would affirm were it to achieve its central aim. As indicated earlier, answering these questions is more challenging than it might seem because each question reveals an ambiguity that must be addressed if we’re to have a firm grip on the aims and claims of the quasi-realist project.

Let’s begin with the first ambiguity. A close look at the passages I’ve quoted from Blackburn and Gibbard reveals that these philosophers formulate their views in importantly different and not clearly compatible ways. Moreover, not only are these formulations different, some incorporate puzzling assumptions about how theories relate to what they are supposed to explain.

Begin with the passage from Blackburn that I’ve already quoted twice. In this passage, Blackburn writes that the quasi-realist finds himself progressively able to mimic “the thoughts and practices supposedly definitive of realism” – where is it clear that by this last phrase Blackburn means the “central elements of our ethical thought and practice” (Blackburn 1993: 4). Elsewhere, Blackburn characterizes quasi-realism differently, writing that the view endeavors to mimic “the realist-sounding discourse within which we promote and debate (our) views,” and to show “how much of the apparently ‘realist’ appearance of ordinary moral thought is explicable and justifiable on an anti-realist picture” (Blackburn 1994: 315; 1993: 151; cf. Blackburn 1993: 77; 1994: 180). Let’s call these formulations, which state that quasi-realism aims to mimic the “realist-sounding” or “realist-seeming” dimensions of ordinary moral thought and practice, the *practice-mimicking* formulations of quasi-realism.

Gibbard’s formulations of quasi-realism are different. When Gibbard characterizes quasi-realism, he typically states that its aim is to mimic, not realist-sounding ordinary moral thought and discourse, but the particular metaethical theses to which realism commits itself (see Gibbard 2011). These theses include claims about how moral concepts work (they represent moral reality) and the existence and character of moral facts (they are among those things that moral concepts represent). Let’s call these formulations, which state that quasi-realism aims to mimic the elements of realism itself, the *theory-mimicking* formulations of quasi-realism.

The practice-mimicking and theory-mimicking formulations of quasi-realism might seem like minor variations on a common theme. But they are not. They are formulations that can diverge in important respects.

Consider the practice-mimicking formulations to start with. Note that quasi-realists could coherently accept these formulations but reject the theory-mimicking formulations

of their view. They might do this, moreover, because they hold that much of what realism says – such as that there is a “world of norms and Forms” – is false (Blackburn 2005a: 59). Since (in their view) what realism says is by-and-large false, there would be no point in mimicking realism’s claims and commitments. This is, in fact, what some of those sympathetic with the view say (PEA Soup 2015).

Now turn to the theory-mimicking formulations. Note that quasi-realists could coherently accept these formulations but reject the practice-mimicking formulations. They might do this for two reasons. First, quasi-realists might deny that ordinary moral thought and practice is “realist-sounding,” holding, for example, that such thought and discourse is largely metaethically neutral (Carter and Chrisman 2012). And, second, quasi-realists might affirm – contrary to what some advocates of the practice-mimicking formulations say – that there is excellent reason for quasi-realism to mimic what realism says. After all, realism commits itself to claims such as:

It is a fact that torture is wrong.

And:

We know that torture is wrong,

which appear to be true. In addition, realism rejects other claims such as:

Torture is wrong simply because we presently disapprove of it.

And:

Were we to approve of torture, then (in those circumstances) we would know that torturing is right,

which we have excellent reason to take to be false. If realism is correct to affirm the first pair of claims and to reject the second pair, then there would be a clear and principled rationale for developing a theory that endeavors to mimic what realism says.

The practice-mimicking and the theory-mimicking formulations of quasi-realism, then, are importantly different. Not only are they importantly different, the reasons for accepting formulations of one sort can be incompatible with the reasons for accepting formulations of the other sort. The fact that these formulations can have divergent and even incompatible implications and justifications raises interpretive questions for those who want to understand the quasi-realist project. And the interpretive questions do not end here. In some cases, they become more acute the more carefully one examines the formulations themselves.

This is especially true of the practice-mimicking formulations of quasi-realism. Recall that these formulations tell us that quasi-realism endeavors to mimic the thoughts and practices that are allegedly definitive of realism or are “realist-sounding.” With our earlier distinction between the core metaethical data and the theses to which metaethical theories appeal when attempting to accommodate and explain these data, consider the practice-mimicking formulations of quasi-realism once again. (I assume that the core metaethical data include, but are not limited to, what Blackburn calls “central elements of our ethical thought and practice.”)

Take the claim that the core metaethical data are “supposedly definitive of realism.” A moment’s reflection reveals that there are no such data. For what is definitive of realism are not the core metaethical data, but the realist theses. Or consider the formulations of quasi-realism which say that there are “realist-seeming” or “realist-sounding” data. Again, we have excellent reason to hold that there are no such data for the simple reason that the core metaethical data do not seem or sound like the metaethical theses that realism affirms (or, for that matter, the theses that any major metaethical theory affirms). For example, there is no interesting sense in which the datum:

- (ii) Not any response to a moral question will do; we can make moral mistakes. Moreover, some answers to moral questions are better than others,

seems or sounds like any of the realist theses. To say it again, realism endeavors to *explain* this datum (and others like it) by appeal to the realist theses. In general, however, there is little reason to suppose that what explains some range of data sounds or seems like the data themselves (or that what is doing the explaining is somehow “built into” the data).

These observations raise the concern that the practice-mimicking formulations of quasi-realism do not admit of a cogent characterization. Think of the concern like this: suppose that advocates of the practice-mimicking formulations were to concede that there is nothing “realist-like” about the core metaethical data, since these data do not seem or sound like the realist theses. Still, proponents of the practice-mimicking formulations of quasi-realism might insist that there are “realist-friendly” data that realism smoothly accommodates and explains, which quasi-realism aims to mimic. But, in fact, quasi-realists do not attempt to mimic data such as (i), which tells us that moral judgments have the marks of descriptive belief. Rather, they attempt to accommodate and explain them from an antirealist and expressivist basis. To accommodate and explain, however, is not *ipso facto* to mimic. Were realism to accommodate and explain (i) in one way and quasi-realism to do so in another way, it would not follow that realism had thereby mimicked these data (or quasi-realism, for that matter).

Suppose, then, we characterize practice-mimicking formulations of quasi-realism as follows: they describe a position that endeavors to accommodate and explain realist-friendly data from a starkly antirealist and expressivist basis. The problem with this formulation is that there appears to be nothing particularly quasi-realist about the view it describes. After all, the view it describes attempts to neither mimic the realist-seeming data (because there are none) nor the realist theses (since it rejects theory-mimicking formulations of quasi-realism). Moreover, the mere fact that this view endeavors to accommodate and explain realist-friendly data from an antirealist starting point does not appear to render it distinctively quasi-realist. After all, rival metaethical views, such as versions of constructivism and error theory, do the same (see Blackburn 1998: chapter 4; 2005a: chapter 5; Olson 2014).

We have been discussing how best to characterize the central aim of the quasi-realist project. When looking at passages in which quasi-realists state their view, we noticed that they characterize it differently. Some endorse practice-mimicking formulations, while others accept theory-mimicking formulations. We’ve also seen that these formulations diverge in important ways. Practice-mimicking formulations, moreover, appear to incorporate some puzzling assumptions about the relations in which theories and the data that

they try to accommodate and explain stand. In fact, I’ve claimed that when we try to free these formulations from these assumptions, it is unclear that what remains is quasi-realist in any interesting sense.

Admittedly, all metaethical theories incorporate ambiguities. Some matter and some do not. In drawing attention to the difference between practice-mimicking and theory-mimicking formulations of quasi-realism, I have tried to identify an ambiguity that matters. It matters because metaethical theories are often accompanied by two sorts of adequacy conditions. One sort of adequacy condition is internal: it specifies the conditions that a theory must satisfy by its own lights. For example, theory-mimicking formulations of quasi-realism state that quasi-realism must thoroughly mimic the realist theses. Another type of adequacy conditions is external: it specifies the conditions that any metaethical theory must satisfy to be worthy of acceptance. Satisfying these conditions, for example, might consist in a theory accommodating and explaining the core metaethical data. We can determine whether quasi-realism satisfies the first sort of adequacy condition, however, only if we know what the aim of quasi-realism is – what it endeavors to accomplish. And this, we’ve seen, is not apparent.

Let us now turn to the second ambiguity, which concerns not quasi-realism’s central aim, but what the view would commit itself to, were it to satisfy its central aim. The concern that quasi-realists ought to be wary of whether their view satisfies its central aim has long animated critics. The concern can be expressed in a dilemma: either quasi-realism successfully mimics realism or it does not. If the project is successful, then quasi-realism ceases to be a distinctive position: it collapses into a version of realism (see Dreier 2004). If the project is not successful, then it fails to satisfy its own aims. Either option is unacceptable. The first is unacceptable because quasi-realism is supposed to be an alternative to realism that is immune to the problems that afflict realism. The second option is also unacceptable because then the view fails to satisfy adequacy conditions of its own devising (see Wright 1988; Dworkin 1996).

While this dilemma deserves consideration, I am concerned not to press it but to explore a question to which it directs our attention, which is this: Suppose quasi-realists were to execute their project successfully. What claims would it thereby commit both itself and ordinary moral agents to? I am going to suggest that, on this issue, the textual evidence pulls us in two different directions. Some texts suggest that quasi-realists commit themselves to a position I call *Thin expressivism*, while other texts suggest that quasi-realists commit themselves to a different and incompatible position that I call *Thick expressivism*.

To appreciate the differences between these views, consider a moral sentence such as:

- (E) “It is a fact that torture is wrong.”

According to *Thin expressivism*, to sincerely utter (E) is:

- (1) to express an attitudinal state (say, that of condemning torture).

And:

- (2) not to state that there is (or otherwise commit oneself to the existence of) a moral fact, namely, *that torture is wrong*.

Thin expressivism holds that an agent would make no mistake by not committing herself to the existence of moral facts when uttering (E), for it embraces:

- (3) either there are no moral facts or it is not possible to state that there are such facts using standard moral (or metaethical) discourse.

In contrast, according to *Thick expressivism*, to sincerely utter (E) is:

- (4) to express an attitudinal state.

And also:

- (5) to state or otherwise commit oneself to the existence of moral facts.

Thick expressivism maintains that an agent would make no mistake by committing herself to the existence of moral facts when uttering (E) because it holds:

- (6) there are moral facts (although they do not explain the core metaethical data).

In what follows, I want to look at some passages in which quasi-realists appear to endorse Thin expressivism. I'll then turn to passages in which they appear to endorse Thick expressivism.

In *Ruling Passions*, Blackburn writes:

The theory I want to defend is one that gives a story about the way in which ethical thought functions. Valuing something, it says, is not to be understood as *describing* it in certain terms, any more than hoping for or desiring something are describing it in particular terms.

(Blackburn 1998: 49)

On the assumption that to value something is to express one's thoughts "in terms of what is good, bad, obligatory, right, justifiable," this passage appears to commit quasi-realism to claims (1) and (2) stated above.

But if there were moral facts, wouldn't this position commit ordinary agents to a mistake of sorts – a mistake that consists not in misrepresenting ethical reality, but in failing to represent it and be guided by it? In various places, Blackburn explains why such agents would commit no mistake:

Protected by quasi-realism, my projectivist says the things that sound so realist to begin with – that there are real obligations and values, and that many of them are independent of us, for example He *affirms all that could ever properly be meant by saying that there are real obligations ...* . It is just that the explanation of why there are obligations ... is not quite that of untutored common sense. It deserves to be called anti-realist because it avoids the view that when we moralize we respond to, and describe, an independent aspect of [moral] reality.

(Blackburn 1993: 157)

Elsewhere, Blackburn states more forthrightly why an agent who sincerely uttered (E) but did not commit herself to there being moral facts would make no mistake. Discussing the issue of whether his view commits itself to the thesis that actions are wrong because we disapprove of them, Blackburn writes:

According to me, there is only one proper way to take the question "On what does the wrongness of wanton cruelty depend?": as a moral question, with an answer in which no mention of our actual responses properly figures. There *would* be an external (i.e., metaethical) reading if realism were true. For in that case there would be a fact, a state of affairs (the wrongness of cruelty) whose rise and fall and dependency on others could be charted. But antirealism acknowledges no such state of affairs and no such issue of dependency.

(Blackburn 1993: 173)

In these two passages, Blackburn appears to affirm (3), stating either that there are no moral facts (as in the second passage quoted) or that talk of such facts must be "heard in an ethical tone of voice" in which we are merely "trying to express and systemize our actual values" (as in the first passage quoted) (Blackburn 1998: 50). In fact, in various places, Blackburn suggests that there is no such thing as distinctively metaethical discourse regarding a range of putatively central metaethical questions and that, in fact, moral realism is not really a philosophical position at all (Blackburn 2010). For any attempt to formulate or defend realism, Blackburn says, is simply a matter of expressing attitudinal states: "this is why any 'cognitivism' or 'realism' that supposes itself to be protecting (moral objectivity) actually marks no philosophical position. It amounts, at best to an ethical commitment: something like an injunction to get involved" (Blackburn 1998: 297; Blackburn 2010).

Thin expressivism is not, however, the only interpretation of quasi-realism that receives textual support. Consider, for example, what Blackburn says in this passage:

Now, if the projectivist adopts quasi-realism, he ends up friendly to moral predicates and moral truth. He can say with everyone else that various social arrangements are unjust, and that it is true that this is so. Once this is said, no further theoretical risks are taken by saying that injustice is a *feature* of such arrangements, or a quality that they possess and that others do not. The first step, in other words, is to allow propositional forms of discourse, and once that is done we have the moral predicate, and features are simply abstractions from predicates.

(Blackburn 1993: 206; cf. 8, 181)

Elsewhere, Blackburn expands upon this idea:

Yes, I am an anti-realist; no, this does not mean that there are no facts of an ethical or normative kind.

Quasi-realism ... refuses to give ethical facts a typical explanatory role. This is already heralded when we turn our back on ethical representation. A representation of something as F is typically explained by the fact that it is F. A

representation *answers* to what is representational. I hold that ethical facts do not play this explanatory role.

(Blackburn 1999: 216)

Far from denying that moral facts exist, these passages find Blackburn affirming that they exist. As such, these passages appear to commit quasi-realism to (5) and (6). In the first passage, Blackburn seems to commit quasi-realism to the claim that uttering a sentence such as (E) commits an agent to things having moral properties (but not necessarily to the claim that the agent thereby describes these properties when engaging in moral thought and discourse). In the second passage, Blackburn affirms that there are moral facts but denies that they play the explanatory roles attributed to them by realists.

In some places, Gibbard appears to go even further than this, writing that, according to quasi-realists, “normative terms like ‘ought’ ... signify natural properties.” Using Sidgwick’s position as an example, Gibbard suggests that “ought” might refer to the property “being *unihedonic*.” Of course this observation is consistent with Thin expressivism; advocates of this view needn’t deny that moral predicates refer to natural properties. However, Gibbard goes on to claim that quasi-realism can affirm that “the property being what one ought to do just *is* the property of being unihedonic” (Gibbard 2002: 56). That is, Gibbard appears to affirm that there is a property *being what one ought to do*, which actions have, and it is identical with a natural property. In saying this, Gibbard appears to commit quasi-realism to a claim that is stronger than (6). The difference between his view and Blackburn’s is that Gibbard does not claim that moral properties are “abstractions” from or “shadows” of moral predicates. Rather, Gibbard affirms that they are simply natural properties that are part of the causal order and, hence, explanatorily robust. (I’ll return to this position at the end of this discussion.)

The question that we are exploring is what quasi-realism would commit itself to were it to successfully journey from its stark antirealist starting point to its more realist-friendly end point. Thin expressivism provides one answer; Thick expressivism provides another. These answers are not compatible. Therefore, any attempt to assess quasi-realism must address the questions of whether one of these positions is the best interpretation of quasi-realism or whether there is no unified quasi-realist program but only versions of quasi-realism – some staying close to the program’s antirealist roots, others venturing very close to (if not collapsing into) a version of realism. Like the first ambiguity we considered, this ambiguity is also important. Not only would it be illegitimate for quasi-realists to slide back and forth between these two versions of their view when responding to objections (as might be feared), any assessment of quasi-realism must also pay close attention to which version of quasi-realism is under consideration. Which version of quasi-realism is under consideration can, moreover, have important ramifications regarding how highly one rates the view’s prospects.

ASSESSING QUASI-REALISM

A natural conclusion to draw from the preceding section is that we are in no position to assess the viability of the quasi-realist program because it is unclear how to understand the program’s aims and commitments. Indeed, one might conclude that there is no

unified view or research program designated by the term “quasi-realism.” There are only a variety of incompatible views flying under the banner of quasi-realism that must be distinguished and considered on an individual basis.

Rather than conclude that we should forestall any assessment of quasi-realism until we settle this interpretational situation, let me try to push the discussion forward by voicing a challenge to the view (or family of views). In voicing this challenge, we need to keep in mind that it may apply only to some versions of quasi-realism.

Quasi-realism’s advocates often present the view as representing a metaethical via media, an approach that strikes a path between equally unattractive metaethical options. On the one hand, quasi-realism is supposed to represent an alternative to the realist’s way of theorizing in which realism simply helps itself “to normative facts at the outset” (Gibbard 2003: 183). On the other hand, quasi-realism is designed to be a viable antirealist alternative to error theory and subjectivism (Blackburn 1993: chapters 8, 9). Critics have charged, however, that the contrast between quasi-realism and these rival antirealist views is much less pronounced than it might seem: *all* these views are subject to the objection that they fail to accommodate and explain some deeply entrenched features of the moral life (see FitzPatrick 2010; Parfit 2011: chapter 28).

Perhaps the best way to articulate the challenge in question is to return to the core metaethical data and, in particular, the datum that:

- (ii) Not any response to a moral question will do; we can make moral mistakes. Moreover, some answers to moral questions are better than others.

Let us dub this the *Mistake datum*. One point of deep dissatisfaction with both error theoretic and subjectivist positions is that they do not accommodate this datum. True, the error theory, which holds that all of our moral judgments are untrue, can accommodate the fact that we make moral mistakes. This is because the view implies that *all* our moral judgments are mistaken. But the view also implies that no moral response is better than any other: all are equally erroneous. Subjectivism, which holds that what makes an act right (for an agent) is that she approves of it and that what makes an act wrong (for an agent) is that she disapproves of it, is (in this regard) the mirror image of error theory. For, given the plausible assumption that we tend to know what we approve of and disapprove of, subjectivism implies that (nearly) all our moral judgments (which concern what we approve of and disapprove of) are correct. It also implies that it is very difficult to make moral mistakes. Most importantly, subjectivism implies that it is impossible for one response to a moral question to be better than another: they are all equally good.

Quasi-realism rejects both error theory and subjectivism. So, critics cannot plausibly charge that it fails to accommodate the Mistake datum for the same reasons that these views fail to. Even so, it may be that quasi-realism fails to accommodate the Mistake datum for different reasons.

Let us pursue this possibility by considering Thin expressivism once again. Thin expressivism offers an account of what it is to say or think that performing an action is wrong: to think or say that an action is wrong is to be in, or to express, an attitudinal state. Thin expressivism also tells us what it is to say or think things such as:

It is a fact that torture is wrong.

And:

It is true that torture is wrong.

Given their deflationary or minimalist commitments, Thin expressivists typically tell us that to say or think these things is simply to say or think:

Torture is wrong,

which is itself an attitudinal state or the expression thereof.

But it is one thing for a position to affirm that we can say or think that it is true (or correct) that torture is wrong, and to explain what it is to say or think such things. It is another thing for a view to imply that it *is* true (or correct) that torture is wrong, and to explain *what it is* for it to be true (or correct) that torture is wrong. Admittedly, Gibbard seems to run the two issues together when he writes:

In one sense there clearly are “facts” of what a person ought to do, and in a sense of the word “true” there is a truth of the matter. That’s a minimalist sense, in which “It’s *true* that pain is to be avoided” just amounts to saying that pain is to be avoided – and likewise for “It’s a fact that.”

(Gibbard 2003: x)

But the issues are different, and we can draw no conclusions about whether there are moral facts or truths by appealing to minimalist or deflationist accounts of what it is to think or say something factual or true, as this passage appears to.

We are now positioned to formulate what I’ll call the *No Mistake challenge*:

- (A) If Thin expressivism fails to accommodate and explain the Mistake datum, then there is an important respect in which it is identical with other antirealist positions such as error theory and subjectivism.
- (B) Thin expressivism fails to accommodate and explain this datum.
- (C) So, there is an important respect in which Thin expressivism is identical with other antirealist positions such as error theory and subjectivism: it fails to accommodate and explain the Mistake datum.

Now suppose we assume:

- (7) If a metaethical position fails to accommodate and explain the Mistake datum, then that is *prima facie* reason to reject it.

It follows that we have *prima facie* reason to reject error theory, subjectivism, and Thin expressivism because none of them accommodates and explains the Mistake datum.

Let’s consider four replies that quasi-realists might offer in response to the No Mistake challenge. These responses range from those that are not at all concessive to those that are quite concessive to realism.

The first reply: The No Mistake challenge assumes that a metaethical position could attempt to accommodate and explain the Mistake datum by appealing to moral facts. Thin expressivism rejects this assumption, at least when it is understood along non-expressivist

lines. Instead, Thin expressivism holds that all talk of moral facts and moral mistakes must be understood simply to be the expression of attitudinal states. The impression that there is a cogent non-expressivist account of such talk is illusory.

Here is an argument for this claim, which might be implicit in some of the passages quoted earlier from Blackburn (see Blackburn 1993: 174). Suppose Thin expressivism were true, offering an accurate account of the character of moral/metaethical thought and discourse across the board. If it were, then moral/metaethical sentences would not even purport to represent moral reality or standards that could render some moral judgments mistaken; they would express attitudinal states. But when non-expressivists, such as realists and constructivists, engage in moral/metaethical discourse, they use moral/metaethical sentences. So, they cannot be stating moral propositions (i.e., moral representational contents) when using these sentences. Instead, they must be using these sentences to express attitudinal states. If this is right, then the claim that all moral/metaethical discourse must express attitudinal states is simply an implication of Thin expressivism.

The argument is problematic. It is problematic not just because it implies that realists and other “factualists” do not know what they are saying when they say that there are moral facts or that we make moral mistakes, which looks deeply uncharitable (see Cuneo n.d.). The argument is also problematic because its premises do not imply its conclusion. To see why, distinguish *standard* from *non-standard* moral/metaethical discourse. Such discourse is standard, let’s assume, just in case it conforms to the ordinary use of moral/metaethical sentences. Otherwise, it is non-standard. Now suppose, for argument’s sake, that Thin expressivism offers us an account of standard moral/metaethical discourse and that its advocates engage in such discourse when developing their position.

Note that these assumptions are compatible with the following two claims. First, realism offers us an account of non-standard moral/metaethical discourse and its proponents engage in such discourse when developing their position. Second, they have good reason to, because standard moral/metaethical discourse is deeply problematic, lacking expressive power. Among other things, such discourse lacks the resources that would enable its users to talk about moral reality (however that is understood) and moral mistakes. It is probably worth adding that if the scenario I am presenting were to describe the actual character of contemporary metaethical debate, it follows that there would be no genuine moral or metaethical disagreement between Thin expressivism and realism (and nearly all other metaethical positions). At best, there would be a meta-semantic disagreement about how to engage in moral/metaethical discourse (see Plunkett and Sundell 2013).

The second reply: The No Mistake challenge overlooks an important sense in which expressivism is different from other antirealist positions. The primary difference is that expressivists employ “the core expressivist maneuver” (Carter and Chrisman 2012). Think of the maneuver as having three steps. The first step is to switch the subject: rather than concern itself with the question of what it is for something to be a moral fact or truth, expressivism concerns itself with what it is to express a normative judgment. The second step is to remind us that expressivism offers a distinctive account of what moral thought and discourse are, which rejects the claim that moral judgments have moral representational content. The third step consists in expressivists concluding that they (in their role as theorists) can “just stop talking” about moral facts or the accuracy of moral judgments and, instead, issue first-order moral judgments “inside the domain of ethics” (Carter and Chrisman 2012: 334; cf. Blackburn 1998: 50; Gibbard 2003: x). By ceasing to

talk about the topic of moral facts or truths, advocates of the core expressivist maneuver needn't claim that these facts or truths do not exist. They can simply note that they are not of ethical interest, and that the view can remain agnostic about whether they exist (Carter and Chrisman 2012).

This reply fails to engage the No Mistake challenge. Remaining agnostic about the existence of moral facts or truths does not accommodate or explain the Mistake datum. It goes no distance toward explaining how agents could make moral mistakes if quasi-realism were true (as opposed to explaining how we can think and say that there are such mistakes). Since it fails to do this, when taken on its own, the core expressivist maneuver does not explain why quasi-realism fares better with regard to the No Mistake challenge than error theory or subjectivism.

The third reply: The No Mistake challenge is directed at Thin expressivism. Quasi-realists should reject this view, opting instead for Thick expressivism. The benefit of doing so is that critics cannot charge quasi-realism with failing to account for the possibility of moral mistakes because quasi-realism denies that moral facts exist. Thick expressivism does not reject but affirms their existence.

This reply also fails to engage the No Mistake challenge. The No Mistake challenge charges Thin expressivism with being unable to explain how we could make moral mistakes. This charge also applies to Thick expressivism. This is because Thick expressivism endorses No Explanation, which tells us that moral facts fail to play explanatory roles, which enables a theory to accommodate and explain the core metaethical data, such as the Mistake datum.

The fourth reply: Quasi-realists should opt for a version of Thick expressivism according to which moral facts *can* explain various elements of the core metaethical data. Since these facts are simply natural facts, this admission does not compromise the integrity of the quasi-realist project.

We have seen that there are passages in which quasi-realists seem to endorse this view (Gibbard 2002; cf. Ridge 2014). I'll make only two observations about it in closing.

The first is that this position is not clearly quasi-realist. It does not "start from an anti-realist position." Rather, it helps itself to the natural world and all its denizens, including among its denizens the moral facts. It is true that this version of Thick expressivism does not invoke these facts in order to do exactly the same explanatory work that they do in a realist scheme. But these facts do explanatory work nonetheless: they explain how we can make moral mistakes. Indeed, if appeal to such "mistake conditions" is necessary to explain how attitudinal states could behave like ordinary beliefs – as Schroeder (2008: chapter 10) argues – then this view would appeal to these facts *at the outset of theorizing* to explain the workings of moral thought and discourse. This, however, is exactly what quasi-realists maintain that their view does not do. To which I'll add, that once moral facts are added to the quasi-realist picture, many of the moves quasi-realists make, in response to certain kinds of objections, would be puzzling at best.

Take, for example, the accusation that quasi-realism cannot account for the fact that kicking dogs for fun would be wrong even if we approved of it. In reply to this accusation, quasi-realists have appealed to metalinguistic/semantic theses about what we are saying when we say that kicking dogs would be wrong even if we approved of it (Blackburn 1994: 218; Dreier 2012). This type of reply would be unnecessary – and puzzling! – if Thick expressivism were true. It is enough for Thick expressivists to point out that the relevant

moral facts are attitude-independent, and they settle the issue of whether kicking dogs would be wrong in those circumstances that we approved of it.

The second observation is that quasi-realists have assiduously avoided delving into such issues as the ontology of moral facts, focusing nearly exclusively on what it is to think or express moral thoughts. In fact, a persistent selling point of quasi-realism is that it can avoid wading into the waters of moral ontology (Blackburn 1993: 171; cf. Schroeder 2010: 8). But if quasi-realists are just reductive naturalists about moral reality, then they will have to address these issues and respond to the many objections raised against reductive naturalism. In a way, embracing this last Thick expressivist reply to the No Mistake challenge commits quasi-realism to a two-flank battle: a battle against those who reject an expressivist account of moral thought and discourse, and those who reject a reductive naturalist metaphysics of moral reality. Given the hard work that remains to be done on both these issues, it may be that the attempt to vindicate quasi-realism – assuming that there is a unified approach that answers to the label – has just begun.

ACKNOWLEDGMENT

Thanks to the editors and Jamie Fritz for their comments on an earlier draft of this chapter.

REFERENCES

- Blackburn, S. (1993) *Essays in Quasi-Realism*, Oxford: Oxford University Press.
- Blackburn, S. (1994) "Quasi-realism," *The Oxford Dictionary of Philosophy*, Oxford: Oxford University Press: 315.
- Blackburn, S. (1998) *Ruling Passions*, Oxford: Oxford University Press.
- Blackburn, S. (1999) "Is Objective Moral Justification Possible on Quasi-Realist Foundations?" *Inquiry* 42: 213–28.
- Blackburn, S. (2005a) *Truth: A Guide for the Perplexed*, Oxford: Oxford University Press.
- Blackburn, S. (2005b) "Quasi-realism Is No Fictionalism," in M. Kalderon (ed.) *Fictionalism in Metaphysics*, Oxford: Oxford University Press.
- Blackburn, S. (2009) "Truth and A Priori Possibility: Egan's Charge against Quasi-Realism," *Australasian Journal of Philosophy* 87(2): 201–213.
- Blackburn, S. (2010) "Truth, Beauty, and Goodness," in R. Shafer-Landau (ed.) *Oxford Studies in Metaethics*, Vol. V, Oxford: Oxford University Press: 295–314.
- Carter, J. A. and Chrisman, M. (2012) "Is Expressivism Incompatible with Inquiry?" *Philosophical Studies* 159: 323–39.
- Cuneo, T. (n.d.) "The Expressivist Gap."
- Dreier, J. (2004) "Meta-ethics and the Problem of Creeping Minimalism," in J. Hawthorne (ed.) *Philosophical Perspectives* 18, *Ethics*, Oxford: Blackwell: 23–44.
- Dreier, J. (2012) "Quasi-Realism and the Problem of Unexplained Coincidence," *Analytic Philosophy* 53: 269–87.
- Dworkin, R. (1996) "Objectivity and Truth: You'd Better Believe It," *Philosophy and Public Affairs* 25: 87–139.
- Egan, A. (2007) "Quasi-Realism and Fundamental Moral Error," *Australasian Journal of Philosophy* 85: 205–219.
- FitzPatrick, W. (2010) "Ethical Non-naturalism and Normative Properties," in M. Brady (ed.) *New Waves in Metaethics*, London: Palgrave: 7–35.
- Gibbard, A. (2002) "Reasons of a Living Being," *Proceedings and Addresses of the Aristotelian Society* 76: 45–60.

- Gibbard, A. (2003) *Thinking How To Live*, Cambridge, MA: Harvard University Press.
- Gibbard, A. (2011) "How Much Realism?" in R. Shafer-Landau (ed.) *Oxford Studies in Metaethics*, Vol. V, Oxford: Oxford University Press.
- Gibbard, A. (2012) *Meaning and Normativity*, Oxford: Oxford University Press.
- Hale, B. (1993) "Can There Be A Logic of Attitudes?" in J. Haldane and C. Wright (eds.) *Reality, Representation, and Projection*, Oxford: Oxford University Press: 337–63.
- Olson, J. (2014) *Moral Error Theory: History, Critique, Defense*, Oxford: Oxford University Press.
- Parfit, D. (2011) *On What Matters*, Vol. II, Oxford: Oxford University Press.
- PEA Soup (2015) "Does Expressivism Have a Knowledge Problem?" March 16. Available at <http://peasoup.typepad.com/peasoup/2015/03/does-expressivism-have-a-knowledge-problem-by-featured-philosopher-terence-cuneo.html>.
- Plunkett, D. and Sundell, T. (2013) "Disagreement and the Semantics of Normative and Evaluative Terms," *Philosophers' Imprint* 13. <http://hdl.handle.net/2027/spo.3521354.0013.023>.
- Ridge, M. (2014) *Impassioned Belief*, Oxford: Oxford University Press.
- Schroeder, M. (2008) *Being For*, Oxford: Oxford University Press.
- Schroeder, M. (2010) *Noncognitivism*, London: Routledge.
- Street, S. (2011) "Mind-Independence without the Mystery: Why Quasi-Realists Can't Have It Both Ways." in R. Shafer-Landau (ed.) *Oxford Studies in Metaethics*, Vol. 6, Oxford: Oxford University Press: 1–32.
- Sturgeon, N.L. (1991) "Contents and Causes: A Reply to Blackburn," *Philosophical Studies*, 61: 19–37.
- Wright, C. (1988) "Realism, Antirealism, Irrrealism, Quasirealism," in P. French, T. Uehling, and H. Wettstein (eds.) *Midwest Studies in Philosophy* 12: 12–49.

41

Metaethical Quietism

Doug Kremm and Karl Schafer

Much like fanatics and hipsters, quietists are seldom quick to describe themselves as such. So it's unsurprising that contemporary metaethicists tend to avoid describing themselves in such terms. Indeed, many of the authors discussed below would actively oppose any such label. And yet analytic metaethics has often been haunted by the specter of various forms of metaethical quietism. For it is distinctive of quietism that it calls into question, if not the very intelligibility of much of metaethics, then at least its point.

In this essay, we begin by discussing how to understand metaethical quietism. In doing so, we distinguish several attitudes towards canonical metaethical questions that might be regarded as quietist. Then we discuss some prominent motivations for quietism—before turning to some recent objections to these views. We conclude by saying a bit about the relationship between these issues and more general questions of philosophical methodology. In discussing these issues, we will attempt to be as charitable as possible to the quietist—not because we are necessarily quietists ourselves—but because we want these views to have as fair a hearing as possible.

CHARACTERIZING QUIETISM

As we understand it, quietism in its various forms can be characterized in terms of the rejection of some range of questions as unworthy of philosophical debate. As a first pass, then, a metaethical quietist can be understood as someone who rejects at least a significant subset of metaethical questions in this way.

So understood, as a question-directed attitude (Friedman 2013), quietism can be made precise in at least two ways: First, with respect to the precise sort of "rejection" it involves, and, second, with respect to the set of questions it rejects.