

Quasi-Realism and Fundamental Moral Error¹

Andy Egan

Australian National University/University of Michigan

Introduction

One of the tasks of metaethical theorizing is to provide a metaphysical underpinning for our ordinary moral thought and practice. This is hard to do, because the desiderata for theories in this area tend to pull against each other. Part of what determines the plausibility of some proposed metaphysical underpinning for some bit of our thought and practice is the extent to which it *justifies* our thought and practice in that area—whether, given that underpinning, our projects, concerns, and activities in the relevant area seem reasonable or not. Other, more straightforwardly metaphysical concerns also contribute to plausibility. This can lead to some tension—the desire to provide an account that’s plausible on *metaphysical* grounds can pull against the desire to provide an account that’s plausible on *justificatory* grounds.

One way to motivate an expressivist metaethics, for example, is by appealing to concerns about the possibility of giving a naturalistically respectable account of moral facts and properties. If you think that the existence of moral facts and properties is difficult or impossible to square with a desirable sort of philosophical naturalism, an account of our moral thought and talk that doesn’t require that there be any such facts or properties will be extremely attractive.²

¹ Thanks to Ralph Wedgwood, Michael Smith, Tyler Doggett, James John, Daniel Korman, Judith Thomson, Alex Byrne, Ned Hall, Ned Markosian, Robert Stalnaker, Terence Cuneo, Peter Railton, the WWU faculty seminar, and audiences and commentators at the University of Colorado Student Philosophy Conference, the LSU Ethics Symposium, the Australian National University, the University of Auckland, and the University of Sydney, and to the anonymous referees for this journal, for helpful comments, suggestions, and objections.

² See for example [Mackie 1977] and [Ayer 1952] for some classic statements of this sort of worry. It’s not clear that the metaphysical concerns about moral facts and properties are really justified. One reason to

Naïve expressivism delivers a nice naturalist metaphysics, but it doesn't do such a great job of justifying ordinary moral thought and practice. We would like our metaethical theories to make sense of the sorts of things we ordinarily do, say, and worry about. We would like an explanation of what we're up to when we engage in moral thought and discourse such that that's a reasonable thing for us to be up to, and an account of what we're concerned about such that that's a reasonable thing to be concerned about. The naïve expressivist account of what we're up to when we're engaged in moral thought and discourse doesn't do that very well. (For example, we have to give up the truth and falsity of moral claims and the possibility of genuine moral disagreement, and we have a big problem with making sense of embedded occurrences of, and inferences involving, moral claims.³)

Simon Blackburn's [1985; 1993b; 1998] quasi-realism is an attempt to keep the clean metaphysics of the expressivist without forcing us to give up big and important chunks of ordinary moral discourse and practice.⁴ After we've become quasi-realists, we should still be able to go in for all of the same realistic-sounding moral thought and talk as before, but with a clean naturalistic conscience.

One of the problems with expressivist theories of moral discourse is in accounting for concerns about moral error. If moral discourse serves only to express various sorts of attitudes, it's not obvious how can there be any serious issue about getting things *wrong*.

oppose anti-realist and quasi-realist metaethical theories is on the grounds that they're ill-motivated, since the problems about the metaphysical respectability of moral facts and properties that they're meant to solve aren't genuine problems. For present purposes, though, let's set this aside.

³ See [Geach 1960; 1965] and the ensuing literature. (To pick a small, semi-random sample: [Blackburn 1993a], [Schueler 1998]; [Hale 1986;1993], [Dreier 1994], and [Unwin 1999].) A similar complaint is also made by Searle [1964], and responded to by Hare [1970].

⁴ See [Blackburn 1984; 1993b; 1998].

But there clearly *is* such an issue. This concern about moral error is one that quasi-realists need to address as well.

Blackburn [1998] offers a promising quasi-realist account of moral error—one that deals nicely with very many cases. In what follows, I will argue that there is a sort of *fundamental* moral error such that if we take Blackburn’s line, we are forced to grant the possibility that *other* people are subject to such errors, but cannot make sense of *first-person* worries about such errors. So if I am a quasi-realist, I will be forced to acknowledge a sort of moral error to which other people are subject, but against which I have an *a priori* guarantee of immunity. If we’re concerned to underwrite our moral discourse and practice as it is—if we want to avoid forcing major revisions to our way of thinking and talking about morality—then this is an unacceptable consequence.

1. Expressivism and Quasi-Realism

The central expressivist claim is that moral language serves not to describe the world, but to express some sort of attitude. When I say, “stealing is bad”, I’m not saying anything about how the world is (such as, for example, that there’s some property of *badness* that’s instantiated by, among other things, thefts), but merely expressing some distinctive con-attitude toward stealing. Different expressivists differ about what the relevant con-attitude is. (For ease of presentation, I’m going to use “disapproval” to stand in for the relevant attitude, which probably isn’t the attitude that we ordinarily call “disapproval”.) At least at the beginning of the expressivist story, moral utterances aren’t true or false, and there aren’t any moral facts or properties. The “at least at the beginning of the story” qualification is important, since things get more complicated in the case of

sophisticated expressivist views like Blackburn's, which give an expressivist account of such claims as, "there's a property of *badness* that's instantiated by, among other things, thefts".

As mentioned above, the trouble with expressivism—at least with naïve, pre-Blackburn expressivism—is that it doesn't adequately respect our ordinary ways of talking and thinking about morality. We have to give up, for example, talk about moral properties, moral facts, and the truth or falsity of moral claims, and we seem to lose the possibility of genuine moral disagreement and error.

Blackburn argues that the right kind of expressivist needn't give up all of our talk about moral facts, moral truth, moral properties, and so on. Expressivists can make sense of this sort of talk by becoming quasi-realists (see [Blackburn 1993b; 1998], etc.). The cornerstone of the move to quasi-realism is deflationism about metaethical claims. According to the quasi-realist, all there is to moral truth is the equivalence demonstrated by the T-schema: '*P* is true iff *P*'. The quasi-realist uses this equivalence to give us an account of what we're up to when we claim that certain moral statements are true: saying, for example, "it's true that stealing is wrong" is just the same as—expresses exactly the same attitude as—saying simply "stealing is wrong".

A similar deflationism about facts, properties, and so forth consumes the more elaborate bits of apparently realist moral talk: "it's a fact that stealing is wrong", "stealing instantiates the property of wrongness", "my claim that stealing is wrong resonates with the eternal verities of the universe", etc., are all just equivalent to "stealing is wrong". So, given the expressivist account of "stealing is wrong", they all simply express, in more or less dramatic fashion, my negative attitude toward stealing.

Talk about moral *belief* is brought under the quasi-realist umbrella, as well, in roughly the following way: Believing that stealing is wrong is just taking it to be true that stealing is wrong. What is it to take it to be true that stealing is wrong? Well, it's to have the attitude that one expresses in sincere utterances of "it's true that stealing is wrong". That is, it's to have the attitude that one expresses in sincere utterances of "stealing is wrong". That is, it's to bear the right kind of negative attitude toward stealing.

In this way the quasi-realist works his way back to being able to talk about moral belief, moral truth, moral properties, moral facts, and so on, while keeping his official ontology free from contamination by naturalistically disreputable entities (like moral facts and properties of the sort advocated by realists).⁵ If the project is successful, this is a very nice result for those with expressivist inclinations, since it removes one of expressivism's more serious theoretical costs.

It's worth pausing at this point to reiterate the importance of preserving the appearances in motivating quasi-realism (this will be important later on). Quasi-realists are concerned to underwrite our ordinary ways of talking and thinking about morality. Becoming a quasi-realist isn't supposed to require one to revise anything substantive in one's first-order moral beliefs, talk, or practice. The conversion to quasi-realism is meant to be a strictly *metaethical* conversion, leaving our *ethical* views intact. Quasi-realism is

⁵ It's controversial whether quasi-realism really delivers the benefits it's supposed to (see [Rosen 1998] for arguments that it does not). The quasi-realist hopes to wind up saying that there are moral properties, that moral claims are true, etc., but without being a realist *simpliciter*, as this is supposed to be fraught with horrible metaphysical problems. But, as Rosen points out, it's hard to see what's *quasi* about quasi-realism at the end of the day, when they endorse all the same claims as the realist. There are responses available to the quasi-realist here, and counters available to the critic of quasi-realism, but I'd like to set this issue aside. It's also controversial whether moral properties, facts, etc. as advocated by realists *simpliciter* really are naturalistically disreputable. (See for example [Boyd 1988; Sturgeon 1988; Brink 1989; Smith 1994].) I'd like to set this issue aside, as well. Let's assume for present purposes that quasi-realism really *does* deliver the goods of allowing us to endorse talk about moral facts, moral properties, and so forth, without incurring some naturalistically objectionable commitments that accompany realism *simpliciter*. There is a problem—a fatal one, I think—about moral error even then.

intended to be a naturalistically respectable way of explaining what it amounts to to have some particular moral view, which does not restrict the sorts of moral views that it makes sense to have.

Quasi-realists, and Blackburn in particular, are absolutely right to be concerned about preserving the appearances. Like naïve expressivism, quasi-realism is a theory about how to understand moral discourse. If a theory of how to understand a certain area of discourse undermines big chunks of it, by making nonsense of much of what we say, or by interpreting concerns and debates that look sensible as really being silly, that's a reason to reject the theory.

2. *Moral Error*

It's an obvious and important fact about our ordinary ways of thinking about morality that moral error is possible. We often attribute moral error to other people and to our past selves, and we often worry about the accuracy of our present moral beliefs. An adequate theory of what we're up to when we engage in moral discourse and deliberation should be able to account for this.

Distinguish three kinds of moral error that I might attribute, or be concerned about: *Third person* error, in which somebody else holds some false moral view; *First person* error at some other time, in which I used to, or will in the future, hold some false moral view; and first-person *present* error, in which I presently hold some false moral view.

Third person error, and first-person error at other times, do not, on the face of it, pose a problem for even the least sophisticated expressivist views. Every expressivist seems to have the resources to account for these sorts of error, by saying the following:

When I claim that you mistakenly believe that stealing is permissible, I (roughly) claim that you approve of stealing, and express my disapproval of stealing. When I say that my past (or future) self falsely believed that stealing was permissible, I say that my past (or future) self approved of stealing, and express my present disapproval. No problem there.

The difficulty is with making sense of concerns that we might, even now, be mistaken about some of our current moral beliefs. My concern about whether my current belief that stealing is wrong is mistaken *isn't* a concern about whether or not I really disapprove of stealing. I could be absolutely certain of my disapproval, but still concerned that perhaps my disapproval not appropriate. Concern about present moral error isn't plausibly characterized in terms of concern about whether we really bear the relevant attitudes to the relevant kinds of things. So the quasi-realist needs to provide us with something else to be concerned about when we're concerned about first person, present moral error.

Put another way: there's a possibility that I'm concerned might be actual when I'm concerned about present, first person moral error. What is it? When I say, "I believe that stealing is bad – I hope I'm right", what am I doing, and how would things have to be for my hope not to be realized? The concern here is that the obvious translations of the moves that seemed to work so well for third-person error, and first-person error at other times, are non-starters in this case. (For example, it's clearly not going to work to say that I'm self-attributing a certain attitude toward stealing with "I believe that stealing is bad", and then expressing my hope that I really do have that attitude with "I hope I'm right".) The same problem arises when we attempt to give a quasi-realist account of

expressions of epistemic modesty like, “I think that torture is permissible in certain ticking-bomb cases, but I might be mistaken”.⁶

This is a pressing issue. Concerns about moral error are concerns about *something*, and they are, at least in very many cases, *reasonable* concerns. I take it to be a constraint on the quasi-realist account of worries about moral error that it agree with common sense that people who are worried about moral error aren’t typically just being silly—there really is something there that it’s sensible to worry about.

Blackburn addresses this difficulty, and offers a convincing response. Here is what he says:

⁶ Notice that this really is a problem specifically for expressivist accounts of epistemic modesty *about the moral*, and not a general problem about epistemic modesty about any subject matter whatsoever. The problem is *not* about giving a general semantics for ‘might’, or for utterances in which ‘might’ is deployed for the purpose of expressing epistemic modesty. (I do have views about these issues – see [Egan, Hawthorne and Weatherson 2005; Egan forthcoming] – but I don’t think that they’re particularly relevant to the issue at hand.)

Note also that it’s not a general problem about Moore-paradoxical claims, or anything of that sort. Statements like “Stealing is bad, but I might be mistaken” are, perhaps, Moore-paradoxical (they are at least Moore-suspicious-looking), and are in that respect in exactly the same boat as “There’s a cat on the mat, but I might be mistaken”. But it’s important to notice that those sorts of Moore-suspicious statements *aren’t* the sorts of statements under discussion. The statements we’re concerned with here are ones such as “*I believe that* stealing is bad, but I might be mistaken”. And these are clearly *not* Moore-paradoxical. We make these sorts of statements all the time. Preparing to pass a slow-moving car on the highway, I ask you to check for traffic on the passenger side. You reply, “I don’t think there are any cars coming, but I could be wrong”. A student asks if there is anything in the literature on subject X. The teacher responds, “I believe there’s a paper by Professor Jones, but I might be mistaken”. These are perfectly routine, and totally unproblematic. So the problem clearly isn’t a general problem about Moore-paradoxicality, or the general puzzlingness of epistemic modesty. Epistemic modesty isn’t, in general, either puzzling or Moore-paradoxical.

One use of these sorts of statements of epistemic modesty is to express a fear that we’re in some unfortunate epistemic situation. The problem for the quasi-realist is how to characterize the unfortunate situation that we’re in if such fears are realized. For subject matters about which we are unabashed realists, this is completely straightforward. It’s the situation in which we believe that P, but P is false. Of course the quasi-realist will want to say this, too. The problem is that it’s not clear that he *can*, given his characterizations of moral truth and moral belief – the quasi-realist’s theory doesn’t seem to leave the right sort of gap between *belief* and present *truth*. If it’s in order for me to say that I believe that P (because I really have the relevant attitude), then it’s in order for me to say that P is true (again, because I really have the relevant attitude). And the quasi-realist can’t, without telling some further story, appeal to a truth/assertability gap in order to find a sense in which it’s *in order* to say that P is true, but nonetheless *mistaken*. (Of course, the way to avoid the problem is precisely to tell such a further story, which is what Blackburn does. We’ll come to this in a moment.)

The problem comes with thinking of myself... that I may be mistaken. How can I make sense of my own fears of fallibility? Well, there are a number of things that I admire: for instance, information, sensitivity, maturity, imagination, coherence. I know that other people show defects in these respects, and that these defects lead to bad opinions. But can I exempt myself from the same possibility? Of course not (that would be unpardonably smug). So I can think that perhaps some of my opinions are due to defects of information, sensitivity, maturity, imagination, and coherence. If I really set out to investigate whether this is true, I stand on one part of the (Neurath) boat and inspect the others. [Blackburn 1998: 318]

Before looking at exactly what Blackburn's proposal is here, let's first file away a premise for a later argument. Blackburn says (quite rightly, it seems) that it would be "unpardonably smug" to take ourselves to be immune to this sort of error that we know others are (or might be) subject to. This is an instance of a quite general symmetry requirement on vulnerability to error—I don't have any special, *a priori* guarantees against kinds of moral error to which others are vulnerable. So the premise to file away is:

NO SMUGNESS: There isn't any sort of moral error to which others are subject, but against which I have an *a priori* guarantee of immunity.⁷

⁷ NO SMUGNESS is only plausible as a prohibition on *a priori* guarantees. It might be that you could get some *a posteriori* guarantee that you're immune to some sort of moral error that others are subject to—you could, for example, get very good evidence that you never go in for a certain sort of fallacious moral reasoning that plagues other people. If you've got the evidence, then you're not being smug (or at least, not unpardonably smug) when you leave open the possibility that I'm, say, assigning too great a weight to the happiness of my immediate family, or affirming the consequent, or mixing up my quantifiers, while ruling out the possibility that *you* are making the same kind of error. Similarly, you're not being smug (or at

We'll come back to this, and deploy it in an argument against quasi-realism, in section 5. Notice first, though, that Blackburn's apparent endorsement of NO SMUGNESS is exactly what we should expect, given that one of the major attractions of quasi-realism is its promise as a way to respect our ordinary moral discourse and practice. A metaethical theory that rejected NO SMUGNESS would require a very dramatic deviation from our ordinary ways of thinking about morality. On our ordinary ways of thinking, nobody has a kind of privileged access to the moral truth that others lack. Or at the very least, nobody knows *a priori* that they have a kind of privileged access to the moral truth that others lack.

Blackburn's proposal for how to account for the possibility of moral error shouldn't be too surprising. This is a familiar sort of move – many authors, both realists and anti-realists, who have wanted to ground moral facts (or moral talk) in our attitudes have appealed to some sort of idealization in order to account for the appearance of a gap between what we presently endorse and what's actually right.⁸

On Blackburn's account, when I'm concerned about whether or not my present moral beliefs are correct, I'm concerned about whether or not some improving change would lead me to revise them. So in a particular case, when I express epistemic modesty about my belief that stealing is bad (by saying, for instance, "I think that stealing is bad, but I might be mistaken"), I'm (a) expressing my disapproval of stealing, and (b)

least, not unpardonably smug) when you, based on years of accumulated evidence of your outstanding mathematical abilities, leave open the possibility that I've made a mistake in calculating the tip while ruling out the possibility that you've made the same kind of error. (Though it's worth noting that even in these *a posteriori* cases, you're not really entitled to a *guarantee*, but only to an extremely disproportionate assignment of probabilities.) Thanks to Daniel Korman for discussion on this point.

⁸ See for example [Timmons 1998; Smith 1994]. See also [Lewis 1989] and [Johnston 1989] for such theories of the valuable.

admitting that I'm not certain that I won't (or that I couldn't) undergo some improving change that would make me stop disapproving of stealing.⁹ When I'm concerned that my belief that stealing is bad might be false, I'm concerned that, while I disapprove of stealing, I wouldn't so disapprove if my attitudes, beliefs, etc. went through some improving change.

So what we're concerned about when we're concerned about moral error is whether or not our system of moral beliefs (in the case of worries about moral error in general), or some particular belief (in the case of worries about whether or not some *particular* moral belief is mistaken), would survive a course of improving changes—whether our present attitudes match up with the ones that we would have after some improvement.

This really does seem like a satisfactory account of the phenomenon. It gives epistemic modesty a legitimate role in our ordinary moral practice. It also sounds like the right sort of thing to be worried about when we're worried about the accuracy of our present moral beliefs. In fact, it seems like the *only* thing to say, if one wants to be a quasi-realist (or any kind of expressivist, for that matter). If we're expressivists, we can't say that concerns about moral error are concerns about whether our moral beliefs or attitudes match up with some externally determined moral facts. (Well, if we're quasi-realists, we can *say* this, but we can't wind up appealing to such facts in our official metaethical theorizing about what we're doing when we're engaging in moral discourse.

⁹ This might not be exactly right – for instance, I might be self-attributing the relevant attitude rather than *expressing* it with the “I think...” clause – but (a) the proposal will be somewhere in this immediate neighborhood, and (b) the exact details won't matter for our purposes, so this should be close enough to be going on with.

We need, at the end of the day, to be able to ground everything in an expressivistically respectable story about the content and function of our moral thought and talk.)

So our concerns about error will have to be, somehow or other, *internal* concerns—concerns about how our beliefs and attitudes match up with each other, or how our current beliefs and attitudes match up with the ones we'd have after suitable revision. Some story about stability under a course of improving changes seems to be the only sensible way to cash this out.

3. Divergence and Fundamental Error

Call a belief *stable* just in case no change that the believer would endorse as an improvement would lead them to abandon it. Call a belief *unstable* just in case it's not stable; that is, just in case it would be abandoned after some change that the believer would endorse as an improvement.

Consider two people, Ned and Ted. Ned stably believes that P, while Ted stably believes that not-P. There are two ways that this could happen. First, Ned and Ted could both endorse the same standards of improvement – some broadly coherentist standard, say – but start in sufficiently different places that, as each seeks coherence, their views don't converge. Second, Ned and Ted could endorse different standards of improvement, such that, even if their initial views were in other respects similar, they would diverge as each implemented revisions that were, by his own lights, improvements.

Ned and Ted disagree about whether P. But this is not the ordinary, garden-variety type of disagreement, which can be resolved by discussion, offering of reasons, doing further research, etc. This is a very dramatic sort of disagreement—call it

fundamental disagreement—in which neither party can be convinced to accept the other's view by any method that they would endorse.

Where there is disagreement, there is error. Where there is *fundamental* disagreement, there is fundamental error. If there's a fact of the matter about whether P, then either Ned or Ted has fallen into an especially bad sort of error, which isolates him from the truth in a particularly serious way: It's impossible for him to arrive at the truth about whether P by engaging in any process of belief revision that he'd endorse as legitimate. He can't get to the truth about whether P under his own doxastic steam.

So long as it's possible for people to start off with quite different views and/or different standards of improvement, this sort of divergence, leading to fundamental disagreement (and therefore fundamental error) will be possible. And of course, it *is* possible for people to start off with quite different views, and quite different standards of improvement.

Since the quasi-realist is committed to the possibility of people holding quite different moral views, and to the possibility of people having quite different standards of improvement, the quasi-realist is committed to the possibility of fundamental moral disagreement. It's possible for people to hold sufficiently different moral views that, even if both take the same sorts of changes to be improvements, no series of such changes would bring them to agree. It's also possible for people to endorse sufficiently different standards for improvement that no series of self-approved changes in either would bring him around to the other's view.

One of the major motivations for the move to quasi-realism is in order to allow for genuine moral disagreement. Where there is genuine disagreement, there must be

error. If our disagreement about whether P is genuine, then one or the other of us must be mistaken about whether P. Where there is *fundamental* disagreement, there must be *fundamental* error—error about one’s stable moral beliefs. If I stably believe that P and you stably believe that not-P, then one or the other of us is *fundamentally* mistaken—mistaken about one of our stable moral beliefs. Fundamental moral error is an especially serious sort of moral error. The agent who is fundamentally mistaken is isolated from the moral truth in a particularly tragic way—no change that they would recognize as an improvement will bring them around. Their own best, most sincere efforts at self-improvement are all doomed to failure; even at the end of their most heroic character-development projects, they will still be villains.

So, there’s an extremely bad sort of moral error—fundamental moral error—that the expressivist must admit is possible. In other words, the expressivist is committed to:

FUNDAMENTAL FALLIBILITY: It’s possible for people’s stable moral beliefs to be mistaken.

That is, it’s possible for people to be in the following bad situation: being such that some of their stable moral beliefs are mistaken.

But just a moment: does the quasi-realist really need to say that where there’s disagreement, there’s error? They do if they want to underwrite our ordinary practices of moral thought and debate. If the quasi-realist says that there can be moral disagreement without moral error, then they force a major revision of our ordinary ways of thinking about morality. Perhaps more importantly, they undermine their claim to have licensed

talk of moral truth: If it's possible to have moral disagreement without moral error, what's at stake in debates about morality can't really be *truth*. Quasi-realism is not supposed to commit us to relativism or dialetheism.

Blackburn seems to agree. Toward the end of *Ruling Passions*, he discusses the following objection raised by Judith Thomson: If *correctness* for moral beliefs is just a matter of some kind of *coherence*, then it could happen that Smith discovers that P, and Jones discovers that not-P. In that case, since 'discover' is factive, we'll be forced to say that P and not-P. And that would be very bad—it would show, again, that whatever is at stake in moral debates, it can't be *truth*.

Blackburn's response is to note that *I* will never have to say that P and not P, since it won't happen that *I* endorse both P and not P, or that *I would* endorse both P and not-P after suitable improvements (by *my* lights). The danger is that the quasi-realist will, because of facts about what coheres with *other* people's beliefs, be forced into endorsing contradictions. The response is that the quasi-realist is only forced to endorse the moral claims that cohere (in the right way) with her *own* moral beliefs.

Thomson's objection can be rephrased as follows: quasi-realists need to deny, in the case of disagreements about morality, that where there's disagreement there's error, and this shows that quasi-realists aren't really entitled to realistic-sounding, objective-sounding talk about moral truth after all. Blackburn's response is to show that a quasi-realist really *can* retain the disagreement-then-error principle. This seems to be the right response—it really would be a disaster for the project of underwriting full-blooded talk about moral truth if the quasi-realist had to admit the possibility of moral disagreement without error.

4. First-Person Immunity

For me to be fundamentally in error, I need to have some moral view that's (a) stable, and (b) mistaken. But given Blackburn's account of moral error, this can't happen. For my moral belief that P to be *stable* is for it to be such that it would survive any improving change (or course of improving changes). For my moral belief that P to be *mistaken* is for there to be some improving change (or course of improving changes) that would lead me to abandon P. So on Blackburn's account of moral error, a moral belief is mistaken only if it's not stable. So for me to be fundamentally in error, I'd need to have some moral view that was (a) stable, and (b) not stable, which I pretty clearly can't have.

So if I'm a reflective quasi-realist, I can know in advance, just by thinking about what moral error is, that *I* can't be fundamentally morally mistaken. And every reflective quasi-realist can go through the same reasoning to get their own first-person guarantee that they haven't fallen into fundamental moral error.¹⁰

So the quasi-realist is committed to:

FIRST-PERSON IMMUNITY: I have an *a priori* guarantee against fundamental moral error.¹¹

¹⁰ Of course, neither of us will recognize the other's guarantee. Whether or not *your* moral beliefs are mistaken (by my lights) depends not on whether or not they'd survive changes that *you* would endorse as improvements, but whether they'd survive changes that *I* would endorse as improvements. (Or perhaps whether they agree with my views, or the views I would have after a suitable course of improving changes, or something of this sort—in any event, the important thing is that, when I attribute moral error to you, it's *my* views, and improvements thereon, that your views are measured against, not your own. More on this in section 6.)

¹¹ Mark Johnston [1989] discusses a similar consequence of his dispositional theory of value. Though he takes the consequence to be a drawback of his account, he rightly does not take it to be as *serious* a problem there as it is in the case of quasi-realism.

Notice that the result here is *not* that I've got an *a priori* guarantee that any *particular* moral belief is correct. To get such guarantees about particular moral beliefs, I'd have to know that they were stable. I typically don't know that, and even if I did, it probably wouldn't be *a priori*. What I've got is a guarantee that *none* of my moral beliefs are *fundamentally* mistaken—that is, stable but incorrect. That is, I've got an *a priori* guarantee that I'm not isolated from the moral truth in such a way that I can't ever come to believe it by any process of revision that I'd endorse. What I know *a priori* is that I'm not in the following bad situation: being such that some of my stable moral beliefs are mistaken.

5. *This is Very Bad*

Now we have the ingredients for an argument against quasi-realism.

We have a very plausible principle, which Blackburn seems to endorse, about moral error:

NO SMUGNESS: There isn't any sort of moral error to which others are subject, but against which I have an *a priori* guarantee of immunity.

We also have two theses about fundamental moral error that follow from the Blackburnian account of moral error discussed in section 2:

FUNDAMENTAL FALLIBILITY: It's possible for people's stable moral beliefs to be mistaken.

FIRST-PERSON IMMUNITY: I have an *a priori* guarantee against fundamental moral error.

But we cannot have all three. FUNDAMENTAL FALLIBILITY and FIRST-PERSON IMMUNITY together tell me that there is a sort of moral error—namely *fundamental* moral error—of the sort that NO SMUGNESS says that there is not. There's a particularly bad epistemic situation—having some mistaken stable moral beliefs—that it's *possible* for people to be in, but which I know *a priori* is not *my* situation.

The fact that these three claims are inconsistent is bad news for the quasi-realist. FUNDAMENTAL FALLIBILITY and FIRST-PERSON IMMUNITY are consequences of quasi-realism plus the Blackburnian account of moral error sketched above. NO SMUGNESS captures an important part of our ordinary way of thinking about morality. The combination of quasi-realism, the Blackburnian account of moral error, and NO SMUGNESS is inconsistent. We must give up one of the three.

This is, once again, intended as a problem specifically for quasi-realism, *not* for expressivism in general. Not all expressivist views get in to this kind of trouble. One reason for this is that not all expressivist views are as ambitious as quasi-realism – not all expressivists are so committed to providing an account that is compatible with our ordinary ways of thinking and talking about morality as Blackburn. (Again, I think that Blackburn is *right* to be as committed to this as he is – if the quasi-realist project of squaring expressivism with ordinary moral practice were successful, that would be a

tremendous point in its favor. What I'm concerned to argue here is that the project, as ingenious and well-motivated as it is, turns out not to be successful.)¹²

For example, other expressivist views could deny FUNDAMENTAL FALLIBILITY, and say that *no one* is subject to fundamental moral error. An expressivist view less ambitious than quasi-realism, and more willing to give up on the 'objective pretensions' of our ordinary moral practice, could do this by embracing subjectivism or relativism. Such views would grant a perfectly *general* immunity to fundamental error, and thus avoid any violation of NO SMUGNESS. (Timmons [1998] contemplates this, and perhaps Gibbard [1990; 2003] embraces it.) I have provided no argument against this view. But this view is not (full-blooded) quasi-realism, because it gives up on the central quasi-realist project: underwriting our ordinary realist-seeming moral thought and practice.

Another option is to deny first-person immunity. This would require a new story about error – probably one that gives up on the idealization approach altogether. But it's hard to see what other approach there is to take, that would allow us to make room for moral error while still being antirealists.

¹² One might be concerned, not just about how general a problem this is for expressivist accounts, but to what extent this is also a problem for a certain sort of *realist* account that relies on similar kinds of idealizations of our existing moral outlooks. Note, though, that this sort of problem does not arise for an approach like Michael Smith's [1994], since Smith's theory includes a built-in convergence constraint. Since he's committed to convergence, Smith is not committed to FUNDAMENTAL FALLIBILITY. And so there is no objectionable asymmetry in our vulnerability to fundamental error.

There might be a similar problem for some constructivist approaches to mathematics or science – if I thought that mathematical or scientific truth was determined by what *I* would accept, in some ideal circumstances, then the same sorts of worries will arise for mathematical and scientific error. I take it, though, that there's typically a convergence constraint built in to mathematical or scientific constructivism, too – it's what *we* (not *I*) would accept, in some ideal circumstances. Perhaps there is still a lingering concern, both for Smith and for constructivists, about the possibility of other *communities* whose end-of-inquiry views diverge from ours. If there is, though, here is not the place to pursue it further. (Thanks to Michael Smith and Robert Nola for discussion.)

Finally, giving up NO SMUGNESS is unattractive. It really does seem extremely implausible that I have a special, *a priori* guarantee against some sort of (very serious) moral error to which you are vulnerable. It's also hard to make sense of a situation in which *everybody* has the same first-person guarantee against a sort of error to which *others* are vulnerable. At best, this is very, very strange. At worst, it is incoherent.

More importantly, whatever else is wrong with giving up NO SMUGNESS, denying it is in sharp conflict with our ordinary ways of thinking and talking about morality. Ordinary, realistic-looking moral views take all of us to be on (roughly) the same footing with respect to moral error – if *you're* potentially subject to a certain kind of moral error, then so am I. (Or at least, then I don't know *a priori* that I'm not.)¹³

Quasi-realism was supposed to provide us with a way of going on with our moral practice as before, but understanding it in a metaphysically innocent way. If quasi-realists are obliged to reject NO SMUGNESS, then they won't be able to just go on as before—being a quasi-realist will force a major revision of our ordinary ways of thinking and talking about moral matters.

This leaves us with two options: abandon quasi-realism, or provide a different quasi-realist account of moral error that avoids the commitment to one or both of FUNDAMENTAL FALLIBILITY and FIRST-PERSON IMMUNITY. If it turns out that we can't provide such an account, we'll be forced to give up quasi-realism. I'm skeptical about the prospects for a substantially different quasi-realist account of moral

¹³ A rhetorical question Michael Smith asked in conversation expresses the point nicely: "What's supposed to make me so [darn] special?" It's part of our ordinary moral practice that I'm *not* so [darn] special. Or at least, that if I am, I don't have an *a priori* guarantee of it, which I can find out about just by thinking about the nature of ethics. (If I am that special, it's because I'm *en rapport* with God, or the Forms, or something, and that's something I'm only going to find out *a posteriori*.)

error. Some story about stability under improving changes really does seem to be the best (perhaps the only) account of moral error that's available to quasi-realists.

6. *Another Worry*¹⁴

On the quasi-realist account of moral error sketched here, my moral beliefs are *mistaken* just in case some improving change would lead me to abandon them; the condition that my moral beliefs have to meet in order for them to be *correct* is that they be *stable*.

Quasi-realists may balk at this point. This account of correctness looks suspiciously like a specification of truth conditions of the sort that the quasi-realist will want to avoid. But the quasi-realist really is committed to this sort of genuinely truth-conditional looking thesis about what moral correctness consists in. Or at least, if they aren't, then they haven't really given an account of moral error that's satisfactory even for ordinary cases of non-fundamental error. There's something I'm concerned about when I'm concerned about my moral beliefs being mistaken. What is it? According to Blackburn, it's that they wouldn't survive some improving change – that is, that they're not stable. The bad situation that I'm in if my concerns are justified is that my moral belief is mistaken – that is, that it's unstable. So what it is for one of my moral beliefs to be mistaken is for it to be unstable. So what it is for one of my moral beliefs to be correct is for it not to be mistaken – that is, for it to be stable.

We can generalize: the condition that a moral belief (mine or someone else's) has to meet in order to be correct is that I *would* stably believe it, after some improving change.

¹⁴ Thanks to Michael Smith and Victoria McGeer for the conversation that led to this section.

This generalization has an important benefit: it allows quasi-realists to tell a unified story about error and correctness for my beliefs and yours. This is important because first- and third- person error interact in ways that seem to require a unified account. If you and I both believe that P, then it had better turn out that you're mistaken iff I am. If you believe that P and I believe that not-P, it had better turn out that you're mistaken iff I'm correct.

But now perhaps we've done too much: stable belief after some improving change has started to look an awful lot like *truth* for moral claims. There's a danger that, in providing an account of moral error, we've given up the expressivism at the core of the quasi-realist project. We seem to be entitled to the following two claims:

CORRECTNESS/TRUTH: The feature that a moral belief has to have in order to be correct is (moral) truth.

CORRECTNESS/IDEAL BELIEF: The feature that a moral belief has to have in order to be correct is: being such that I'd stably believe it after some improving change.

It follows from CORRECTNESS/TRUTH and CORRECTNESS/IDEAL BELIEF that what it is for a moral claim to be *true* is for it to be such that I'd stably believe it after some improving change; moral truth is stable belief by improved versions of me.

There are two problems with this outcome. First, once we've accepted this, we seem to be forced to give up either the expressivism or the deflationism that form the foundations of quasi-realism. The following triad is inconsistent:

EXPRESSIVISM: When I say that stealing is bad, I'm simply expressing my disapproval of stealing.

DEFLATIONISM: When I say that it's *true* that stealing is bad, I'm doing exactly the same thing as when I just say that stealing is bad.

TRUTH AS IDEAL BELIEF: When I say that it's true that stealing is bad, I'm saying that I would stably believe that stealing is bad after some improving change.

The source of trouble here is that EXPRESSIVISM and DEFLATIONISM are core quasi-realist doctrines, and it's hard to see how to run Blackburn's account of moral error as instability without incurring a commitment to TRUTH AS IDEAL BELIEF.

The second problem is that it now looks as if we're committed to a sort of subjectivism: when I call something right, I say that I'd stably approve of it after improving changes. When you call something right, you say that *you* would stably approve of it after improving changes. This is bad for all the reasons that subjectivism is bad. (For example, it looks like we're now talking past each other in moral disputes. When I say that stealing is wrong_{me}, and you say that it's right_{you}, we haven't succeeded in disagreeing with each other.) But it's worse than that. Whatever we think of the plausibility of subjectivism, accepting quasi-realism wasn't supposed to *commit* us to being subjectivists. Quasi-realism was supposed to be compatible with continuing to say all of the things that a die-hard moral realist says, and one of the things that die-hard moral realists say is that no version of subjectivism is true.

So generalizing Blackburn's account of moral error brings with it two more problems: it seems to force the quasi-realist to give up either his expressivism or his deflationism about truth, and it seems to commit quasi-realists to a sort of subjectivism. If we don't generalize the account, though, we're left without an account of the connections between first- person and third- person moral error, which is unacceptable.

Conclusion

A common first reaction to expressivist and quasi-realist theories is the thought that, if these theories are right, there's some objectionable sense in which we can't be wrong about morality. This worry turns out to be surprisingly difficult to make stick—an account of moral error as instability under improving changes provides the quasi-realist with the resources to explain many of our concerns about moral error. The story breaks down, though, in the case of *fundamental* moral error. This is where the initial worry finally sticks—quasi-realism tells me that I can't be *fundamentally* wrong about morality, though others can.

There is also a danger that providing a successful quasi-realist account of even *ordinary* moral error winds up undermining the quasi-realist project, by forcing the quasi-realist away from her original expressivism and deflationism, and obliging her to adopt a sort of subjectivism.

References:

- Ayer, A.J. (1952) *Language, Truth, and Logic*, Dover.
- Blackburn, Simon. (1984) *Spreading the Word*, Clarendon.
- (1993a) Realism, *Quasi* or Queasy?, in *Reality, Representation, and Projection*, J. Haldane and C. Wright, eds., Oxford University Press.
- (1993b) *Essays on Quasi-Realism*, Oxford University Press.
- (1998) *Ruling Passions*, Oxford University Press.
- Boyd, Richard. (1988) How to Be a Moral Realist, in *Essays on Moral Realism*, G. Sayre-McCord, ed., Cornell University Press.
- Brink, David. (1989) *Moral Realism and the Foundations of Ethics*, Cambridge University Press.
- Dreier, Jamie. (1996) Expressivist Embeddings and Minimal Truth, *Philosophical Studies* **83**: 29-51.
- Egan, Andy, Hawthorne, John and Weatherson, Brian. (2005) Epistemic Modals in Context, in *Contextualism in Philosophy*, G. Preyer and G. Peter, eds., Oxford University Press.
- Egan, Andy. (forthcoming) Epistemic Modals, Relativism, and Assertion, forthcoming in *MIT Working Papers in Linguistics and Philosophical Studies*.
- Geach, Peter. (1960) Ascriptivism, *Phil Review* **69(2)**:221-225.
- Geach, Peter. (1965) Assertion, *Phil Review* **74(4)**:449-465.
- Gibbard, Allan. (1990) *Wise Choices, Apt Feelings*, Harvard University Press
- Gibbard, Allan. (2003) *Thinking How to Live*, Harvard University Press
- Hale, Robert. (1986) The Compleat Projectivist, *Philosophical Quarterly* **36**: 65-84.
- Hale, Robert. (1993) Can there be a Logic of Attitudes?, in *Reality, Representation, and Projection*, J. Haldane and C. Wright, eds., Oxford University Press.
- Hare, R. M. (1970) Meaning and Speech Acts, *Philosophical Review* **79(1)**: 3-24.
- Johnston, Mark. (1989) Dispositional Theories of Value, *Proceedings of the Aristotelian Society*, Supplement **63**:139-174.
- Lewis, David. (1989) Dispositional Theories of Value, *Proceedings of the Aristotelian Society*, Supplement **63**: 113-137
- Mackie, John. (1977) *Ethics: Inventing Right and Wrong*, Penguin Books.
- Rosen, Gideon. (1998) Review of *Essays on Quasi-Realism*, *Nous* **32(3)**:386-405.
- Schueler, George. (1988) *Modus Ponens* and Moral Realism, *Ethics* **98**: 492-500.
- Searle, John. (1962) Meaning and Speech Acts, *Philosophical Review* **71(4)**:423-432.
- Smith, Michael. (1994) *The Moral Problem*, Blackwell.
- Sturgeon, Nicholas. (1988) Moral Explanations, in *Essays on Moral Realism*, G. Sayre-McCord, ed., Cornell University Press.
- Timmons, Mark. (1998) *Morality Without Foundations*, Oxford University Press.
- Unwin, Nick. (1999) *Quasi-Realism*, Negation, and the Frege-Geach Problem, *Philosophical Quarterly* **49**: 337-252.