

The Refined Moral Problem

Michael Smith

As Alexander Miller and Seth Whittington say in their "A Note on the Error Theory and the Refined Moral Problem" (2023), the title of *The Moral Problem* (1994) refers to the following three prima-facie inconsistent propositions that I present in the first chapter of the book:

- (1) Moral judgements of the form 'It is right that I Φ ' express a subject's beliefs about an objective matter of fact, a fact about what is right for her to do.
- (2) If someone judges that it is right to Φ then, ceteris paribus, she is motivated to Φ .
- (3) An agent is motivated to act in a certain way just in case she has an appropriate desire and means-end belief, where belief and desire are, in Hume's terms, "distinct existences" (Smith 1994: 12)

I claim that this prima-facie inconsistent triad of propositions is the central organizing problem in meta-ethics because various standard metaethical theories can be seen as different ways of responding to it. Traditional non-cognitivists like A. J. Ayer (1936) accept (2) and (3), and so reject (1); traditional externalists like Philippa Foot (1972) accept (1) and (3), and so reject (2); and anti-Humeans about motivation like John McDowell (1978) accept (1) and (2), and so reject (3).

Having said that, my own view, as Miller and Whittington also note, is that the inconsistency is merely prima-facie. There is a view in metaethics, moral rationalism, that comes in many versions, and while some versions can be represented as rejecting (3), but accepting (1) and (2), other versions accept (1), (2), and (3). However, explaining why this is so, and what the take-home lessons are for metaethicists, requires us to make and attend to a number of important distinctions. This is why the argument unfolds throughout the rest of the book and only becomes explicit in the final two chapters.

One of the important distinctions that needs to be made concerns talk of reasons for action. Such talk is, I suggest, systematically ambiguous between talk of motivating reasons—these are the psychological states that cause our actions—and normative reasons—these are the considerations that justify our actions.¹ The distinction between motivating and normative reasons for action is important because moral rationalism is a view about the relationship between the deontic statuses of actions—being impermissible, permissible, and obligatory—and normative reasons for action. According to moral rationalists, claims about actions being impermissible, permissible, and obligatory entail corresponding claims about normative reasons for action.² What moral rationalists should say about motivating reasons is one of the main issues addressed in the book. Somewhat surprisingly, it turns out that they should accept (3).

Another important distinction concerns moral rationalism itself. The doctrine can be understood in either relativistic or non-relativistic terms (Smith forthcoming). Consider Gilbert Harman's (1975) view that when you

¹ There has been a great deal of discussion about this distinction in the thirty years since *The Moral Problem* was published. Some have argued that I was wrong to think of motivating and normative reasons as belonging to different categories: psychological states versus considerations (see especially Dancy 2000). In their view, motivating reasons should also be thought of as considerations, and more specifically as those considerations we have in mind when we are moved to act. Sometimes when we act, they point out, the considerations we have in mind are the very considerations that in fact justify our actions, so in these cases our motivating reasons are our normative reasons. So far, so good. The trouble is that this isn't always the case. Sometimes we have false beliefs about the considerations that justify our actions, and we have these in mind when we act; sometimes—for example when we are feeling especially perverse—we have in mind considerations that *dijustify* our actions, to use the excellent term coined by Michael Stocker (1979), and so act in a way that is unjustified; and sometimes—for example when we are on autopilot (for me this is roughly the first hour after waking up!)—we have no considerations at all in mind when we act. It is therefore implausible to suppose that all reasons are considerations. Moreover, given that what makes a bodily movement an action is the fact that it is a movement of our body undertaken for reasons, and given that these reasons must be reasons of a kind that are present whenever we act, whether we have in mind certain considerations or not, and if we do, whether the considerations we have in mind justify our actions or not, these reasons cannot be considerations. Since the only things that are present whenever we act are those psychological states that cause our actions, it therefore seems best to reserve the term 'motivating reasons' for them. Actions are bodily movements undertaken for reasons, where these are just bodily movements caused by motivating reasons where motivating reasons are psychological states.

² In *The Moral Problem* I was thinking of being right and being the right thing to do as the same property, and that these are both the same property as being obligatory. However, I have since been persuaded that being right is better understood as being permissible. All sorts of complications for my statement of the moral problem ensue, but I will ignore these in what follows. I will pretend that the view I held when I wrote *The Moral Problem* is correct, and that an action's being right is the same as its being the right thing to do.

say of someone that they have moral obligation to act in a certain way, you are saying something that entails that acting in that way accords with the cooperative intentions that you share with them, or perhaps with the cooperative intentions that you are both committed to sharing given the desires each of you has. Given Harman's Humean view that someone has a normative reason for action only if they have some desire that will be satisfied by their acting on that normative reason, it follows that Harman's is a relativistic version of moral rationalism. To see why, ask yourself whether some agent has a moral obligation to act in a certain way in a particular situation, and hence a normative reason to act in that way. According to Harman, the answer depends on whether that person either shares, or is committed to sharing, cooperative intentions to act in that way in that situation *with you*. Does this mean that Harman's view is consistent with (1), (2), and (3)? It doesn't if, as I argue in *The Moral Problem*, the most plausible reading of (1) is the reading where 'objective' entails 'non-relativistic'. On that reading, Harman's view is inconsistent with (1).

Contrast Kant's (1785) view that the concept of a moral obligation is the concept of a categorical imperative, that is, the concept of a reason for action that people have simply in virtue of having the rational capacities they have as agents. This is a non-relativistic version of moral rationalism. To see why, ask yourself once again whether some agent has a moral obligation to act in a certain way in a particular situation. According to Kant, you can answer this question without knowing anything about that agent's cooperative intentions concerning that situation, or about any of the desires they have, or about any of the intentions or desires that you have. This is because the answer follows from the mere fact that the situation has the nature that it has, together with the fact that they are an agent. Agents in the same situations all have the same moral obligations, according to Kant. It is this non-relativistic version of moral rationalism that is front and center in *The Moral Problem*.

As Miller and Whittington point out, the view I argue for in the book is that the deontic concepts are all analysable in respondent-dependent terms; that the response-dependent analysis I give of the deontic concepts entails that the view I argue for is a version of non-relativistic moral rationalism; and that that version of non-relativistic moral rationalism entails (1), (2), and (3). To spell out these conclusions in a little more detail, moral judgements turn out to be expressions of beliefs about the normative reasons for action we have simply in virtue of being agents, and these in turn are representable as beliefs about the desires we would all converge on having if we were fully rational, which are beliefs about a non-relative matter of fact (thus (1)); there is a wide-scope rational requirement connecting such beliefs with corre-

sponding desires (thus (2)); and an agent with such beliefs is only motivated when they conform to that wide-scope rational requirement by having a corresponding desire that hooks up with some distinct means-end belief (thus (3)). The inconsistency of (1), (2), and (3) is thus merely prima-facie. Mine is a version of moral rationalism that accepts all three claims.

Note, however, that the argument that I give for these conclusions is based entirely on the nature of the concepts in play. The upshot is that even if the arguments are successful, it doesn't follow that any such normative reasons for action exist. Showing that would require showing that the concepts in play are instantiated. In other words, it would require demonstrating that we would indeed all converge on certain desires if we were fully rational. In the final chapter of *The Moral Problem*, I address this concern and close on an optimistic note. I say that we can proceed on the assumption that the concept of the desires we would all converge on if we were fully rational is instantiated, and that we can profitably spend our time looking for the arguments that show this to be so.³ But I also acknowledge that if no such demonstration is to be had, then we will have to conclude that the concept of the desires we would all converge on if we were fully rational is not instantiated. Our talk about actions being impermissible, permissible, and obligatory would then be based on a false presupposition of convergence. This, I say, is where the error theory is best located. The error theorist and I agree that our deontic concepts are analysable in terms of the concept of a non-relative normative reason for action, and this suffices for both of us to reconcile (1), (2), and (3). We still disagree, however, because I say that this concept is instantiated, whereas the error theorist says that it isn't.

Let's now turn to Miller and Whittington's complaints. Without taking a stand on whether there is a central organizing problem in metaethics, Miller and Whittington argue that the refined moral problem they propose is a better candidate for being that central organizing problem, if there is one, than the three prima-facie inconsistent propositions I propose. The trouble with my proposal is that:

...it is unclear how the topography provided by Smith's version of "The Moral Problem" can accommodate error theories of the sort most famously advocated by J.L. Mackie (Mackie 1977), and we trace this limitation to a failure to distinguish between two distinct types of what Smith terms "Moral Nihilism". (Miller and Whittington 2023: 110)

³ My most recent attempt to provide such an argument appears in Smith 2020.

The Refined Moral Problem

Given what I've just said about how I accommodate Mackie's error theory, I must confess I was surprised when I read this. Here is Miller and Whittington's explanation of why my attempted accommodation is inadequate.

Smith himself describes the error theory...as involving “moral nihilism” (1994: 11), and he also characterizes moral nihilism as involving the idea that no proposition in the trio of (1), (2) and (3) can be justifiably rejected to solve the “Moral Problem” (1994: 13). Perhaps, then, the error theory does appear in Smith’s cartography: it’s the view that the “Moral Problem” can’t in fact be solved by justifiably rejecting one of the propositions which gives rise to it. This won't do, however, as the “moral nihilism” that concedes that the “Moral Problem” can’t be solved would amount to the view that “the very idea of morality [is] altogether incoherent” (1994: 5), incoherent in the sense of containing a conceptual contradiction or irresolvable tension of the sort apparently generated by the conjunction of (1), (2) and (3). This is not a good fit for Mackie’s error theory: Mackie’s view, after all, is not the view that moral practice is in some deep sense incoherent, but rather the view that although moral practice is coherent, nothing in fact corresponds to it in reality. If moral practice were “altogether incoherent” as per Smith’s characterization of moral nihilism, it would presumably be impossible to be anything other than an *eliminativist* concerning it. And as we saw above, Mackie’s error theorist is not an eliminativist. One way of framing this point would be to characterize Mackie’s error theory as a form of *weak moral nihilism* and the view that the Moral Problem cannot be solved as a form of *strong moral nihilism*. The criticism of Smith would then be that he conflates the strong and weak forms of moral nihilism. (Miller and Whittington 2023: 115)

This requires some unpacking.

Mackie's (1977) remarks about the objectively prescriptive nature of moral properties, and his reasons for thinking no such properties exist, are murky at best. His initial characterization of them as properties whose recognition would bring motivation with it makes it sound like he thinks their existence would require an anti-Humean psychology. For moral properties to exist, he seems to be thinking, (3) would have to be false, but (3) is true. However, his later suggestion that his talk about the objective prescriptivity of moral properties, and Kant's talk about moral judgements expressing

categorical imperatives, both amount to the same thing suggests a quite different objection. To believe that someone is under a moral obligation to act in a certain way in a certain situation is, on this way of thinking about the objective prescriptivity of moral properties, to believe that they have a normative reason to act in that way in that situation that isn't grounded in any intention or desire they have in that situation, but is instead grounded in the fact that the situation has the features that it has and the agent in question is a rational agent. On this way of thinking about objective prescriptivity, I agree with Mackie that moral properties are objectively prescriptive. We agree that moral judgements are beliefs about categorical normative reasons for action. We only part company when he goes on to argue that there are no such normative reasons.

The problem with thinking that there are categorical normative reasons, as Mackie sees things, is not that the actual world is missing some contingent feature, and that in other possible worlds where that contingency is true there are such normative reasons. The problem is rather that there is no compelling argument available in any possible world that would convince the arbitrary rational agent in that situation to acquire a desire to act in the relevant way if they didn't have such a desire already. His view is therefore that it is a necessary truth that there are no categorical normative reasons, and that that necessary truth is a priori knowable. Contrary to what Miller and Whittington say, Mackie is therefore committed to talk of moral obligation being incoherent in the sense of being inconsistent. But the inconsistency is not the inconsistency of (1), (2), and (3). The inconsistency is that between the proposition that gives the content of the belief in (1)—the proposition that we would all converge on a desire to act in a certain way if we were fully rational—and something that, at least according to Mackie, is knowable by anyone who has that belief on the basis of reflection alone, namely, that that belief is false.

So what should we make of Miller and Whittington's suggestion that "Mackie's view...is not the view that moral practice is in some deep sense incoherent, but rather the view that although moral practice is coherent, nothing in fact corresponds to it in reality" (Miller and Whittington 2023: 115)? Here is what Miller and Whittington say in support of this interpretation of Mackie.

Mackie is not an eliminativist about moral judgement: although moral judgements are systematically false, some moral judgements are justified in the sense that their adoption assists groups of humans to

garner the benefits of social co-operation (1977: Chapter 5). (Miller and Whittington 2023: 114)

When Miller and Whittington say that Mackie's view is that moral practice is coherent, it turns out that all they mean is that he thinks that it is justified in the sense of being beneficial to engage in. But note that Mackie's thinking this is consistent with his also thinking that talk of moral obligation commits the one doing the talking to inconsistency. The upshot is that while Miller and Whittington are right that Mackie is not an eliminativist about moral obligation talk, they are wrong that this tells against supposing that the error made by those who engage in moral practice is the error of inconsistency.

Does anything that's just been said about Mackie's error theory show that my own characterization of the error theory in *The Moral Problem* is problematic? Are Miller and Whittington right that my characterization is one according to which the error theory is eliminativist? The answer is no to both questions. Since I take no stand on whether it would be beneficial to continue to talk about moral obligations, even if we were to become convinced that all such talk commits us to inconsistency, my characterization is neutral on whether an error theorist should be a strong or a weak moral nihilist. Miller and Whittington's claim that I conflate the strong and weak forms of moral nihilism is therefore plainly false. As I see things, strong and weak moral nihilists agree that talk of moral obligation commits the one doing the taking to inconsistency, but disagree about whether a cost-benefit analysis of continuing to talk in that way, notwithstanding the inconsistency, comes down on the side of continuing to talk in that way. Weak moral nihilists like Mackie say that it does, strong moral nihilists say that it doesn't. So while strong moral nihilists are eliminativists, since I take no stand on the issue that divides them from weak moral nihilists, my own characterization of the error theory is silent on the issue of elimination.

Up until this point I have said nothing about the following Refined Moral Problem which, Miller and Whittington argue, is a better candidate for being the central organizing problem in metaethics than the three prima-facie inconsistent propositions I present in the first chapter of *The Moral Problem*.

- (1_{RMP}) There is at least one moral judgement of the form “It is right that I Φ ” such that this judgement
- a. Expresses a subject’s belief about an objective matter of fact, a fact about what is right for her to do.
 - b. Is true.
 - c. Is justified.

- (2_{RMP}) It is a conceptual truth that: if someone knows that it is right to Φ and is practically rational then, *ceteris paribus*, she is motivated to Φ .
- (3) An agent is motivated to act in a certain way just in case she has an appropriate desire and means-end belief, where belief and desire are, in Hume's terms, "distinct existences".

Note that the Refined Moral Problem is also supposed to be a *prima-facie* inconsistent trio of propositions, and that only the third is the same as in my statement of the problem (Miller and Whittington 2023: 116). The question is whether the Refined Moral Problem, which replaces (1) with (1_{RMP}) and (2) with (2_{RMP}), is an improvement on my suggestion.

As Miller and Whittington point out, one attraction of (1_{RMP}) is that it allows us to easily locate the difference between the error theorist, the non-cognitivist, and the realist. The non-cognitivist rejects (a), and because he rejects (a) he also rejects (b) and (c). The error theorist accepts (a), but he rejects (b), and he may or may not accept (c). This will depend on whether he thinks that someone's failing to engage in the *a priori* reasoning that leads to the error theory suffices for any moral beliefs they have to be unjustified. The realist accepts all three parts of (1_{RMP}). (1_{RMP}) has another attractive feature as well. It prompts us to ask whether there are positions in metaethics beyond those we've been talking about. For example, is there a view in metaethics according to which (a) and (b) are true, but (c) is false? This would be a realist view, but one according to which we have no reliable methods for getting moral knowledge. Interestingly, Sharon Street seems to think that all robust realists are committed to a view like this by what she calls the Darwinian Dilemma (Street 2006)—the view defended by Thomas M. Scanlon in the first chapter of his *What We Owe To Each Other* seems to have been her main target.

These attractions of the Refined Moral Problem suggest that it teaches us more about the topography of metaethics than the three *prima-facie* inconsistent propositions that I present in the first chapter of *The Moral Problem*. But the Refined Moral Problem also has some negative features. I said earlier that non-cognitivists accept (2) and (3), and so reject (1); that externalists accept (1) and (3), and so reject (2); and that anti-Humeans about motivation accept (1) and (2), and so reject (3). (1), (2), and (3), and their *prima-facie* inconsistency, thus play an explanatory role in metaethics. They capture metaethicists' reasons for embracing the metaethical theories they embrace. The question is whether (1_{RMP}), (2_{RMP}), and (3), together with any

prima-facie inconsistency between them or their parts, similarly explains why metaethicists embrace the metaethical theories they embrace.

Focus on non-cognitivism and (1_{RMP}). Does the Refined Moral Problem explain why traditional non-cognitivists reject (a), and hence reject (b) and (c) as well. The answer is that it does not. Traditional non-cognitivists reject (a) because they are judgement internalists and Humeans: that is, because they accept (2) and (3). They do not reject (a) because they accept (2_{RMP}) and (3). Indeed, it isn't clear that traditional non-cognitivists could think that a putative conceptual connection between moral knowledge and motivation of the kind claimed in (2_{RMP}) could explain anything at all, given that they don't think there are any moral beliefs, and hence no possible worlds in which there is any moral knowledge either. What's lost in Miller and Whittington's abandonment of (2) in favour of (2_{RMP}) is thus the crucial role played by judgement internalism in motivating non-cognitivism in the history of metaethics.

All things considered, it therefore seems to me that my original presentation of the moral problem in *The Moral Problem* is a better candidate for being the central organizing problem in metaethics than Miller and Whittington's Refined Moral Problem. But having said that, we should be grateful to them for drawing our attention to the ways in which different metaethical theorists respond to the different parts of (1_{RMP}), and for asking us what they might have to say about (2_{RMP}). We are likely to add significantly to our stock of metaethical knowledge by following up on their suggestions.⁴

References

- Ayer, A.J. 1936. *Language, Truth and Logic*. London: Gollancz.
- Dancy, Jonathan 2000. *Practical Reality*. Oxford: Oxford University Press.
- Enoch, David 2011. *Taking Morality Seriously*. Oxford: Oxford University Press.
- Foot, Philippa 1972. "Morality as a System of Hypothetical Imperatives" reprinted in Philippa Foot, *Virtues and Vices*. Berkeley: University of California Press, 1978. 157–73.
- Harman, Gilbert 1975. "Moral Relativism Defended" in *The Philosophical Review*, Vol. 84, No.1, 3–22.
- Kant, Immanuel 1785. *Groundwork for the Metaphysics of Morals*. Oxford: Oxford University Press, 2019.
- Mackie, J.L. 1977. *Ethics: Inventing Right and Wrong*. New York: Penguin.

⁴ Thanks to an anonymous referee for helpful comments on a previous draft.

Michael Smith

- McDowell, John 1978. "Are Moral Requirements Hypothetical Imperatives?", *Proceedings of the Aristotelian Society Supplementary Volume*. 13–29.
- Miller, Alexander and Seth Whittington 2023. "A Note on Error Theory and the Refined Moral Problem" in *Filosofiska Notiser*, Årgång 10, Nr. 1, 110–120.
- Smith, Michael 1994. *The Moral Problem*. Oxford: Wiley Blackwell.
- _____. 2020: "The Modal Conception of Ideal Rational Agents: Objectively Ideal Not Merely Subjectively Ideal, Advisors not Exemplars, Agentially Concerned Not Agentially Indifferent, Social Not Solitary, Self-and-Other Regarding Not Wholly Self-Regarding" in *Explorations in Ethics*. David Kaspar (ed). New York: Palgrave Macmillan. 59–79.
- _____. forthcoming: "The Bearing of Moral Rationalism on Moral Realism" in *Oxford Handbook of Moral Realism*. Paul Bloomfield and David Copp (eds). New York: Oxford University Press. 84–105.
- Stocker, Michael 1979. "Desiring the Bad: An Essay in Moral Psychology" in *Journal of Philosophy* Vol. 76, No. 12, 738–53.
- Street, Sharon 2006. "A Darwinian Dilemma for Realist Theories of Value" in *Philosophical Studies*, Vol. 127, No. 1, 109–166.

Michael Smith
Princeton University
msmith@princeton.edu