

Moral Realism and Arbitrariness

(published in *Southern Journal of Philosophy* 43 (1):109-129 (2005);
<http://www.blackwell-synergy.com/doi/abs/10.1111/j.2041-6962.2005.tb01946.x>)

Moral realism is extremely attractive. It seems to capture several of our important moral intuitions. It provides for the cognitivism which seems implicit in most commonsense moral discourse. It allows for genuine moral disagreement. Perhaps its most attractive feature is that it appears to offer us a way for us to say that (e.g.) the Nazis were objectively wrong - not just from our culture's perspective, or because of typical human emotional responses, but *truly* wrong. They are out of accord with the facts of the world. We can show them that they are wrong, and that if they came to understand the universe properly they would see that their actions are immoral and, further, that this gives them reason to stop performing them. They would be just as mistaken as if they declared that $2+2=5$, that water is HCL, or that whales have gills.

In this paper I argue that the appeal of moral realism, particularly with respect to this last virtue, is largely illusory. I focus on moral realism, but intend the arguments raised to apply equally to other normative realisms. I argue (i) that choosing to abide by such normative facts would be as arbitrary as choosing to abide by the mere preferences of a God (a difficulty akin to the Euthyphro dilemma raised for divine command theorists); in both cases we would lack reason to prefer these standards to alternative codes of conduct. I further develop this general line of thought by arguing in particular (ii) that we would lack any non-circular justification to concern ourselves with any such realist normative standards.

What is moral realism? This, of course, is something of a vexed question. I wish to work with a rather broad, intuitive notion. I take moral realism to involve primarily a metaphysical claim, that there exist certain mind-independent moral facts or properties, and that these

determine the truth values of our moral claims.¹ The moral status of an action (state of affairs, character trait, etc.) is not fundamentally determined by the attitudes of any agents (be they fully-informed humans, ideal contractors, God, or what-have-you).² Russ Shafer-Landau usefully characterizes moral realism as follows:

[T]he moral standards that fix moral facts are not made true by virtue of their ratification from within any given actual or hypothetical perspective.³

Of course the attitudes of living beings will often be relevant to the moral status of an action (or state of affairs, etc.) in other ways. For example, consider a moral realist form of act-utilitarianism. If a certain action will produce a great deal more happiness than any other alternative, this will obviously be morally relevant, and the action morally right. But note that the action is right according to such a realist theory because it maximizes happiness; not because people in our culture have pro-attitudes towards actions that maximize happiness, or because God prefers actions that maximize happiness, or what-have-you.

¹ For classic examples of American ('Cornell') moral realism, see Richard Boyd, "How to Be a Moral Realist," in Geoffrey Sayre-McCord (ed.), *Essays on Moral Realism* (Ithaca: Cornell University Press, 1988), 181-228, David O. Brink, *Moral Realism and the Foundation of Ethics* (Cambridge: Cambridge University Press, 1989), and Peter Railton, "Moral Realism," *Philosophical Review* 95 (1986), 163-207. For British moral realism, see Jonathan Dancy, *Moral Reasons* (Oxford: Blackwell, 1993), John McDowell, "Values and Secondary Qualities," in Sayre-McCord (1988), 166-80, and David Wiggins, "Truth, Invention, and the Meaning of Life," in Sayre-McCord (1988), 127-65.

² The British realists do typically emphasize the reactions of agents to moral facts. Still, on such views these reactions themselves must be appropriate (in some sense) to the moral facts, where this is not simply a matter of agent-attitudes. Thus on such theories the moral facts are not fundamentally grounded in the attitudes of agents, and the theories qualify as varieties of moral realism.

³ Russ Shafer-Landau, *Moral Realism: A Defence* (Oxford: Clarendon Press, 2003), 15.

I - The Euthyphro Dilemma and Moral Realism

The standard Euthyphro dilemma is often thought to show a fatal flaw in divine command theories of ethics.⁴ Peter Singer presents the problem as follows:

If all values result from God's will, what reason could God have for willing what he does? If killing is wrong *only* because God said: "Thou shalt not kill," God might just as easily have said: "Thou shalt kill." Would killing then have been right? To agree that it would have been right makes morality too arbitrary; but to deny that it would have been right is to assume that there are standards of right and wrong independent of God's will.⁵

Either God simply arbitrarily chooses certain actions to be good, evil, etc., in which case it is not at all clear why we should abide by the whims of some being (even if extremely powerful) – except perhaps out of mere prudence, or God appeals to some other (independent, objective) standard for judging actions, in which case it is not God's commands which give an action (or what-have-you) its moral status.

Consider again the first horn of the dilemma, the arbitrariness of God's commands. If we can object to such arbitrariness, then we can object to realist moral facts. Real, objective moral facts can be seen to be as 'arbitrary' (in the sense of not being chosen according to a further, independent standard) as God's commands. Better put: the realist moral facts would be equally unexplained. No one chooses the moral facts; they simply appear as part of the world. Just as a God would have no reason for choosing norms, the universe would have no reason for 'choosing' (as it were) the realist norms. Shafer-Landau, in defending moral realism, writes,

⁴ Plato, *Euthyphro*, 9e-11b.

⁵ Peter Singer, *The Expanding Circle: Ethics and Sociobiology* (New York: Farrar, Strauss & Giroux, 1981), x.

If we assume that at least some laws are best construed realistically, then the truth of these standards (e.g. those of mathematics and chemistry) is not to be explained by their having been accepted or created by anyone. *They just are correct.* Now if mathematical or physical laws do not always require lawmakers, then perhaps moral laws do not require lawmakers.⁶

Realists have no explanation of what makes their favoured property identities or supervenience relations true. [...] Realists will claim that there is no intelligible and plausible answer to such a question. Indeed, to insist that the fundamental laws or principles of ethics require something to *make them true* just presupposes the falsity of moral realism.⁷

In both cases, then, we end up positing brute, base-level facts. God simply prefers some rules or actions over others, and there is no further justification for why he so chooses; or it simply is the case that various moral rules and standards are true and correct, and there is no further justification or explanation for why this is so. If this is a problem for one position, it seems it would also be a problem for the other; the realist is just as susceptible as the divine command theorist here. That is, if it is a flaw in divine command theories to ultimately rest on a God's unexplained, brute preferences, then it seems similarly a flaw in moral realism to ultimately rest on unexplained, brute supervenience relations (etc.). On the other hand, if it is acceptable for the moral realist to rest with such unexplained, brute moral facts, then it would also seem acceptable for the divine command theorist to rest with a God's unexplained preferences.⁸ As an initial

⁶ Shafer-Landau, *Moral Realism*, 46 (emphasis added).

⁷ Ibid., 96-7.

⁸ There is a potential disanalogy here. In the case of the divine command theory, there is a lack of *justification* (in the sense of a reasoned-basis) for God's preferences, while it appears that the moral realist simply lacks an

point then, note that versions of the *Euthyphro* dilemma would seem to apply equally to moral realist and divine command theories insofar as both are committed to unexplained, brute moral posits. The two will sink or float together on this count.⁹

In itself, the unexplained (or, indeed, inexplicable) basis of God's preferences or the realist moral facts may not be terribly worrying; after all, explanations must come to an end at some point. We won't enter into the issue here. But more troubling is a deeper problem that can be seen as emerging from this initial discussion: our choosing to act either on God's commands or the moral facts that happen to obtain in the world would be arbitrary, at least without further arguments. We have yet to be given any reason to think that God's commands or the moral facts that obtain in this world are in any sense superior or more choiceworthy than any rival ethical

explanation of why the supervenience relations, etc. that obtain do so. However, notice that the realist would be equally unable to provide a justification of the sort lacked by the divine command theorist. On the realist view there is no agent consciously choosing the fundamental moral supervenience relations, etc.; as such there is no justification for why precisely these fundamental moral facts and rules obtain, and this is no better than the position of the divine command theorist. In any event, we must ask why the divine command theorist would be required to provide a *justification* of the moral facts, if this is not also being demanded of the moral realist. Without an explanation it seems we are illegitimately placing a heavier burden upon the divine theorist – or an unduly light burden on the realist. (Perhaps this heavier burden would be appropriate if we hold divine command theories to be normative, rather than metaethical theories [as metaethical theories might seek only to explain, not justify morality]. Still, I take it that most divine command theorists in fact hold the theory as a metaethical position). Thanks to Josh Glasgow for discussion – see the weblog PEA Soup, <http://peasoup.typepad.com/peasoup>, entry on July 20, 2004.

⁹ Paul Copan, "Can Michael Martin Be a Moral Realist? *Sic et Non*," *Philosophia Christi* 2 (1999), 45-72 draws attention to a similar worry for (non-theistic) moral realists, though in a rather undeveloped form; see 61-3. Copan does not consider any plausible realist responses to the proposed problem, and still embraces a modified divine command theory. Note that the current paper is not a defense of a divine command theory; rather, it is intended to show (in part) that moral realism is subject to worries akin to those raised against divine command theories.

theories. In both cases (divine command theory or moral realism) we arrive at brute moral posits (God's preferences or the moral laws that obtain in the world) with no explanation – or justification – for taking these to be superior to any alternative codes of conduct available to us. We could choose to act in accordance with any number of broadly ethical systems; why restrict ourselves to that system embodied in the realist's moral facts?

For example, suppose it were to turn out that the actual realist morality is a form of utilitarianism. Still, it seems there would be possible worlds with broadly deontological moral facts, or worlds in which the realist moral facts would be virtue-based (with correspondingly different supervenience relationships between moral and base properties). The moral rules obtaining in the actual world are in no obvious way superior to these possible rivals. They are merely the moral rules which happen to obtain here, and we are given no reason to prefer this local morality. We thus arrive at the fundamental arbitrariness problem for moral realism. We need some explanation of why moral facts that simply (and without explanation) obtain in a world are to be seen as worthy of following rather than any rival ethical systems (particularly if the equally unexplained commands of a God are not). Without this, our choosing to act in accordance with the realist moral facts, rather than their rivals, is arbitrary.

Some might not accept that morality could vary across worlds. Again - this variation seems quite possible, just as it seems that the laws of physics could be quite different in other possible worlds. But let us suppose that morality involves a form of, e.g., utilitarianism across all possible worlds. We will not allow the possibility that morality radically changes across worlds. Now the question simply shifts to the following: why favour morality over a Kantian morality*, or a virtue-based morality**? Even if the range of the moral is significantly limited,

we still have competitors for guiding our actions, and we still have no reason to prefer morality. Arbitrariness again enters – why choose to follow the moral facts and not the moral* facts?

An exasperated moral realist will, of course, object: ‘We cannot provide an explanation of why the laws of physics that obtain in the actual world do so, any more than we can provide such an explanation for why the realist moral laws that obtain in the actual world do so. But we obviously don’t complain about obeying the laws of physics! These are the laws of physics that obtain in this world; end of story. We might not be able to explain why these laws obtain in this world, but the fact remains that they do so obtain. And the physical facts that obtain in the world are genuine facts (if you don’t believe this then try testing them by walking into the path of a speeding bus!). Similarly then, with realist moral facts; we might not be able to explain why they obtain in this world, but they do. They are still genuine facts, and we have reason to act in accordance with them, just as we have reason to act in accordance with the physical facts.’

This objection has a strong intuitive pull, but it fails insofar as it rests on a significant disanalogy. With respect to physical laws in the actual world, we have no choice in following their dictates. That is, we necessarily act according to the laws of physics (assuming some form of physicalism). To obtain a proper analogy, we would need to have rival systems of physics that we could choose to follow instead of the physics that obtains in this world. For example, imagine that just by our choice, we could instead live under physics*, with slightly weaker gravity than that which we find in actual physics, or under physics** with slightly stronger gravity. Or, more grandly, what if we could choose to live under cartoon physics? We could walk off of cliffs without falling, so long as we did not look down to the ground below, and so on. In the case of moral realism, we have the choice of acting in accordance with the realist moral rules or choosing an alternative. Now if we truly had a range of possible physics available

to us, the arbitrariness worry would also apply. Why choose 'physics' instead of cartoon physics? What would make actual physics superior to physics*, such that we should choose to concern ourselves with it over physics* (or physics**, or any other rival broadly physical system) – given that we have a genuine choice? Under such conditions, it seems that to blindly choose physics without further reason is an arbitrary choice on our part. And so it is actually with moral realism, given the alternatives actually available to us.

Compare the position of the moral realist to that of the cultural relativist. Imagine a person in a given society wishes to be a good person, in a broadly moral sense. However, she notices that there are other societies with very different moral rules from those which obtain in her society. If she questions why she (and her society) should follow the rules which currently obtain in her specific society, most of us will not find the answer 'Well – these are the *moral* rules which obtain *here*, so of course you should follow them' to be satisfactory. This person has noticed a certain arbitrariness which cannot be dismissed so easily. Now imagine moral realism is correct, but a person recognizes that there are possible worlds with very different moral rules, and that it is open to her (and her society) to choose to act in accordance with these other rules instead. Why would following these actual moral facts be any less arbitrary than following those simply found in some culture, given that we recognize that we have a wide range of options genuinely available to us?

If the moral realist then insists that *moral* laws could not vary in this way across possible worlds, we still arrive at a problem. Compare the cultural relativist who says 'True – there are different codes of conduct in other cultures, but they aren't really *morality*. Morality is the precise system we have in our culture.' Fine, but the natural question to ask is: why should we prefer this *morality* over the other codes of conduct available to us? Similarly, even if the moral

rules do not change across possible worlds, why should we choose to act in accordance with these moral rules, as opposed to moral* rules, etc.?

Some might worry that the present discussion has assumed a certain realism about possible worlds, and that without such an assumption the arguments may lose their force. But even if we put aside the entire issue of possible worlds, the basic problem remains. Terry Horgan and Mark Timmons have argued against certain forms of moral realism by appealing to a scenario involving a Moral Twin Earth. As they write,

One might suppose that there is an important disanalogy between Putnam's Twin Earth scenario about water and our Moral Twin Earth scenario: on earth, water is composed of H_2O , whereas on Twin Earth, what they call 'water' is composed of XYZ; but in the moral case both consequentialist and deontological properties are 'out there'. [...] On both earth and Moral Twin Earth there are too many natural properties available to serve as referents of moral terms like 'good' and 'right' – specifically, the functional property definable as the 'good'-role property by a consequentialist theory T^C , and the functional property definable as the 'good'-role property by a deontological theory T^D . [...] It looks like appeal to referential intentions that the two groups share, involving notions like flourishing and impartiality, is too weak to pin down the referents of moral terms; moral indeterminacy results.¹⁰

¹⁰ Terry Horgan and Mark Timmons, "Copping Out on Twin Earth," *Synthese* 124 (2000), 145. See also Terry Horgan and Mark Timmons, "New Wave Moral Realism Meets Moral Twin Earth," *Journal of Philosophical Research* 16 (1991), 447-65; "Troubles for New Wave Moral Semantics: The Open Question Argument Revived," *Philosophical Papers* 21 (1992), 153-75; "Troubles on Moral Twin Earth: Moral Queerness Revived," *Synthese* 92 (1992), 221-260; and "From Moral Realism to Moral Relativism in One Easy Step," *Critica* 28 (1996), 3-39.

Horgan and Timmons are here concerned with the (in)determinacy of reference of moral terms. What I wish to draw attention to is that even on earth alone, we are faced with a variety of possible broadly moral theories, including a wide range of consequentialist and deontological theories (among others). There are consequentialist properties out there, available to us, but there are also deontological properties (out there, available to us). Putting aside issues of reference, we can still ask ‘why should we choose to act in accordance with any one of these theories and its facts?’. In particular, even if some determinate theory is picked out by some reference-building relation as ‘morality,’ why should we choose to follow what is dictated by this theory, rather than choosing to follow the dictates of a slightly different theory (perhaps requiring a slight change in ordinary language use)? Without further information, the choice – as a choice of moral system - seems arbitrary. There might be good reasons, from the point of view of the philosophy of language, to prefer one system over the other. But this hardly seems to justify one system over the other, *qua* moral system.

Further Objections and Responses

A realist could claim that we have been assuming that moral realism cannot allow for any form of vagueness, and that this in turn leads us to believe we can plausibly choose some other (broadly moral) system. For example, Horgan and Timmons attempt to force moral naturalists (like Boyd, Brink, and Railton) into a dilemma, the first horn of which is:

that the putatively reference-fixing relation R might fail to fix *determinate* reference-relations between moral terms and certain natural properties that satisfy the constraints of R . For instance, perhaps the R -constraints are satisfied by a class of natural properties –

Finally, Timmons also draws upon the Moral Twin Earth scenario in his *Morality Without Foundations* (New York: Oxford University Press, 1999).

functional properties, say – that collectively satisfy some consequentialist moral theory T^C , and yet the R -constraints are *also* satisfied by another class of natural properties – also functional properties, say – that collectively satisfy some deontological moral theory T^D .¹¹

Thus, Horgan and Timmons are taking such indeterminacy of reference to pose a significant difficulty for such theories. But if a moral realist can allow for indeterminacy and vagueness, she can simply embrace the result that there are several moral theories that could satisfy the general constraints of morality (and the reference fixing relation, R). I will not enter here into the issue of whether the particular theories of Boyd, Brink, and Railton could allow for such indeterminacy.

Let us simply assume that there are forms of realism that *can* allow that morality is vague, that there are some determinate moral facts, and that there are some indeterminate cases; Shafer-Landau and Nobis have both argued for this claim.¹² Realists could thus allow both that it might be unclear (epistemically) whether lying in a given case is morally permissible; and that it could be that there is not a determinate right action in the given case (metaphysical vagueness). In this way moral realism could account for those cases where it seems that many alternative actions might be right (and upon which our discussion has relied) – via recognition both of the vagueness of moral properties and concepts, and our limited epistemic abilities.

But even if we allow that the realist can accommodate various forms of vagueness the problem remains. We now have the determinate realist moral facts, and a particular range of

¹¹ Horgan and Timmons, “Copping Out on Twin Earth,” 140.

¹² See Russ Shafer-Landau, “Vagueness, Borderline Cases and Moral Realism,” *American Philosophical Quarterly* 32 (1995), 83-96, and Nathan Nobis, “Vagueness, Borderline Cases and Moral Realism: Where’s the Incompatibility!?,” *Philosophical Writings* 14 (2000), 29-39.

vague cases. But why accept this system, and not morality* which is similar, but with slightly different dictates in several cases, and with a somewhat different range of vague cases? Why concern ourselves with the properties picked out by the reference-fixing relation R rather than the reference-fixing relation R^* ? Consider the virtue of generosity. Suppose that morality requires a certain amount of giving on the part of a person (relative to her means), a certain frequency of giving, and a non-grudging attitude in giving. This will presumably lead to vague cases, as we find people who give less frequently, but give more when they do, or people who give significant amounts frequently, but with a slightly flawed attitude, and so on. Now why prefer morality, and its balancing of elements for the virtue of generosity, instead of a morality*, which might weigh the frequency of giving somewhat more (for generosity*) than morality, while downplaying the amount given?

What makes being generous better than being generous*? Of course it might be *morally* better to be generous; but it will be *morally** better to be generous*. We still lack any independent reason to choose the unexplained, brute moral facts and their dictates. So, even if we allow for vagueness in a realist morality, we can still question why we should prefer this particular realist system with its particular range of vague cases over its alternatives. And in the absence of a compelling, non-circular reason to prefer the realist morality, it seems just as open to us to prefer one of its rivals. This is not yet to say that we have any reason to prefer any morality* over morality; we simply lack a reason to place morality in a privileged position.

A related objection would hold that we have been relying on a certain sort of relativism, and quietly assuming that moral realism is unable to account for any relativism. But perhaps moral realism *can* allow for significant forms of relativism, as has been argued by Sayre-

McCord, and Oddie.¹³ Thus, to the extent that there seems to be a range of plausible actions in a given case (or plausible variations on moral rules governing lying, etc.), this can be accommodated by moral realism. There is no need to abandon a realist morality.

Our response can be rather similar to our response concerning the vagueness objection. Let us grant, as seems correct, that the moral realist can allow for significant forms of moral relativism (allowing, for example, that the morally right action in many cases will depend on local expectations and contracts – perhaps assuming that these expectations and contracts themselves meet certain objective, realist criteria.) The question arises: why prefer this particular realist moral scheme with its range of acceptable relativism to any rival moralities, each of which has somewhat different ranges of acceptable relativism? We have a wide range of possible alternatives, and we do not yet have a reason to focus on the realist morality. Our choice of a particular system again seems arbitrary.

II - The Regress of Normative Standpoints

Suppose that there were a set of realist normative facts – perhaps moral facts. These facts would tell us what objects possess value (or disvalue), what actions to perform in various situations, and so on (from a certain standpoint). But suppose that we ask ourselves ‘Should we act in accordance with these normative facts?’. Why take on the moral point of view (and its standards), the aesthetic standpoint, or any other such point of view? When I speak of ‘taking on a standpoint’ I mean, roughly, choosing to concern ourselves with the norms in question, and attempting to have these norms guide our behaviour (i.e. attempting to act in accordance with the norms). David Copp captures the underlying idea in terms of subscribing to a standard:

¹³ See Geoffrey Sayre-McCord, “Being a Realist About Relativism (in Ethics),” *Philosophical Studies* 61 (1991), 155-76, and Graham Oddie, “Moral Realism, Moral Relativism, and Moral Rules (A Compatibility Argument),”

One subscribes to a standard just in case

- (1) one intends to conform to it and is disposed to conform to it;
- (2) one tends to have a favorable attitude toward oneself for complying with it, and towards one's compliance;
- (3) other things being equal, one tends to have a negative response towards oneself, or one's failure, if one (intentionally) fails to conform; and
- (4) one regards such failures as creating a presumption of liability to a negative response.¹⁴

The question of why we ought to take on the moral point of view (perhaps a rule utilitarianism), or subscribe to this standard is pressing as we could instead take on the moral* point of view (some particular deontological theory), or the moral** point of view (a consequentialism that values happiness, beauty, and knowledge), and so forth.

It seems that any sort of justification for taking up a given normative standpoint could only be given in one of two ways: either (i) from that standpoint itself, or (ii) from some other standpoint. Thus, it might be a moral fact that we should take on the moral point of view. However, this provides no justification of the standpoint for someone who does not currently accept it. The standpoint simply endorses itself.¹⁵ On the other hand, perhaps morality is justified from some other standpoint. It could be that it is aesthetically good to take on the moral standpoint, or rational, or perhaps there is a meta-normative standpoint which endorses the moral

Synthese 117 (1999), 251-74.

¹⁴ David Copp, *Morality, Normativity, & Society* (New York: Oxford University Press, 1995), 84.

¹⁵ Compare a book which declares itself to be the word of god. This declaration, in itself, gives us no reason to accept the claim. Perhaps there are astrological star charts that tell us (via the use of astrology) that we should subscribe to the standards of astrology in forming our beliefs.

point of view.¹⁶ But then we can again raise the original question of justification for this new standpoint. Why should we take on the endorsing standpoint? We are thus left with either circular justifications of any given normative standpoint, or a regress of standpoints endorsing standpoints.¹⁷

It might be suggested that the reason why we should concern ourselves with such normative facts is simply that *they are facts (after all) and we want true beliefs*. They are genuine facts, so of course we have reason to concern ourselves with them, just as we have reason to be concerned with any other facts.

The key point in response is that even if we value having true beliefs, we have only to take note of the normative facts without being motivated by them. We could form the true belief that torturing cats is morally wrong, but our epistemic concern with true beliefs would not yet require that we attempt to avoid acts of cat-torturing. Even if our true belief in this case naturally leads us to be motivated to avoid such acts, our epistemic concerns would not give us reason to act on this sort of motivation; it would still be open to us to try to overcome this motivation. The epistemic standpoint would only require us to have true normative beliefs.

Notice that there will also be true moral* facts, moral** facts, etc., available to us. Suppose morality is a form of deontology. There will also be facts about what act-utilitarianism (here, ‘morality*’) would prescribe for us. For that matter, there will also be facts about what an act-disutilitarian morality** would prescribe (performing whatever action maximizes unhappiness). From an epistemic standpoint, it would be good for us to have true beliefs about

¹⁶ Copp focuses on chains of meta-normative standards when considering the justification of standards. See Copp, *Morality, Normativity & Society*, 37-40.

¹⁷ Both Korsgaard and Copp raise a similar dilemma (though Korsgaard applies it only to naturalistic theories). See Korsgaard, *The Sources of Normativity*, 30, and Copp, “The Ring of Gyges,” 101-3.

the moral facts, the moral* facts, and the moral** facts. But our epistemic concerns do not dictate acting in accordance with the moral facts; nor do they dictate acting in accordance with the moral* facts, or others. In other words, merely having a concern with true normative beliefs will not justify acting in accordance with the standards embodied in the corresponding facts.

How else might we attempt to privilege a realist morality and its dictates over rival standpoints? Perhaps morality and moral properties are distinct, metaphysically, from rival standpoints and properties, and in turn this might give us grounds for preferring morality. For example, suppose that morality is a form of deontology, and that actions in accordance with moral duty have a special non-natural property that supervenes upon them – of a kind entirely lacked by morality* and other rivals. Though there are facts about what act utilitarianism would dictate in the same cases, no special non-natural property would supervene on those actions that would maximize happiness. In this way, then, morality could stand out from amongst its rivals, given its unique metaphysical characteristics.

Ultimately, the proposal does not help us. Suppose that we grant that moral facts are so unique – we might imagine morally right actions to have little supervening haloes (be they natural or non-natural), while morally wrong actions would have horns and tails (and the moral*, moral**, and so on, would lack any similar properties; that is, there would be no moral* haloes, etc.). We arrive at the same questions. Why concern ourselves with the moral haloes and horns? Less flip, why should we concern ourselves with these distinctive moral properties, even if they are metaphysically unlike any possessed by rival schema? If these properties themselves dictate that we should concern ourselves with them (i.e., we morally ought to concern ourselves with these properties), we are in a familiar circle. And if instead there is some external standpoint which holds that we ought to concern ourselves with such unique properties, we then need to ask

why we ought to abide by this external standpoint. Merely being metaphysically distinctive would not constitute adequate grounds for preferring morality over rival codes of conduct.

Next, note that the problem is not merely that an agent who is not currently concerned with a given normative standpoint will lack an internal reason for taking on the standpoint; the problem is deeper. There is no objective, external reason for her to take on the given standpoint. This is because whatever apparent objective reason she would have (regardless of her internal perspective) would only be a reason given some further normative standpoint. And again – when we then examine that further standpoint we find ourselves with a circle or a regress. We have no independent way of showing that an agent should enter into the circle or regress of justifications. And without this, there is no objective, external reason for the agent to take on the original standpoint.¹⁸

We might approach this latter problem in a slightly different way. Imagine that taking on the moral perspective is justified by the rational standpoint (i.e. it is rational to be moral), and that the rational standpoint is justified from a certain epistemic standpoint (being rational will produce a good ratio of true-to-false beliefs).¹⁹ But compare – there could be a wide range of near rational standpoints which could endorse a wide range of near-moral standpoints. And

¹⁸ Thus, the argument of this paragraph could be seen as further developing the worries about external reasons raised by Williams. See Bernard Williams, “Internal and External Reasons,” in Bernard Williams (ed.), *Moral Luck* (New York: Cambridge University Press, 1981), 101-13. Williams shows that it is very hard to imagine what an external reason could be. Here I argue that an agent's concern with a range of prime candidates for such status (external, objective normative facts) can only be justified in circles or by a regress; and neither form of justification is adequate. Therefore, such facts do not provide an agent with external reasons.

¹⁹ I do not endorse this proposal; it is intended only as an example of how chains of justifications using various standpoints could be developed.

these near rational standpoints would be justified from alternative epistemic standpoints. For example, in the first case our epistemic goal might be to maximize our true beliefs. This might lead to a view of rationality which endorses gambling. This in turn could lead to a utilitarian or libertarian political or moral position. But there could also be a broadly epistemic view which would have us minimize our false beliefs, and which would lead to a risk-averse form of practical rationality, which might lead to a more egalitarian political or moral position. And there are, of course, other possibilities. In every case we can produce a chain of justifications for a given standpoint. But why favour one over the others? This would require appeal to further standpoints, which would themselves require justification. We thus have various chains of potential normative standpoints, but there is nothing to secure one of these chains as *the* correct chain (without some sort of circularity).

This possibility of infinite numbers of infinite chains of justification points to a difficulty for Copp's response to the basic problem of trying to justify why we should subscribe to a realist moral standard. He writes:

I am not proposing that we adopt the infinitist position in order to escape normative skepticism. According to the infinitist position, the fact that a standard has a place in an infinite hierarchy of standards is sufficient to show it to be justified. But the point I am insisting on is that the fact that a standard has a place in an infinite hierarchy of standards is *not* sufficient to show it *not* to be justified.²⁰

We can agree, at least for the sake of argument, that the mere fact that a standard has a place in an infinite hierarchy of standards is not sufficient to show it to be unjustified. So far, so good. But a new problem arises when we recognize that there are multiple infinite hierarchies of this

²⁰ Copp, *Morality, Normativity and Society*, 43.

sort. For a being faced with the choice of abiding by one of the chains, there is nothing to justify any of the chains as superior to the others (at least, not without circularity). As such, it seems that any justification that might accrue to a standard via its place in a hierarchy will only be relative to this particular hierarchy; and without any reason to embrace this hierarchy, we have no independent justification for embracing the standard.

Finally, we can consider an important objection. We have been asking ‘Why *should* we subscribe to a realist moral standard?,’ and finding that we end up in regresses or circles of justification. But what standpoint is it that we are appealing to as we assess these other standards – what is the nature of the ‘*should*’ in the above question? Furthermore, what privileges this standpoint – why do we use it in assessing the moral standpoint?

All I have been assuming here is a standpoint that is difficult to define, but that is familiar enough; we might call it the commonsense, armchair standpoint. I do not take this standpoint to be specially privileged in any way. Copp writes

[W]e cannot take up a standpoint that is external to all of our beliefs. Yet the issue here is merely whether we can take up a standpoint that is external to our *moral* views, in the sense that we do not use any moral standards in evaluating moral standards. It seems to me that we can do this. We can evaluate legal standards without making use of any legal standards. [...] The kind of justification of moral standards that is required is not one that would be external in any extraordinary sense. It is simply a justification that is not question-begging, one that does not presuppose that any moral standards are justified.²¹

In the same way, I am attempting only to examine our moral standards, and our subscription to them, from a commonsense, armchair standpoint. Can we arrive at a non-question-begging

²¹ Ibid., 48-9.

justification of our subscription to such standards simply by positing the existence of mind-independent moral facts? I have been arguing that we cannot.

But I would also allow that we can question this armchair standpoint – ‘Why should I concern myself with my armchair, commonsense standards?’ In fact, such questioning simply reinforces the problems we face. For any standpoint we attempt to use in assessing others, we can raise the problem of justification; and we again find regresses or circles.

Motivational Internalism

Perhaps there is a way of distinguishing certain normative standpoints as superior to others. It seems that we are attracted by the moral standpoint, the rational standpoint, and others, while we are repelled by the malevolent standpoint and so on. We might posit a certain motivational internalism, such that the proper recognition of moral facts will naturally lead humans to be motivated to act in accordance with them (and to subscribe to the moral standpoint, more broadly). Could this natural attraction serve as a method of justifying us in choosing the moral-rational set of standpoints over others?

Imagine a group of moths, blackflies, and mosquitoes who have all been attracted to a porch light on a summer evening. They are naturally attracted to the light – their awareness of the light motivates them to approach it. It is also true that there is a light on the porch. However, we can well imagine a gifted moth asking ‘Wait a minute – does anyone know why we’re doing this?’. The mere fact that a group of creatures is naturally attracted to (or motivated by) certain states of affairs does not yet show that they *should* concern themselves with these states of affairs. In the case of humans, even if there are moral facts, and even if we are naturally motivated by them, it seems that we can still reasonably ask whether we should concern

ourselves with these facts. Perhaps we should instead train ourselves to overcome this motivation (just as we learn to control our tempers).

Furthermore, it is far from clear that we all are in fact motivated by the moral facts (should such exist). Consider the stark differences in intuition (and motivation) that we find between utilitarians, deontologists, and virtue theorists concerning many particular situations. We humans might be motivated by what we believe to be ‘moral’; but it seems that we are tracking very different systems in these beliefs. In other words, it seems that some of us might be tracking morality*, others morality**, and so on. Some of us are motivated by the moral* facts, others by the moral** facts. And we as yet have no reason to abandon these views in favour of morality – whatever this turns out to be. We just have different sets of standards, each of which can be justified in light of some further normative standpoint(s); but with nothing (non-circular) to secure any such chain as superior to the others.

Reasons Internalism

In the previous section we considered whether moral facts which have a certain attraction for us (humans) would be enough to justify our taking on the moral standpoint. We can now turn to a rather different attempt to show that moral facts are worthy of our attention. Essentially, the claim is that morality and moral facts provide us with reasons for action regardless of our interests, desires, or nature.²² They provide us with categorical reasons that bind us regardless of

²² Richard Garner argues that such ‘reason-giving’ facts would be metaphysically queer, and that Brink’s distinction between motivational and reasons internalisms ultimately does not solve the queerness problem for moral realism. I tend to agree with Garner, but will assume in this paper that such objectively prescriptive facts are possible. See Richard Garner, “On the Genuine Queerness of Moral Properties and Facts,” *Australasian Journal of Philosophy* 68 (1990), 137-46. For the original queerness objection to moral realism, see J.L. Mackie, *Ethics: Inventing Right and Wrong* (New York: Penguin Books, 1977), 38-42.

our interests, though these reasons need not be seen as categorical in the sense of being overriding. This seems to capture a key intuition which leads many to moral realism: that there are moral facts in the world, and that they are binding on all agents – be they human, Martian, or other. To ask for a justification for following the dictates of morality is not legitimate – they simply are intrinsically binding, without reference to some further justification.

Unfortunately, the proposal is subject to strong objections. To begin, we can return to the analogy drawn earlier between divine commands and realist moral facts. The focus in the earlier discussion was the problem of arbitrariness. Here we can shift to the problem of authority. A common objection to divine theories is that there can be no genuine basis given for why we should obey God's commands.²³ If God simply commands us to obey his commands, we have no real justification. Why listen to this additional command any more than we would have listened to the others? It has no more authority than his other commands. Perhaps we should obey God's commands because he is our creator, and we should obey the commands of those who create us. But where does this supposed obligation come from, and what justifies it?

Parallel points can be raised with respect to moral facts (and their associated reasons). On what authority can they command us? How are they any different from a God who commands us? If they command us to obey them, we have no real justification. In other words, even if it is a moral fact that we ought to act morally, this seems to hold no more force than God commanding us to obey his commands. On the other hand, if there is some further rule 'Agents simply should obey moral facts,' we can ask what justifies this further rule. This leads to a regress parallel to that described above.

One might attempt to defend a modified position, according to which God's commands are commands that all agents should obey, simply *because they are God's commands*; we would need no further justification than this. In other words, this person would not claim that we ought to obey God's commands *because* he commands us to do so, or *because* he is our creator, etc. Rather, she would claim that it is simply a brute fact that we ought to obey divine commands.

In the realist analogue, the claim would be that moral facts are binding on us by their very nature, that we ought to abide by them, and that there is no need for further justification. Shafer-Landau makes a proposal along these lines, suggesting that "moral facts are themselves intrinsically reason-giving, i.e. supply reasons for action regardless of the content of specific moral demands and their relation to other intrinsically or necessarily reason-giving kinds of considerations".²⁴ While he acknowledges that such intrinsically reasoning-giving facts may seem mysterious, Shafer-Landau attempts to show that we do, in fact, embrace such in other domains, particularly epistemology:

I think that intrinsic normativity is ineliminable. To see this, consider the parallels between conditions of epistemic and moral assessment. We say that agents, if they have reason to believe anything at all, have reasons to believe the truth, and to conform their reasoning to truth-preserving schemas, even if believing the truth is not conducive to the goals they set themselves.²⁵

²³ See Christine Korsgaard, *The Sources of Normativity* (New York: Cambridge University Press, 1996), 29 for an example of such an argument applied to the divine command position of Pufferndorf, and Hobbes' appeal to a sovereign.

²⁴ Shafer-Landau, *Moral Realism*, 204.

²⁵ *Ibid.*, 205.

I believe that there is intrinsic reason to think that two and two are four – *the fact itself provides one with reason to believe it*. [...] The basic idea here is that certain things can be intrinsically normative – reason-giving independently of the value actually attached to them by agents.²⁶

Shafer-Landau's example initially seems quite compelling – after all, don't we have reason to believe the truth *simply because it is the truth* (even if this can be overridden in some cases by pragmatic, moral, or other concerns)? But compare another case where we find intrinsic reasons. If you desire that p, then it seems that you have reason to promote p, other things being equal. That is, the very fact that you desire something gives you at least a *prima facie* reason to promote or pursue it; this is the nature of a desire. But now notice – there is an intrinsic reason here that arises out of the psychological state of desiring, not because of some external fact about the object of the desire. Simply insofar as you *desire* that p, you will have reason to promote p; we need not appeal to any additional, special normative properties of p itself. Similarly, in the case of belief, note that to believe that p *is precisely to believe that p is true*; it is inherent to the psychological state of belief that it aims at truth.²⁷ So while we have reason to believe the truth in general, and specifically that two plus two equals four, this is because the psychological state of believing by its very nature aims at truth (and it is true that two and two equal four). As in the case of desiring and its reasons, intrinsic reasons here arise out of the nature of the psychological state, and we need not posit any intrinsic reason-givingness that attaches to all true propositions as Shafer-Landau proposes. Thus, Shafer-Landau's apparent appeal to epistemic normativity fails, because he is actually focusing on the normativity of belief – and here the normativity is a

²⁶ Ibid., 206, emphasis added.

matter of this psychological state, and not because of external reason-giving facts of the kind that Shafer-Landau is positing in the case of moral realism.

One could claim that just as beliefs aim at truth inherently, perhaps all actions must aim at moral rightness inherently. But this is dubious; there are many valid kinds of reasons for our actions, and morality does not play the same role for action as truth does for belief. We can properly perform actions for aesthetic, epistemic, pragmatic, and other non-moral reasons. For example, a musician can appropriately perform the action of tuning his guitar for purely aesthetic reasons. Again: to believe that *p* is to believe that it is true that *p*, and no such relation exists between acting and moral rightness. More broadly, even if (mysteriously enough) all actions were somehow to aim at moral rightness, why not instead perform actions*, which aim at moral rightness* inherently?

Returning to the analogy with the divine command theorist, further questions arise for realist proposals where we simply ought to act in accordance with the moral facts. What is the nature of the realist brute ‘ought’ or reason? If it is a moral ‘ought,’ we are no better off than in a case where God commands us to obey his commands. If it is not a moral ‘ought,’ what kind of an ‘ought’ is it, and why should we occupy ourselves with it? These basic ‘oughts’ must have a nature, and we can legitimately inquire into this. And when we do, we are again left with the dilemma of circular or regressive justifications.

We might worry that the discussion has focused on too narrow a conception of normative facts and standpoints. It might make sense for us to balance moral demands against epistemic demands or aesthetic demands, etc. We might also lack non-question-begging grounds for

²⁷ Consider the incoherence of sincerely asserting that “It is true that *p*, but I do not believe that *p*”. For a defence of these claims concerning belief, see Jonathan E. Adler, *Belief’s Own Ethics* (Cambridge: MIT Press, 2002).

declaring these standpoints superior to rival moral*, epistemic*, etc. standpoints. But what if there were an overarching, overriding normative standpoint?

Owen McLeod proposes an account of an ‘ought’ that ‘is not relative to any particular normative framework [.] This “ought” would not be a moral, prudential, legal, aesthetic, or religious “ought,” but rather an *unqualified* or *just plain* “ought.”’²⁸ We would thus have an overall standpoint, one that would adjudicate conflicts of moral and aesthetic duty, and so on. McLeod elaborates further:

What, one might ask, makes a JPO [“just plain ought”] matter? In a word, nothing. It just matters. Its mattering is intrinsic to it. [...] On the view I have in mind, according to which the relative “oughts” matter only if they entail a JPO, there is only one intrinsically mattering “ought”: namely, JPO. This view promises a unification of the normative realm – the cosmos of duty – and hence a resolution, at least in principle, of conflicts among the “oughts” that matter²⁹

On this account then, there is an overarching normative standpoint that delivers overall ‘just plain oughts,’ and that matters intrinsically, simply by its very nature. However, the same questions arise (*mutatis mutandis*). What could justify this overarching normative standpoint? For example, what if you decide to abide by JPOs, but I decide to abide by JPO*s, i.e. the dictates of a similar but rival overall standpoint? Perhaps I just plain ought to concern myself with what I just plain ought to do; but then again I just plain ought* to concern myself with what

²⁸ Owen McLeod, “Just Plain ‘Ought,’” *The Journal of Ethics* 5 (2001), 269. Note that McLeod stresses that he is simply presenting a possible account of such ‘oughts’, and does not necessarily endorse the account. Further, he does not specify the account in terms of normative realism. Still, his proposal will appeal to many who would consider themselves normative realists.

²⁹ McLeod, “Just Plain Ought,” 274-5.

I just plain ought* to do (as should* you). What secures one as superior to the other? McLeod suggests that JPOs 'matter' intrinsically. So he might suggest that we should concern ourselves with JPOs because they matter, while JPO*s do not. The response should now be predictable – why should we care about mattering, and not mattering* (or mattering**, for that matter***)? If there is a further standpoint, we can again demand a reason to abide by its dictates. If not, the categorical standpoint merely endorses itself, and we will again be unable to provide non-circular reasons to take on this standpoint to those who do not already do so.

Finally, some might feel that there must be a position available to the moral realist analogous to a foundationalism in epistemology - a position which avoids a regress or circle of justifications. Note that in the case of epistemology we can claim that certain beliefs are justified as foundational, but this also requires us to appeal to an additional, external standard - truth. Foundational beliefs are those which, because of their content, or how they were formed, are likely to be true. So the epistemological foundationalist relies on a background value system centred on obtaining truth. There are no bare 'justified' beliefs - it is justification in light of trying to achieve true beliefs. As such, a parallel position in the metaethical domain would require a background value system. And if this is so, the dilemma arises again as we need to inquire into the background value system which is being used to evaluate the moral facts.³⁰

I take it then, that the mere existence of realist normative facts and standpoints is not adequate to show that we ought (in some broad, overall sense) to abide by them. Note also that this is a strong result. It is not merely that the moral realist would be unable to convince an extreme moral sceptic to become moral. Rather, even for those who wish to be broadly 'good'

people, and treat others well, there is no reason to take on the specific norms dictated by any realist moral facts over rival moral schemes. Consider again the stark differences that we find between utilitarians, deontologists, and virtue theorists concerning many particular situations. It seems that some of us are tracking morality*, others morality**, and so on. And we as yet have no reason to abandon these views in favour of morality – whatever this turns out to be. We simply have different sets of standards, each of which can be justified in light of some further normative standpoints; but with nothing (non-circular) to secure any such chain as superior to the others. The realist norms are just one possibility within a range of broadly moral options.³¹

³⁰ And note that the same regress would apply to the epistemologist. Why choose the epistemic standpoint (which, for example, values acquiring true beliefs over avoiding false beliefs) rather than the epistemic* standpoint (which emphasizes avoiding false beliefs rather than acquiring true beliefs)?

³¹ Ancestors of this paper, some distant, were read at Brown University (2000), the Canadian Philosophical Association (2000), the MidSouth Philosophy Conference (2001), the Society for Realist/Anti-Realist Discussion (with the Eastern APA, 2001), and at Brooklyn College (2003); I would like to thank all of these audiences for their comments and suggestions. In particular, I would like to thank Susan Dimock and Pascal Massie (my commentators at the CPA and the MidSouth), Robert Binkley, Dan Brock, Eric Dayton, Jamie Dreier, Jeremy Fantl, Darren Hibbs, Jennifer Lackey, Max Lechien, David Matheson, Baron Reed, and Ernie Sosa. Finally, I would like to thank Nancy Simco and the anonymous referees of this journal for many helpful suggestions.

