Philosophers' Imprint

volume 8, no. 6
july 2008

Metaethics & the Autonomy of Morality

Tristram McPherson

University of Minnesota Duluth

© 2008 Tristram McPherson
<*www.philosophersimprint.org*/008006/>

1. Introduction

Since the publication of G. E. Moore's *Principia Ethica* it has become commonplace for philosophers to distinguish between questions in metaethics and those in normative ethics.[1] A sympathetic characterization of the century of self-consciously metaethical research that followed would emphasize the extraordinary development both in our understanding of the central metaethical problems and in the sophistication of the theories elaborated to meet them. However, some are not so sympathetic. In this paper, I examine one source of distrust in metaethical research: its apparent tension with the notion that morality is autonomous.

To begin, I briefly sketch how I am thinking of metaethics, of the autonomy of morality, and of the tension that can appear to exist between them. One traditional conception of metaethics takes it to concern only the analysis of moral language.[2] However, contemporary philosophers typically use the term more expansively.[3] Here, I use the term to pick out elements common to these contemporary discussions. This common core encompasses moral ontology and moral psychology as well as moral semantics. By contrast, normative ethics (sometimes also called 'substantive ethics') concerns the structure and content of the correct moral evaluation of agents, states of affairs, and actions. Normative ethical theories typically offer accounts of moral value and moral reasons, of virtuous character traits, of rightness, and of the relationships between these.

1. The word 'metaethics' came into regular philosophical usage much later. Moore's discussion emphasized the distinction between two kinds of questions about morality, which in turn encouraged the development of independent research programs in what would later be called "metaethics" and "normative ethics".
2. Obvious worries about this conception arise from Quinean concerns about analyticity. Such worries led Zimmerman to argue for an "ontological turn" in metaethics (1980, 653).
3. Various broader conceptions of metaethics are suggested in Darwall, Gibbard, and Railton (1997, 7), Miller (2003, 2), and Sinnott-Armstrong (2006, 6).

A domain of inquiry is autonomous in the sense I intend if results in other domains do not contribute to the justification of theories in that domain.[4] The *autonomy of morality* is thus the claim that:

AUTONOMY: Non-moral theses are irrelevant to the justification of normative ethical theories.

I will sometimes refer to this as the *autonomy thesis*.[5] The idea that normative ethical theorizing is autonomous may appear intuitively attractive. Thus, it may seem on reflection that the only legitimate grounds for abandoning a basic moral principle will involve tension with other moral commitments.[6] Further, the autonomy thesis promises to explain the plausibility of Hume's dictum that one cannot derive an 'ought' from an 'is': if morality is autonomous, it will always be an error to derive a moral conclusion from strictly non-moral premises.

4. My use of the term 'autonomous' here is inspired by Harman's distinction between "naturalism" and "autonomous ethics" in moral theory (2000, 79). However, I use it to make a slightly different distinction. Harman suggests that any approach to ethics not dominated by a concern to place values in the natural world discovered by science is autonomous (2000, 79). On my gloss, science is only one domain with respect to which morality might be autonomous. For example, I also take divine command theory to be incompatible with the autonomy of morality, as it would render evidence of God's will probative for normative ethical theory.

5. The restriction of the scope of the autonomy thesis to theories, understood as consisting of fundamental principles, is crucial to the plausibility of the autonomy thesis, because it is implausible that evidence from other domains is not relevant to the justification of derivative principles or of particular moral judgments. For example, consider the moral principle "It is wrong to fire loaded guns at innocent people" and the particular moral claim "Lee's firing his gun at John is wrong." A natural way to justify such claims would appeal in part to more basic normative principles, and in part to the typical or predictable causal consequences of firing loaded guns at people. The autonomy thesis should thus be understood as compatible with a role for non-moral theses in explaining why a normative moral principle applies to a particular case, or supports a derivative principle.

6. Compare Harman's description of the autonomous ethicist's reasoning about a case: "[W]e should continue to believe [that a criminal has sufficient reason not to harm his or her victims] unless such continued belief conflicts with generalizations or other theoretical principles internal to ethics that we find attractive because they do a better job at making sense of most of the things we originally believe" (2000, 87).

The apparent tension between metaethics and the autonomy of morality derives from the plausibility of taking metaethical theorizing to have two features. First, metaethics can seemingly bear on the justification of normative ethical theories. Shelly Kagan makes an especially strong version of this claim that metaethics is relevant to the justification of normative theory: "[I]n the course of defending a given theory about the foundations of normative ethics, when we try to explain why it is that the various features of that theory should seem attractive and plausible, inevitably the claims we make will themselves simply be metaethical claims" (1998, 6). Kagan's claim is controversial. However, even if it does not seem inevitable that metaethics plays this sort of foundational role in normative theorizing, the idea that our best accounts of moral properties or moral concepts couldn't count for or against the justifiability of a normative theory might seem puzzling.

Because metaethical claims are themselves claims about morality, this point does not by itself conflict with the autonomy thesis. The apparent problem arises from conjoining this point with reflection on the task of metaethical theory. On a plausible gloss, the central task of metaethics is to explain how morality fits with the other elements of our broader conception of the world. (Note that this description of the task of metaethics is consistent both with conceiving of these other elements in exclusively naturalistic terms, and with thinking that our best overall conception of the world must augment or transcend the scientific worldview.) This suggests that coherence with well-justified general accounts of ontology and semantics counts in favor of a metaethical theory. Hence, it seems natural for metaethical theorizing to involve appeal to non-moral premises.

Together, these two apparent features of metaethical theorizing appear to conflict with the autonomy of morality. Thus, if non-moral premises are relevant to the justification of accounts of moral semantics or ontology, and these accounts in turn play a role in justifying foundational normative theories, then non-moral premises can play an indirect role in justifying normative ethical theories. This conclusion appears to be inconsistent with the autonomy thesis.

If one is sympathetic to the autonomy thesis, this apparent tension

might suggest that there is some error implicit in metaethical theorizing. Some philosophers have recently offered diagnoses of this alleged error. Thus Christine Korsgaard suggests that the dialectic between metaethical realists and antirealists rests on a mistaken conception of the function of moral concepts. She claims that when our thinking about the subject is corrected, "we will not be inclined to think that there is a difference between doing 'meta-ethics' and doing 'normative' or practical ethics" (2003, 121 n. 44). Ronald Dworkin also finds error in much apparently metaethical theorizing, arguing that it rests on mistaking substantive ethical claims to be metaphysical or "philosophical" (1996, 100).

This paper examines three ways to address the apparent tension between metaethics and the autonomy of morality. In section 2, I examine the attempt to address this tension by rejecting the claim that metaethics can be relevant to the justification of normative ethical theories. I evaluate this strategy in part by examining the reasoning that leads Dworkin to the diagnosis mentioned above. In section 3, I examine variants on the proposal to vindicate the autonomy of morality by appealing to a deep contrast between practical and theoretical reason. Korsgaard's diagnosis is best understood as following from such a commitment. Finally, in section 4 I suggest what I take to be the most promising way to resolve the apparent tension between metaethics and the autonomy of morality.

2. Neutrality and Quietism

One way to resolve the apparent tension between metaethics and the autonomy of morality would be to endorse what I call the *neutrality thesis:*

> NEUTRALITY: No metaethical theory is relevant to the justification of normative ethical theories.

The neutrality thesis resolves the tension between metaethics and the autonomy of morality by denying the first of the apparent features of metaethical theses mentioned above: their possible normative implications. Thus, if metaethics was necessarily neutral between normative theories, the sensitivity of metaethical theory to non-moral considerations (the second apparent feature of metaethical theories) could not impugn the autonomy of normative ethical theorizing.

One could think of the neutrality thesis in either of two ways. First, one could think of it as a *hypothesis* about metaethical theories, to be evaluated by examining such theories. However, it is easy to find counterexamples to this neutrality hypothesis, in the form of metaethical theories with direct normative implications. Consider a simple example: the analytic utilitarianism that Moore (1903, 17–18) attributes to Jeremy Bentham. This semantic theory states that the word 'right' means "conducive to general happiness". Analytical utilitarianism is an account of the semantics of a central piece of moral vocabulary, and hence a paradigmatic metaethical theory. However, it also transparently has implications for the content of the correct normative theory. The neutrality thesis thus seems implausible if read as a hypothesis concerning the possible content of theories that we intuitively count as metaethical.

One could instead think of the neutrality thesis as a constraint on a theory's counting as metaethical. Thinking of the neutrality thesis as a hypothesis runs into trouble where seemingly paradigmatic metaethical theories have normative implications. By contrast, a neutrality *constraint* would entail that the same data showed such theories not really to be metaethical. Thus, accepting such a constraint might seem to promise to resolve the tension between metaethics and the autonomy of morality. In the remainder of this section, I examine the implications of accepting this neutrality constraint, focusing on issues highlighted by Dworkin's paper "Objectivity and Truth: You'd Better Believe It". This discussion will show that accepting the neutrality constraint is objectionable in virtue of obscuring philosophically important distinctions between metaethical theories.

Dworkin's main target is what he calls the "neutral Archimedean skeptic": a philosopher who combines first-order normative ethical commitment with a deflationary metaethics such as noncognitivism

or error theory (1996, 92ff.). However, part of Dworkin's central argumentative strategy seemingly involves deploying the neutrality constraint. Here I focus on this use of the constraint and its implications for metaethics.[7] Dworkin applies a two-stage test to an array of central metaethical proposals in his paper. First, test whether the seemingly metaethical claim can be interpreted as having normative import. Second, test whether it can be interpreted as having distinctively metaethical import (Dworkin 1996, 97).[8] If the result of the first test is positive and the result of the second is negative, then the seemingly metaethical claim will be shown to be a disguised normative claim. I consider these tests in turn, and explain how the neutrality constraint is indispensable to Dworkin's application of the second test.

Dworkin's first test examines whether we can interpret pieces of seemingly metaethical discourse as expressing normative ethical claims. He suggests that people typically use seemingly metaethical adverbs, such as 'subjectively' and 'really', to clarify their normative ethical opinions. He thus notes that someone might use the word 'subjective' to signal that his claim that soccer was bad meant that he had a reason to avoid soccer that is not shared by those who have a taste for the game (1996, 98).

In order to properly perform the first test, we need to distinguish *agent-relativity* from *evaluator-relativity*. Agent-relativity is a feature of value or reasons in a normative theory. Very roughly, think of a state of affairs as having agent-relative value if not all agents have a reason to bring it about.[9] Dworkin's gloss on his example points out that the word 'subjectively' might be used to claim that the value of soccer is only agent-relative, the relevant agents being those who share a taste for the game. Only such soccer-enjoyers have reasons to bring it about that they watch soccer. On the other hand, if everyone had reason to bring it about that people watch soccer, then the value of watching soccer would be agent-*neutral*. Agent-relativity is a normative-level feature because it describes a distinctive structure in the relation between value and reasons, two normative notions. Dworkin's exemplary use of 'subjective' is thus best understood as making a normative claim, as he suggests.

Nonetheless, there appears to be a problem with Dworkin's use of this test: some uses of the word 'subjective' indicate the *evaluator*-relativity of a moral term. Evaluator-relativity is an apparently metaethical notion. Thus the semantics of 'right' would be evaluator-relative if the truth of statements of the form 'Φ-ng is right' depended upon who was making or evaluating the statement. Consider an example: people sometimes say "values are subjective" as a way of endorsing a simple subjectivist metaethics. According to simple subjectivism, utterances of the form 'Φ-ing is right' are true in a person's mouth just in case and because she approves of Φ-ing. Such uses of 'subjective' thus appear to fail Dworkin's first test.

Other vocabulary also exhibits the possibility of both first-order and meta-level uses. Consider another of Dworkin's examples, the word 'really'. If I say "That book isn't really red", I might be attempting to convey one of (at least) two things. First, I might be attempting to get you to see that you've gotten the color of the book wrong: it's more of an orangey-brown. Second, I might be attempting to convince you on metaphysical grounds that colors are not real properties of physical objects like books. The possibility of such vocabulary's being used to express both metaethical and normative judgments thus appears unsurprising, being an instance of a general phenomenon.

Dworkin's first test thus appears less than fully successful. He has

7. My concern here is to tease out the implications of accepting the neutrality constraint. I thus focus on elements of Dworkin's discussion that appear to pursue this strategy. Dworkin's argument in (1996) is complex and not always clear, and I do not claim to offer a fully adequate reading of it in this brief discussion.

8. Here and below, this gloss must be read extensionally. While Dworkin's leading targets in this paper are clearly metaethical theories, he generally prefers to talk of collapsing the distinction between "first-order" evaluative claims and "philosophical" claims about morality (*e.g.* 1996, 100).

9. The locus classicus for this distinction is Nagel's discussion of subjective and objective reasons (1970, Ch. X). Here I abstract away from the significant complexities that a fully adequate account of agent-relativity would face.

shown that a variety of *vocabulary* seemingly used to express metaethical claims can be used to express normative claims. However, he hasn't shown this to be true for such vocabulary as it is used in contexts where we would intuitively take it to have metaethical import. This suggests that Dworkin's second test is the crucial one. The goal of Dworkin's second test is to show that, for a range of putatively metaethical vocabulary, no distinctively metaethical interpretation is in fact available. Were this test successful, it might be possible to accommodate the apparent counterexamples to the first test just discussed as instances of diverse *normative-level* uses of 'subjectively' and 'really'.

Dworkin's leading strategy in applying his second test to apparently metaethical claims is to argue that those claims entail normative commitments (1996, 99–112, cf. especially 100). The example of analytic utilitarianism, mentioned above, suggests that at least some apparently metaethical theories have such implications. However, Dworkin repeatedly moves from the premise that an apparently metaethical theory has normative implications to the conclusion that it is therefore actually a normative theory. This inference seems to presuppose that no genuinely metaethical theory may have normative implications. In other words, it presupposes the neutrality constraint on counting as a metaethical theory.

Consider Dworkin's treatment of an example of an apparently metaethical thesis with obvious normative implications: the thesis that the property of rightness is identical to the property of maximizing happiness. Call this thesis *metaphysical utilitarianism*. This appears to be a metaethical thesis about moral ontology that entails a normative theory. This normative theory, *simple act utilitarianism*, states that an act is right just in case and because, among the options available to the agent of the action, it would bring about the most happiness in its consequences. However, Dworkin rejects the claim that metaphysical utilitarianism and simple act utilitarianism are distinct theories:

> The identity claim about rightness is the upshot, in a parallel way, of a substantive moral thesis — utilitarianism — and it is exhausted, in the same way, by that substantive thesis. There is no difference in what two people think if one thinks that the only thing that can make an act right is its maximizing power, so that it makes no sense to evaluate rightness in any other way, and the other thinks that the property of rightness and the property of maximizing power are the very same property. The second opinion uses the jargon of metaphysics, but it cannot add any genuine idea to the first, or subtract any from it. It sounds more philosophical but it is no less evaluative. [1996, 100–1]

I read this series of claims — that the ontological claim is "exhausted" by the normative thesis, that "there is no difference in what two people think" if one thinks the normative thesis is true and the other thinks the ontological thesis is true, and that the ontological thesis merely "uses the jargon of metaphysics" to "sound" more philosophical — to be forceful ways of claiming that, while the two claims might appear distinct, they are in fact identical in content.

Dworkin argues that an extremely broad range of apparently metaethical theories have significant normative consequences (1996, 98ff.).[10] This in turn has striking implications for our understanding of metaethics. Discussing a series of seemingly metaethical theses that he calls "further claims", he says, "[W]e are trying to decide, not whether the further claims can be translated to make them seem more philosophical or metaphysical but whether we can understand those philosophical translations as themselves anything but first-order evaluative claims" (1996, 100).

Dworkin intends his second test to show that the answer to this question is 'no'. This suggests that, just as with metaphysical utilitarianism, a surprising range of seemingly metaethical claims can

10. It is not entirely clear whether Dworkin intends his conclusion to be perfectly general, or whether he would be prepared to countenance the existence of *some* genuinely metaethical claims. Dworkin's discussion prompted Dreier (2002) to offer a careful articulation of a metaethical theory that is claimed to have no normative implications.

be revealed to be normative ethical claims because they fail his second test. Dworkin thus appears to defend a kind of quietism about metaethics.[11]

I suggested at the beginning of this section that accepting the neutrality constraint seemingly promises to resolve the apparent tension between metaethics and the autonomy of morality. The discussion of Dworkin shows that accepting the constraint would also entail that many philosophers who take themselves to be addressing distinctively metaethical questions are deeply confused. This is because, according to the view sketched, many of these philosophers are actually addressing first-order normative questions rather than metaethical questions.

While Dworkin's argument clarifies the consequences of applying the neutrality constraint, his own examples help to illustrate why there is good reason to reject it. Consider again Dworkin's identification of metaphysical and normative utilitarianism, and its implication that the apparently metaphysical character of the former theory is illusory. The problem with this identification can be traced to one of the central motivations for distinguishing metaethical from normative theories. This is that it is possible to agree about the correct normative theory while disagreeing about the correct metaethical theory.[12]

The deep motivation for this claim can be brought out especially clearly by considering the disagreement between two theists who agree about which acts are right. Suppose, for example, that they both accept simple act utilitarianism. Suppose further that they also agree that, necessarily, an act is right just in case God approves of it. These theists can still disagree about the truth of the divine command metaethics: the claim that the above biconditional is true because God's approval *makes* actions right. The natural lesson to draw from this example is that distinct metaethical theories can entail the same normative ethic, because coextension — even necessary coextension — is not sufficient to settle metaethical theory identity. Semantic, determination, and explanatory relations are also of paramount importance.

The defender of the neutrality constraint is thus faced with a dilemma. On the one hand, Dworkin's strategy of identifying apparently metaethical theories like metaphysical utilitarianism with intuitively normative theories like simple act-utilitarianism fails. This is because, as we have just seen, simple act utilitarianism (Dworkin's exemplary normative theory) is compatible with metaethics that are incompatible with metaphysical utilitarianism. Indeed, it is even *entailed* by metaethical theories that deny the identity of the properties of rightness and happiness maximization. For example, it is entailed by some versions of the theory that identifies the property of rightness with the property of being endorsable by an ideally informed impartially sympathetic observer.

On the other hand, it will not help to insist that, in virtue of their implications, these apparently metaethical theories are nonetheless normative theories. Consider again the imagined theorists canvassed thus far: the analytic utilitarian, the metaphysical utilitarian, the divine-command utilitarian, and the ideal-observer utilitarian. These theorists disagree about morality. The first theorist disagrees with the rest about a semantic thesis: whether 'morally right act' and 'act that maximizes happiness' are synonymous. The rest disagree amongst themselves concerning a metaphysical issue: namely, whether the property of being right is identical to the property of maximizing happiness, to the property of being approved of by God, or to the property of being approved of by an ideal observer. Because these disagreements concern moral semantics and metaphysics, they are naturally understood as metaethical disagreements.

Simply insisting that these disagreements should be described as normative ethical disagreements would not avoid the apparent tension

11. In a reply to critics, Dworkin rejects this label, saying: "I don't like the term 'quietist' for the reason I mentioned: it suggests that some more 'robust' sense of objectivity makes sense but is wrong" (1997). I intend my use of 'quietism' to pick out exactly this denial of the intelligibility of distinctively metaethical content.

12. This point has been suggested across the history of discussion of metaethics and normative ethics. See for example Frankena (1951, 45), Zimmerman (1980, 659), and Kagan (1998, 5), who mentions but does not explicitly endorse this possibility.

between metaethics and the autonomy of morality. Recall that this apparent tension is generated in part by the pressure to reconcile our accounts of moral semantics and ontology with our general semantic and ontological theories. No matter what we *call* them, the semantic and metaphysical views identified here face such pressure.

Dworkin's use of the neutrality constraint on metaethical theories seemingly promised to undercut the apparent tension between metaethics and the autonomy of morality, by showing that the apparent metaphysical character of many metaethical theories is an illusion. If sound, this use of the neutrality constraint would vindicate the current of distrust in metaethical research mentioned at the beginning of this paper. This is because this conclusion suggests that most metaethicists are badly confused, either addressing pseudo-questions or doing normative ethics without realizing it.

In this section, I have argued against the attempt to use the neutrality constraint to resolve the apparent tension between metaethics and the autonomy of morality. I first argued that a metaethical theory that entails some normative thesis should not be conflated with that normative thesis (or any other). This is because metaethical theories with normative implications also have further distinctive content. I then argued that the appeal to the neutrality constraint either is objectionable in virtue of obscuring this content, or fails to resolve the apparent tension between metaethics and the autonomy of morality. This suggests that we should look elsewhere to resolve the apparent tension between metaethics and the autonomy of morality. In the next section, I examine a cluster of strategies that seek to vindicate the autonomy of morality by appealing to a contrast between practical and theoretical reason.

3. The Practicality of Morality

I suggested in the Introduction that the autonomy of morality can appear intuitively attractive. However, even if one is sympathetic to the autonomy thesis, it would be more satisfying to have a plausible explanation of why morality must be autonomous than to take it as a brute fact. In this section, I consider variants on an important strategy for explaining why morality is autonomous.

This strategy begins by noting that normative ethical principles are paradigmatically practical. Thus, determining that an action would be wrong plays a distinctive and at least often decisive role in deliberation regarding whether to perform it. This point is common ground: the distinctive motivational and normative aspects of morality are central explananda in current metaethics. However, one distinctive and influential line of thought in contemporary ethics insists that unlike theoretical reasoning, moral theorizing must proceed from the perspective of first-person practical deliberation. This *deliberative requirement* might seem to make the divide between theoretical and practical reasoning deep enough to preclude theoretical reasoning from being relevant to the justification of moral theories. Accounts that develop this contrast thus promise to vindicate the autonomy of morality.[13]

I begin this section by sketching an apparent counterexample to the autonomy of morality, and the basic form of a response to this counterexample that appeals to the deliberative requirement. I then consider three ways that this response could be developed. The first two alternatives are variants of the idea that practical and theoretical reason deliver distinctive evaluations of a given proposition. The final alternative satisfies the deliberative requirement by characterizing a distinctive notion of practical content. Finally, I summarize the consequences for metaethics of these approaches. As the issues surrounding these proposals are both deep and complex, I do not evaluate or even fully develop them. Instead, I sketch just enough structure to ascertain the relevant consequences of such views.

Derek Parfit's central argument in Part Three of *Reasons and Persons* suggests an apparent counterexample to the autonomy of morality. Parfit argues that the metaphysical conception of personal identity

13. Darwall, Gibbard, and Railton note something like this link in their characterization of what they call "practical reasoning theories": "It is ethics' intrinsically practical character, its hold on us *as agents* that explains the open question and, they say, marks ethics off from science" (1997, 9).

that matters for morality is very thin. He then uses this metaphysical premise to argue for the parity of moral significance between intrapersonal intertemporal considerations, and interpersonal considerations (1984, 199–350).

Korsgaard claims that Parfit's argument misses its mark, because the conception of personal identity relevant to deliberation is one required by practical reason, not one identifiable by metaphysical inquiry (1989b, 112). She argues further that we all have sufficient practical reasons to view ourselves as unified persons. These reasons are twofold. First, there is the "raw necessity of eliminating conflict". Korsgaard claims that unless one's parts — both one's subpersonal processes and one's temporal slices — coordinate in such a way as to allow one to act as a unified agent, anything resembling action would become impossible, and practical paralysis would result. Second, the deliberative standpoint requires that a deliberator see himself as something beyond a mere collection of desires or other subpersonal states. Thus, Korsgaard observes that one cannot deliberate about what to do merely by waiting to see which of one's desires wins out. One can only choose from a practically unified standpoint (1989b, 110–1).

Suppose that Korsgaard is correct, and that the conception of identity that we have practical reason to accept is distinct from the conception that Parfit defends. This would suggest a test case for the three ways of developing the deliberative requirement that I will now sketch: Parfit's metaphysical arguments suggest one thing about personal identity and normative ethics, while practical reasoning presupposes another. If the accounts to be sketched are to vindicate the autonomy of morality, they must explain why Parfit's arguments are not relevant to normative theorizing.

Consider first the idea that one can have divergent practical and theoretical reasons concerning the acceptance of a given proposition. This distinction suggests a possible contrast between the metaethics for which we possess the best evidence, and the metaethics that we have decisive practical reasons to accept. How one develops this proposal depends upon whether one takes practical reasoning to require agents to believe in a metaethics, or merely to have some species of nondoxastic commitment towards it. I will call these the "two norms" and "two attitudes" proposals, respectively.

The minimal form of the two norms proposal suggests that the practical and theoretical perspectives deliver distinctive norms for belief in a given proposition. This claim must be augmented in two ways if it is to vindicate the autonomy of morality. First, the autonomy of morality is supposed to mark a contrast between the justification of normative ethical principles and the justification of ordinary beliefs. However, circumstances can probably make it *useful* (for example) to have arbitrary beliefs on any topic.[14] The special role for practical norms thus needs to be constrained in its scope. I set this problem aside here.

Second, the autonomy of morality requires that relevant practical norms for belief must always trump theoretical norms — in Rawlsian argot, practical norms must have lexical priority. Briefly consider two alternatives. On the one hand, epistemic and practical norms cannot be (even roughly) commensurable. This is because such an alternative would allow the possibility that practical reasons in favor of believing a normative principle could be outweighed by countervailing theoretical considerations, thereby violating the autonomy thesis. On the other, the possibility that it could be fully rational to believe a proposition in virtue of practical considerations, while also fully rational to disbelieve it for epistemic reasons, is also unappealing. Besides its intrinsic peculiarity, such a view would fail to vindicate the autonomy of morality because it would suggest that it could be fully rational to believe a moral proposition on theoretical grounds.[15]

Suppose then that there are distinctive practical and theoretical norms that apply to the evaluation of beliefs with moral content, and that the practical norms have lexical priority over the theoretical

14. Think here of iterations of a science fiction case in which a billionaire with a brain scanner will reward you if you believe such-and-such.

15. Nelkin (2000) considers a variant of this approach that suggests that we can rationally believe obviously contradictory propositions, provided that we do so "from different standpoints". Nelkin persuasively emphasizes the difficulties that follow from such a proposal. I set this view aside, noting only that the upshot of such a view for metaethics would be broadly similar to that of the two norms view discussed in the text.

norms. This account can grant that Parfit's arguments might give us good *epistemic* reasons to believe in a thin account of personal identity and in the alleged consequentialist normative implications of this account. However, the practical necessity of accepting a stronger account of identity would nevertheless require us to believe that the stronger account is correct. This proposal would not vindicate the autonomy thesis: for all it says, Parfit's arguments epistemically justify normative ethical theses. However, it does entail something recognizably similar to the autonomy thesis: that it is never correct to form a moral belief on theoretical grounds, where practical considerations count against it.

This proposal has interesting implications for normative and metaethical thinking. On the one hand, it is clearly consistent with the existence of true metaethical and normative theories, each discoverable by theoretical reasoning. On the other hand, it raises the perhaps disturbing possibility that we might be practically required to disbelieve the true moral theory, or to refrain from theoretical inquiry into morality.[16]

The two attitudes proposal differs from the two norms account in virtue of the psychological state that it claims to be rationalized by practical norms. According to this account, while epistemic norms rationalize belief, practical norms rationalize a distinct state of practical commitment towards a proposition. This idea again cries out for more detailed treatment. However, this sketch suffices to suggest the implications relevant to this paper. Like the two norms account, the two attitudes account can grant the *theoretical* force of arguments like Parfit's. However, it claims that, in virtue of conflicting with practically necessary suppositions, these arguments tell us nothing about the distinctive practical commitments that we ought to accept.

This proposal is thus consistent with the existence of true moral theories discoverable by theoretical reason, and hence again does not vindicate the autonomy thesis. However, it would vindicate the autonomy of our moral *commitments* from theoretical reasoning. This is because it makes theoretical reasoning the wrong kind of inquiry to determine the content of the morality that one has practical grounds to accept. If ordinary moral talk expresses our practical commitments, then this proposal suggests a species of fictionalism about our moral discourse, coexisting uneasily with unexpressed moral beliefs. For example, it suggests that one might reasonably believe an action to be permissible, while being practically committed to its being wrong.

Despite vindicating a kind of autonomy for moral discourse, this account again does not impugn metaethics. Like the two norms account, it is consistent with theoretical inquiry into ethics. Further, developing this sort of view would in part involve metaethical inquiry. Consider three central elements of an adequate development of such a view. First, such development would require a properly developed defense of the alleged deep contrast between the practical and theoretical perspectives. Second, it would require a plausible account of the psychological state of practical commitment. And finally, it would require a fictionalist semantics that explained how we express such a state, rather than our beliefs about morality, in our ordinary moral utterances. The second and third of these tasks are recognizably metaethical projects.

We have seen that, inasmuch as we are seeking a deep vindication of the autonomy of morality, the two norms and two attitudes proposals are vulnerable to an important objection. While each of these accounts makes space for a domain of autonomous moral reasoning, neither proposal fully vindicates the autonomy of morality. This is because each is compatible with there being a true normative ethical theory discoverable by theoretical reasoning.[17]

16. The two norms account will also seem objectionable to some in virtue of presupposing a controversial view concerning the ethics of belief: that there could be practical requirements not just on the actions that lead to beliefs but on beliefs themselves. This has seemed wrong to some, in virtue of the thought that believing is involuntary, and that practical requirements can apply only to voluntary action.

17. Theoretical inquiry into ethics would, however, lose a standard motivation on these accounts: to guide action. This is because it would be silent concerning the norms that we are practically required to accept. We might even hesitate to call investigation shorn of this practical point "ethical inquiry", despite its being aimed at the identification of true moral propositions.

The last way to develop the contrast between practical and theoretical reason that I consider avoids this objection. This account claims that practical and theoretical reasoning involve distinctive types of content. I thus call this the "two contents" account. Such an account suggests that nonderivative theoretical reasoning about morality would involve a category mistake. This in turn ensures that there cannot be a true moral theory discoverable by purely theoretical inquiry. The two contents account thus promises to more fully vindicate the autonomy of morality than the other proposals sketched. Because of this, I consider the two contents account to be the most interesting development of the contrast between practical and theoretical reason that grounds the deliberative requirement. Consequently, I consider it in greater detail than the others.

The most pressing question facing such an account is what exactly the notion of "distinctively practical content" could come to. Here, I pursue a helpful suggestion offered by Korsgaard.[18] She claims that practical concepts do not function "to describe reality" (2003, 105). Rather, the essential function of a practical concept is to "refer to the solution to a [practical] problem" (2003, 115).[19]

Korsgaard offers as an example of this model a reading of Kant's derivation of the Universal Law formulation of the categorical imperative. The practical problem in this case is that the will must adopt a principle for itself on pain of heteronomy. Further, that principle must be a law, because (according to Kant) for the will to be a cause, it must conform to a law. Korsgaard claims that the Universal Law formulation of the categorical imperative ("Act only on a maxim that can be willed to be a universal law") can be extracted from these constraints. The lesson is that, since these constraints are simply descriptions of the details of the relevant practical problem, we are in a position to extract normative principles from the statement of a practical problem alone (2003, 113–5). Korsgaard concludes: "The categorical imperative is a principle of the *logic* of practical deliberation, a principle that is constitutive of deliberation, not a theoretical premise that is applied in practical thought" (2003, 115).

Korsgaard suggests that this sort of "practical logic" constitutes the correct moral methodology: in moral inquiry we attempt to discover moral principles, exactly by understanding these inescapable practical problems and the structure of their solutions (2003, 115–6).[20] She claims that this account of practical concepts also provides truth-conditions for propositions with practical content. On this account, normative principles, understood as solutions to practical problems, function as guides to the correct application of the concepts that they name. Korsgaard thus suggests that a proposition with practical content is true if (*inter alia*) it conforms to the application conditions of the relevant practical concept (2003, 117).

This account of practical concepts fleshes out the two contents account by suggesting what distinctively practical conceptual content could be. The practical logic methodology that follows from it also explains how the two contents account would vindicate the autonomy of morality. This is because the practical logic conception of moral methodology implies that theoretical premises are simply the wrong sorts of things to enter into normative ethical theorizing.[21] Thus, if the

18. It is not entirely clear whether Korsgaard consistently endorses what I call the two contents account. For example, Nelkin suggests reading Korsgaard as adopting the "two standpoints" proposal mentioned in the previous note (2000, 567–8). However, Nelkin expresses uncertainty regarding Korsgaard's ultimate view. I intend here to apply Korsgaard's account of practical content, while making no claim to have identified her considered opinion on these topics.

19. Compare also her claim that morality requires "a different conceptual organization of the world" from that required by the theoretical perspective (1989a, 37).

20. This conception of practical methodology is, of course, controversial. One central question is how much can be extracted from the logic of our practical concepts. Gibbard (1999), for example, is pessimistic.

21. This requires that we think of the output of moral theorizing in terms of general principles or maxims, just as my gloss of the autonomy thesis requires. However, *pace* Korsgaard's claim mentioned above, it does seem that these maxims would constitute a kind of moral knowledge that we would need to apply in action in light of our beliefs about the relevant features of our options (*e.g.* the belief that Φ-ing would cause injury, or that Ψ-ing would be the telling of a lie).

two contents account were correct, the defender of the practicality of morality would be in a position to say something much stronger in response to Parfit than was permitted by the other accounts. On this account she could convict Parfit of making a fundamental mistake, because moral concepts and concepts of personal identity are practical, and hence not appropriate objects for the sort of theoretical reasoning that he attempts to apply to understand them.

The relation between the two contents account and metaethics is less obvious than that of the previous accounts. One might think that the proposed distinctively practical character of moral concepts was incompatible with metaethics. The worry is that, since metaethics is a theoretical project, the metaethicist inevitably makes a category mistake on this account, like a chess player deciding that making himself a coffee would be the best move. However, this diagnosis would fail to appreciate the scope and flexibility of theoretical reasoning. I will argue that Korsgaard's account of practical concepts can rather be best understood as entailing a species of response-dependent metaethical realism.

To begin, consider the following schema for a response-dependent account of a property *F*:

> RD *X*s are *F* just in case and because *Y*s respond *R*-ly to *X*s in conditions *C*.[22]

This is a schema for response-dependent accounts of a property, because it is some response of *Y*s that metaphysically determines that an *X* is *F*. For example, a crude response-dependent theory of humor would claim that for a joke to have the property of being funny just is for people to tend to respond with amusement upon hearing the joke in ordinary conditions.

It is possible to use the Korsgaardian account of practical concepts to construct a response-dependent metaethics by appealing to a kind of semantic deference. Just as one can defer semantically to a person or a group of experts, I propose that a theoretical reasoner would be able to defer to the practical perspective. This, I suggest, would allow such a reasoner to construct a response-dependent metaethical account of moral properties. Thus, take the terms 'agent' and 'agency' to be used deferentially, to pick out the meanings of these terms given by the logic internal to the practical perspective. These terms can be used from within the theoretical perspective to construct an instance of the response-dependence schema for permissibility, *practical response-dependence*:

> PRD Options are permissible just in case and because agents would judge them to conform to norms that they accept, when they are informed about, and responsive to, the conditions necessary to fully maintain their agency.

This account grants that moral concepts can be nonderivatively grasped only from within the practical perspective.

According to practical response-dependence, facts derived from the practical perspective are necessary and sufficient to fix whether an option is permissible for an agent. However, practical response-dependence is nonetheless an identifiably metaethical view.[23] This is because it offers a non-trivially explanatory account of the property of permissibility, graspable by a purely theoretical reasoner.

If this reading of the metaethical upshot of the Korsgaardian account of practical concepts is correct, then the two contents proposal, like the two attitudes proposal, turns out to entail metaethical commitments.[24] However, unlike the two attitudes proposal, this account does

22. Johnston introduced the term 'response-dependence' in his (1989). I set aside a variety of important issues here. Most importantly, I assume the ontological respectability of response-dependent properties. Readers attracted to a more austere ontology should translate what follows into a defense of the claim that the two contents account permits (derivative) "theoretical" moral concepts.

23. Practical response-dependence is also obviously incomplete; a full response-dependent metaphysics of morality would have to account for the plurality of normative properties. However, this complication is orthogonal to whether the most promising version of the two contents account entails a response-dependent metaethics.

24. To say that the two contents account entails a response-dependent metaethics

not permit conflict between the evaluations suggested by practical and theoretical reason. This is because of the deferential nature of the theoretical reasoning about morality permitted by this account.

In this section, I have drawn on Korsgaard's description of the basic tension between Parfit's argument and the deliberative requirement, and I have used some of her discussion of practical concepts to develop what I take to be the most promising version of that requirement. However, in the Introduction I noted that Korsgaard took a correct understanding of morality to entail that errors infect standard metaethical practice. Whatever the details of her view, Korsgaard is clearly sympathetic to some variation of the view that there is a deep contrast between theoretical and practical reason. Having explored three ways of developing this contrast, I now examine two striking implications that she takes it to have for metaethics.

First, Korsgaard claims that metaethical theorizing is distorted by the presupposition that all "authentically cognitive" concepts function to describe the world (2003, 105). In other words, most metaethicists do not see the possibility of the sort of practical concepts that she suggests. This is arguably true. However, the practical response-dependence account that I suggest shows that this account of practical concepts is not incompatible with metaethics, but rather suggests a determinate metaethics. So understood, Korsgaard's first criticism is best understood as a (quite deep) methodological criticism within metaethics, rather than a diagnosis of an error inevitably implicit in the domain as a whole.

Second, Korsgaard suggests that once her account is understood,

> [W]e will not be inclined to think that there is a difference between doing 'meta-ethics' and doing 'normative' or practical ethics. The attempt to specify the meaning and reference of an ethical concept will point fairly directly to practical ramifications. [2003, 121 n. 44]

The practical logic account of methodology that she suggests does entail that a correct understanding of practical concepts will entail determinate normative implications. However, this does not threaten the distinction between metaethics and normative ethics, as Korsgaard may appear to suggest. The lesson of the discussion in section 2 of metaphysical utilitarianism and simple act utilitarianism was precisely the need to distinguish metaethical and normative theories, even if there are entailments between them. Just as in that case, it is conceivable that a metaethics that made no mention of distinctively practical concepts could entail an account of the relation between actions, motives, and rightness structurally identical to the relations that Korsgaard claims to be entailed by the practical logic. We would want to say that such a metaethics was distinct from the Korsgaardian account, but entailed the same normative theory.[25]

In this section, I have explored three ways of developing the "deliberative requirement" — the claim that moral theorizing must take place from within the perspective of first-person practical deliberation — with the idea that such strategies might vindicate the autonomy

is not to imply that in developing the former, one would necessarily be aiming to solve metaethical problems. Rather, the point is simply that such an account permits a kind of derivative theoretical reasoning about morality. By engaging in such reasoning, I claim, one could determine that the practical response-dependence metaethics follows from the two contents account.

25. I have emphasized that I am using some of Korsgaard's helpful suggestions to elucidate the two contents account, without claiming it to be her considered view. One of her claims about metaethics is both inconsistent with the two concepts account and independently puzzling.

She claims that both metaethical expressivism and realism are *true*, albeit in a "way that makes [them] boring" (2003, 118; 122 n. 49). In elucidating this point, she again appeals to the contrast between the practical and theoretical perspectives. Expressivism is true from the theoretical perspective, she suggests, while realism is true from within the practical perspective. Perhaps this is best read as a rhetorical overstatement. Note first that this appears inconsistent with Korsgaard's own discussion of practical concepts. For if, as she repeatedly insists, practical concepts do not function to "describe reality" (2003, 105), then (as I have suggested above) correct theoretical reasoning about the practical cannot conflict with the practical perspective. Note second that, as truth is arguably essentially perspective-independent, Korsgaard's claim is extremely unappealing, because it makes truth itself perspective-relative. Helpful further discussion of this claim can be found in Hussain and Shah (unpublished manuscript, 33ff.).

of morality, and suggest a resolution to the apparent tension between the autonomy thesis and metaethics. I argued that each of the three accounts that I examined does vindicate some approximation of the autonomy of morality. However, I suggested that only the two contents account is fully satisfying in this respect: the other accounts were consistent with the true moral theory's being discoverable by theoretical reasoning. I also argued that each of these theories was compatible with metaethical theorizing, and suggested that this discussion permitted a clarifying response to Korsgaard's criticisms of metaethics.

4. Resolving the tension

In the Introduction, I explained an apparent tension between metaethics, as it is ordinarily conceived, and the notion that morality is autonomous. In sections 2 and 3, I have examined two relatively radical strategies for vindicating the autonomy of morality. In this section, I suggest a different strategy for resolving this tension. Recall that two of the accounts examined in section 3 involved or entailed particular metaethical theories. This suggests a general insight that can be applied to resolve the apparent tension between metaethics and the autonomy of morality: rather than being necessarily inconsistent with metaethical research, the autonomy of morality could be vindicated in part by a particular metaethical theory.

The fictionalist and practical response-dependence metaethics introduced in section 3 can be understood as instances of this phenomenon. Thus, the two attitudes and two contents proposals were introduced as means of vindicating the autonomy thesis. However, I suggested that these proposals each involved commitment to views in metaethics. Some metaethics would vindicate the autonomy of morality more directly. A metaethics could vindicate the autonomy of morality by characterizing the semantics or ontology of morality in such a way as to secure the irrelevance of non-moral theses to the justification of moral theories. Consider two examples of how a metaethics might achieve this.

First, consider the thesis that the moral facts are determined by the outcome of idealized autonomous normative ethical inquiry. Such a view could be developed into a determinate metaethical theory.[26] However, by making moral facts parasitic on what is by stipulation autonomous moral inquiry, this account would secure the autonomy thesis. Second, while the discussion of section 2 suggests that it is a mistake to make neutrality between normative theories a constraint on counting as a metaethical theory, a metaethics could secure the autonomy of morality *by* being neutral between normative theories. Some contemporary sensibility theories and expressivisms are perhaps best understood as being neutral in this way.[27]

It might be objected that a metaethics cannot vindicate the autonomy of morality, because the justification of such a metaethics will itself involve appeal to non-moral considerations. As I emphasized in the Introduction, some of the central desiderata on metaethical theories appear to concern the fit between morality and the other elements of our best accounts of the world. The objector insists that, far from vindicating the autonomy of morality, the metaethical theories suggested would thus be inconsistent with it.[28]

The plausibility of the apparent tension between metaethics and the autonomy of morality rests on exactly the confusion that motivates this objection. An analogy will help to illustrate why the objection is misguided. Suppose that there is a correct epistemic rule of inference *Double Negation*:

26. Compare Sayre-McCord, who suggests that moral kinds are those groupings deemed significant by "the best moral theory" (1997a, 315; cf. also 1997b).

27. Thus the sensibility theory sketched in Dreier (2002) is constructed to be scrupulously neutral between normative theories. Gibbard's norm-expressivism also plausibly vindicates the autonomy thesis. In his discussion of moral inquiry, Gibbard identifies characteristic psychological pressures that a self-conscious norm-expressivist will be under in her normative thinking (1990, 274ff). On Gibbard's account, the appropriate response to those pressures could only be settled autonomously, by whether one took them to be normatively salient.

28. This line of objection suggests that the autonomy of morality can only be defended as a basic commitment. Such a view might explain Dworkin's otherwise puzzling disinterest in justifying the neutrality constraint.

DN If you believe that not-not-*P* then you ought to believe that *P*.[29]

Double Negation entails that when one believes that not-not-*P*, one should believe that P. However, otherwise reasonable but very unlucky inquirers might acquire misleading evidence that Double Negation is not a correct rule of inference. And that might lead them astray, when they think about what to conclude given not-not-*P*. However, if Double Negation is a correct rule, it just settles what they should conclude; the justification of Double Negation is itself a separate matter.

This contrast between the content of a rule and its justification also holds for the autonomy of morality. Recall that the autonomy of morality is the thesis that non-moral premises are irrelevant to the justification of normative ethical theories. Now, suppose that a metaethical theory *M* is true, and that *M* entails that non-moral premises are irrelevant to the justification of normative ethical theories. Then the truth of *M* just settles that morality is autonomous. The objector balks because the justification of *M* involves non-moral premises. This can only appear to be an objection if one takes the autonomy thesis to apply not just to the justification of normative ethical theories, but also to *its own* justification. However, such self-application is no part of the content of the autonomy thesis.

Nor could a plausible autonomy thesis ensure that it is always reasonable to believe that morality is autonomous. Compare again Double Negation: it might be unreasonable for our imagined unlucky inquirers to accept or apply Double Negation, given their misleading higher-order evidence. However, this does not entail that Double Negation is not a correct epistemic norm. An autonomy thesis that attempted to preclude the possibility of such misleading evidence would simply be obviously false. The objection thus fails, and with it the primary motivation for resisting the simple resolution of the apparent tension between metaethics and the autonomy of morality. Thus, it

29. This supposition is implausibly simple. At minimum, a double negation elimination rule of inference likely needs to read: "If you believe that not-not-p then you ought either to believe that p or cease to believe that not-not-p."

is plausible that the autonomy of morality depends for its truth on metaethical fact.

It is important to distinguish this claim about explanatory dependence from the epistemological relationship between the autonomy thesis and metaethical theorizing. Thus, suppose that one finds the autonomy of morality intuitively plausible. On some epistemologies, such plausibility would entail that compatibility with the autonomy of morality would contribute to justifying the acceptance of a metaethical theory. Vindicating the autonomy thesis might thus serve as a desideratum on one's metaethical theorizing.[30]

In section 3, I examined one strategy for vindicating the autonomy thesis: developing an account of the deliberative requirement on normative theories. This section suggests another strategy: defending a metaethics that directly vindicates the autonomy thesis. There is an important contrast between these strategies. While I have argued in this section that the autonomy of morality is not in tension with metaethics per se, it is in tension with a variety of metaethical theories. This suggests that unless an account of the deliberative requirement such as those examined in section 3 can be successfully defended, the defensibility of the autonomy thesis rests on the metaethical case that can be made for autonomy-vindicating metaethics, such as those mentioned above.

Conclusion

I began this paper began by bringing out an apparent tension between metaethics and the autonomy of morality. I then considered two radical strategies for resolving this tension and securing the autonomy of morality. I suggested that these strategies entailed a variety of significant consequences for metaethics. Thus, some ways of securing the autonomy of morality would, if correct, suggest that contemporary

30. One possibility is that such a desideratum might explain an important part of what is worth salvaging from Moore's notorious "open question argument". This possibility was suggested to me by an anonymous reader for *Philosophers' Imprint*.

metaethical theorizing typically involves a fundamental error. I also argued that some of these strategies were independently unattractive. In the final section, I suggested that the apparent tension between metaethics and the autonomy of morality can be dissolved: far from putting metaethics into question, the autonomy of morality may most simply be secured by defending a metaethical theory that explains why non-moral theses are irrelevant to normative ethical theorizing. While the approaches discussed in sections 2 and 3 at least initially suggested a diminished role for metaethical theorizing, the resolution proposed here suggests that metaethical theorizing has an indispensible role to play in the evaluation of the autonomy thesis.[31]

31. I am grateful to many people for helpful comments and discussion, especially Michael Smith, Tom Kelly, Gil Harman, Philip Pettit, Nadeem Hussain, Sarah Buss, Mark Schroeder, Brendan Ritchie, Amy Shuster, John Maier, Nick Stang, Rochelle Edinburg, Ryan Davis, Sam Arnold, and Sari Kisilevsky. I am especially indebted to two anonymous readers for *Philosophers' Imprint*, who each made many extremely useful suggestions. I would also like to thank Princeton's University Center for Human Values for a fellowship that helped to support my work on this paper.

Works Cited

Darwall, Stephen, Allan Gibbard, and Peter Railton. "Towards *Fin de Siècle* Ethics: Some Trends". *Moral Discourse and Practice*. Ed. Stephen Darwall, Allan Gibbard, and Peter Railton. New York: Oxford University Press, 1997. 3–47.

Dreier, James. "Metaethics and Normative Commitment". *Philosophical Issues* 12 (2002): *Realism and Relativism*. Ed. Ernest Sosa and Enrique Villanueva. 241–63.

Dworkin, Ronald. "Objectivity and Truth: You'd Better Believe It". *Philosophy & Public Affairs* 25.2 (1996): 87–139.

Dworkin, Ronald. "Replies to Critics: Zangwill". *Brown Electronic Article Review Service*. Ed. Jamie Dreier and David Estlund. Posted 4/9/1997. <http://www.brown.edu/Departments/Philosophy/bears/homepage.html>

Frankena, William K. "Main Trends in Recent Philosophy: Moral Philosophy at Mid-Century". *Philosophical Review* 69.1 (Jan. 1951): 44–55.

Gibbard, Allan. "Morality as Consistency in Living: Korsgaard's Kantian Lectures". *Ethics* 110 (October 1999): 140–64.

Gibbard, Allan. *Wise Choices, Apt Feelings*. Cambridge, MA: Harvard University Press, 1990.

Harman, Gilbert. *Explaining Value*. Oxford: Oxford University Press, 2000.

Hussain, Nadeem, and Nishi Shah. "Metaethics and Its Discontents: A Case Study in Korsgaard," unpublished manuscript.

Johnston, Mark. "Dispositional Theories of Value". *Proceedings of the Aristotelian Society* Supp. Vol. LXIII (1989): 139–74.

Kagan, Shelly. *Normative Ethics*. Boulder, CO: Westview Press, 1998.

Korsgaard, Christine. "Morality as Freedom". *Kant's Practical Philosophy Reconsidered*. Ed. Yirmiyahu Yovel. Boston: Kluwer, 1989a. 23–48.

Korsgaard, Christine. "Personal Identity and the Unity of Agency: A Kantian Response to Parfit". *Philosophy and Public Affairs* 18.2 (Spring 1989b): 101–32.

Korsgaard, Christine. "Realism and Constructivism in Twentieth-

Century Moral Philosophy". *Philosophy in America at the Turn of the Century*. APA Centennial Supp. to *Journal of Philosophical Research* (2003): 99–122.

Miller, Alexander. *Introduction to Contemporary Metaethics*. Cambridge: Polity, 2003.

Moore, G.E. *Principia Ethica*. Cambridge: Cambridge University Press, 1903.

Nagel, Thomas. *The Possibility of Altruism*. Princeton: Princeton University Press, 1970.

Nelkin, Dana. "Two Standpoints and the Belief in Freedom". *Journal of Philosophy* 97.10 (2000): 564–76.

Parfit, Derek. *Reasons and Persons*. Oxford: Clarendon Press, 1984.

Sayre-McCord, Geoff. "Different Kinds of Kind Terms". *Philosophical Issues* 8 (1997a): 313–23.

Sayre-McCord, Geoff. "'Good' on Twin Earth". *Philosophical Issues* 8 (1997b): 267–92.

Sinnott-Armstrong, Walter. *Moral Skepticisms*. New York: Oxford University Press, 2006.

Zimmerman, David. "Meta-Ethics Naturalized". *Canadian Journal of Philosophy* 10.4 (December 1980): 637–62.