

Unifying Moral Methodology

Note: this is the author's penultimate manuscript.

The official version will be published in
Pacific Philosophical Quarterly 93 (2012): 523–549.

Tristram McPherson
Virginia Tech
dr.tristram@gmail.com

Abstract:

This paper argues that the best way to pursue systematic normative ethical theorizing involves metaethical enquiry. My argument builds upon two central claims. First, I argue that plausible metaethical accounts can have implications that can help to resolve the methodological controversies facing normative ethics. Second, I argue that metaethical research is at least roughly as well-supported as normative ethical research. I conclude by examining the implications of my thesis. *Inter alia*, it shows that the common practice of engaging in systematic normative theorizing uninformed by metaethical commitments offers a significant and troubling hostage to metaethical fortune.

Introduction

Consider two central projects in philosophical ethics. The core project of *normative ethics* is to systematically explain what we ought to do and why we ought to do it. The core project of *metaethics* is to explain how ethical judgment and discourse, and the ethical facts (if any), fit within our best general account of the world. The striking difference between these two projects might appear to suggest that they should be pursued separately. In this paper, I argue that this appearance is mistaken. Rather, despite the differences between these projects, the best method for normative ethical theorizing involves engaging in metaethical theorizing. I call this claim (made more precise in §1 below) the *unity thesis*.

Intuitive motivation for the unity thesis can be given by example. First, suppose that the correct metaethic would state that there are ethical facts, and explain the nature of those facts. This would prompt us to consider which methods of normative ethical investigation are best suited to discovering facts with this nature. Second, suppose that the correct metaethic would tell us that there are no ethical facts. This would raise questions about what exactly we are doing when we argue for normative ethical theories, and how we might best go about this task. Either way, metaethical theories promise to have deep significance for the methods we should employ in normative ethics.

Despite this intuitive motivation, the unity thesis is a revisionary claim in contemporary philosophical ethics, both in methodological theory and in practice. John Rawls' method of reflective equilibrium is the most widely accepted methodology in ethics. It is thus striking that part of Rawls' stated ambition for the method was to underwrite the methodological independence of normative ethical enquiry. Rawls' ambition has been widely honored in practice: normative ethicists often develop and defend their views without engaging in significant metaethical enquiry. If correct, the unity thesis raises serious questions about the wisdom of this practice.

In this paper I sketch the outline of a powerful argument for the unity thesis (§2). I then defend the premises of this argument in more detail, thereby addressing central challenges to this thesis (§§ 3-6). Finally, I reflect on the significance of the unity thesis for philosophical work in ethics (§7). I begin by clarifying both the central question of the paper and the unity thesis.

1. The unity thesis explained

Precisely stated, the unity thesis I seek to defend is this:

Unity The best way to pursue the project of normative ethical theorizing involves integrating that theorizing with metaethical theorizing.

The question addressed by the unity thesis is unfamiliar, and this formulation requires significant unpacking. I thus beg the reader's patience in this section as I clarify the central question of the paper and the answer provided by the unity thesis. I first briefly explain how I understand each of the central concepts invoked in this thesis: *normative ethics*, *metaethics*, *optimality*, and *integration*. I then use these materials to distinguish the unity thesis from other claims with which it can easily be confused.

First consider normative ethics. Normative ethicists seek to construct and justify the correct normative ethical theory. This theory would consist in a set of general principles that tell us which actions, states, character traits (etc.), fall under normative categories like rightness, obligatoriness, goodness, virtuousness (etc.), and why.¹ Consider an example: a simple version of utilitarianism consists in the claim that an action is right just in case and because, among the options available to the agent, it maximizes happiness. This theory purports to give an exhaustive and explanatory characterization of the rightness of actions in terms of a very simple explanatory ground: happiness maximization. In this paper, I treat the scope of normative ethics as restricted to relatively general theories.² Utilitarianism, contractualism, and virtue theory, for example, fall within the scope of the thesis. By contrast, theories with restricted application – such as theories of consent or of just war – fall outside of its scope.

Next consider metaethics. Metaethicists seek to understand how our practices of ethical judgment and discourse, and the ethical facts (if any), fit within our best general account of the world.³ The first part of this project is the central task of moral psychology and moral (meta-)semantics. Metaethicists disagree over whether the correct moral psychology and moral semantics commit us to the existence of moral facts or properties. Those who believe that it does face an additional metaethical task: the metaphysical task of determining what such properties would be like, and if any are actually instantiated.⁴

I assume here that the projects of metaethics and normative ethics, as just sketched, are conceptually distinct from each other.⁵ The unity thesis argues that despite

this conceptual difference, it is *best* to pursue these two projects together. I intend this claim of optimality to be evaluated relative to the constitutive goal of normative ethical theorizing: identifying and justifying the correct normative theory. This goal contrasts with more modest theoretical goals that particular normative ethicists might have. For example, some ethicists may aim in their work simply to make some contribution to our thinking about normative ethics (more obviously irrelevant are goals such as *getting tenure*, or *defending the indefensible*).

One might worry that the goal of identifying and justifying the correct normative theory is too ambitious to take seriously. In reply, it may be helpful to compare the goals of perfecting one's golf swing or tennis serve. As in these cases, normative ethicists can and should pursue their goal by seeking to approximate it to the best of their abilities, given their available resources. This suggests that we can imagine relative success in pursuit of this goal to be constituted by one's more closely approximating the correct theory. For example, applied to physics this goal suggests that Feynman was more successful than Newton, who was more successful than Aristotle. This is so even if Newton's *contribution to the discipline* is intuitively greater than Feynman's.

I also intend the goal to be understood as applying to generic philosophical enquirers. For example, suppose that Hermes was deeply humiliated early in life by a metaethicist, and finds himself unable ever to think about metaethics without flying into a rage. Doing metaethics will evidently not be useful to *his* work in normative ethics. My argument is generic in the sense that I do not intend to address such biographical quirks.

The unity thesis states that the best way to pursue the goal of normative theorizing involves *integrating* our normative theorizing with systematic metaethical enquiry. In short, this involves pursuing the answers to metaethical and normative questions together. In slightly artificial detail, this can be understood as a three-stage process. The first stage involves sketching available arguments and plausible candidate answers to the central questions in both metaethics and normative ethics. On one conception of evidence, this stage involves acquiring and synthesizing a body of evidence within each domain. The second stage involves identifying the entailments and epistemic implications between well-justified metaethical theses and candidate normative theories. Finally, the third stage of integration involves identifying the normative theories that are best justified in light of their

relation to the well-supported theses in both domains (and the global desiderata on theory choice).

These glosses put us in a position to distinguish the unity thesis from stronger and weaker methodological claims in the vicinity. Most importantly, the unity thesis is *not* the strong claim that it is best to first do metaethics, and only then do normative ethics. The integrated approach to normative theorizing requires bringing our best justified claims in both domains to bear on the justification of systematic normative theories. This does not imply that work in metaethics has some sort of *priority* in that process.⁶ In fact, I am sympathetic with the symmetrical view that the best way to do metaethics may involve integration with systematic normative theorizing. However, the unity thesis itself is *silent* concerning this question about the methodology of metaethics: it is a claim about the best way to do normative ethics, and makes no claim about how to do metaethics. (I explain this silence at the end of §6).

At the other extreme, the unity thesis is not merely the weak claim that metaethical research could in principle be relevant to normative ethics.⁷ This misunderstands the claim of optimality made by the thesis. There are many sorts of work that could in principle be relevant to the justification of normative ethical theories. The unity thesis claims that, unlike many of these, *metaethical* work is related to work in normative ethics in a way that makes integrating metaethical work into our systematic normative enquiry part of the best methodology of normative ethics.

The significance of the unity thesis is that it condemns other approaches to normative ethics as inferior. Most importantly, this includes the attempt to construct and justify normative ethical theories *without* investigating which metaethical views are well-justified, or what implications those metaethical views might have for normative ethics.

2. An outline of the case for the unity thesis

Having clarified the unity thesis, I now introduce my argument for it. I suggested in the introduction that it is plausible that progress in metaethics could enable progress in normative ethics. I begin by explaining the mechanism that makes this possible, before building an argument for the unity thesis on the foundation of that mechanism.

My case for the unity thesis rests upon the distinctive relations that the projects of metaethics and normative ethics each bear to moral epistemology. The normative ethicist's distinctive interest in moral epistemology is methodological: an informative moral epistemology will tend to have implications for how best to construct and justify normative ethical theories. By contrast, the metaethicist's distinctive interest in moral epistemology arises because metaethical theories tend to have implications *for* moral epistemology. Metaethicists thus face an instance of what Christopher Peacocke (1999) has dubbed the 'integration challenge': they need to ensure that their views in moral semantics or metaphysics do not imply that moral knowledge is impossible, implausibly easy, or bizarrely acquired. These contrasting relations to moral epistemology provide the core *mechanism* that I take to support the unity thesis: well-supported metaethical theories might have significant implications for moral epistemology, which might in turn have implications for how the normative ethicist should evaluate candidate normative theories.

Two premises need to be true for this mechanism to have methodological bite:

1. There are live methodological questions about the justification of normative theories.
2. Metaethical research can help to resolve those methodological questions.

If premise 1 were false, then the normative ethicist's methodological project would be close to complete, and the mechanism would be a mere curiosity. If premise 2 were false, then metaethical research would not be methodologically useful to the normative ethicist, and again the mechanism would fail. Consider the case for these premises in turn.

A cursory exploration of the relevant literature provides striking support for the first premise: the methodological controversy facing normative ethics is both deep and broad. Perhaps the most familiar approach to normative theorizing involves the appeal to intuitions about possible cases.⁸ However, there is vigorous disagreement regarding the role that such intuitions should play in our theorizing. Is it unavoidable?⁹ To be avoided at all costs?¹⁰ Should we be careful to appeal only to real or realistic cases? Or can we appeal with abandon to scenarios involving evil demons, brains in vats, or credulity-straining causal chains?¹¹ Are appeals to cases distinctively more or less problematic than appeals to the plausibility of principles, or to the intuitive moral significance of certain kinds of properties, or appeals to the intuitive importance of certain broad theoretical desiderata?¹² Finally, can we augment or even supplant the appeal to intuition by

producing empirical results that bear on normative ethical theses, as a growing number of researchers claim?¹³

The second premise of my argument claims that methodological controversies like those just identified can be addressed at least in part by thinking about the epistemic implications of our best metaethical accounts. It is important to distinguish a weaker and a stronger reading of this claim. On the weaker reading, the claim is that our best metaethical accounts would, *if true*, have implications for these methodological controversies. The stronger reading would include the additional claim that metaethical theories that do this are true or at least well-justified. Premise 2 of my argument intends *only* the weaker reading. The stronger reading includes claims that will be taken up by later parts of the argument.

So understood, this premise can be motivated by noting the contrasting methodological implications of three leading contemporary metaethics. First, Russ Shafer-Landau argues for a very close link between his non-naturalist metaethic and a broadly rationalist epistemology, according to which certain fundamental moral principles are (synthetic) a priori (e.g. at 2003, 247). By contrast, Richard Boyd argues that his version of naturalist realism entails an empiricist, a posteriori moral epistemology. Indeed, he plausibly suggests that the methodology made appropriate by his theory is approximately that of the human sciences (1997, 122-4). Finally, Allan Gibbard claims that adopting his expressivist metaethic will typically suggest two central sorts of pressure on one's methodological commitments.¹⁴ The first is a sort of intuitionistic pressure towards taking one's existing normative judgments to be presumptively legitimate.¹⁵ In striking contrast to Shafer-Landau, however, Gibbard claims that this pressure towards intuitionism must compete with a pragmatic pressure that tends to discredit any set of evaluative attitudes that make things go badly for their bearer and those she cares about (1990, 282-4). Gibbard has lately sought to flesh out the distinctive consequences of the intersection of these two pressures, concluding that they support (but do not entail) an approach to normative theorizing that itself favors a utilitarian normative ethic (2008).

The premises just defended suggest an important role for the mechanism sketched at the beginning of this section. However, two further premises are required to complete a valid argument for the unity thesis.

3. The best way to pursue the project of normative theorizing involves integrating relatively epistemically well-supported research that could help to resolve methodological questions about the justification of normative theories.
4. Metaethical research is epistemically well-supported relative to normative ethical research.

Premise 3 is my central methodological thesis, and the necessity of defending premise 4 follows from that thesis. Consider these premises in turn.

The third premise introduces a new idea: relative epistemic support. A domain of research is *epistemically well-supported* relative to normative ethics just in case the best systematically ambitious research in normative ethics tends not to be significantly better justified than the best research in this domain. (Notice that a domain that is equally well-justified as normative ethics will count as well-supported by this standard.)

The third premise claims that if research is well-supported in this sense and it could help resolve methodological questions in normative ethics (as premise 2 argued that metaethics can), then the best method for normative theorizing involves incorporating it. The idea that relative epistemic support is both necessary and sufficient in this context can be motivated by example.

On the one hand, premise 3 states that research *must* be comparatively well-supported in order to play this role. Consider an example where this condition fails. If my Vaguely Intuitive Metaphysic (VIM) is inconsistent with the truth of quantum mechanics, it thereby bears on the justification of quantum mechanics, in one clear sense: a sense analogous to the intended ('weak') reading of premise 2 of my argument. However, if we suppose that VIM is significantly less well epistemically supported than the central theses of quantum mechanics, there is little to be said for physicists taking time out from their usual work to explore the intuitive case for VIM relative to its metaphysical competitors.

On the other hand, relative epistemic support also appears sufficient to make integration optimal. Suppose that Clio is seeking to assess a normative principle, and has assembled two conflicting bodies of putative evidence, E1 and E2, drawn from different sources, that are apparently relevant to the justification of the principle. Suppose next that Clio possesses relatively good (but not decisive) evidence for a metaethical theory. Suppose finally that this metaethic entails that only the source of E1, and not that of E2, is likely to provide reliable information about the ethical facts. Because the metaethic is relatively well-supported, it would surely be less than optimal for Clio to simply disregard the

implications of this metaethic while assessing whether she is justified in accepting the normative principle.

I take these cases to motivate premise 3: if a domain of research is relevant to the target of your enquiry, and well-justified relative to the other considerations that you have to draw upon, the best approach to your enquiry will involve drawing upon that domain of research.¹⁶ Premise 4 addresses the question of whether metaethical research meets this standard of justification.

As in the third premise, the fourth premise concerns relative, and not absolute, epistemic quality.¹⁷ Denial of this premise requires that research in normative ethics is much better supported than research in metaethics. This in turn requires either extraordinary optimism about the epistemic quality of normative ethical research pursued independently of metaethics, or profound pessimism about metaethics. There is a straightforward reason for rejecting both of those grounds for doubt, however. This is that leading metaethical and normative research is conducted within an overwhelmingly shared intellectual context. Indeed, there is more diversity in the methods applied within each of these domains than there is methodological difference between them. This point casts strong *prima facie* doubt on the idea that any contrast in epistemic quality between them is radical.

The premises that I have introduced and briefly defended in this section constitute a powerful argument for the unity thesis, which I will call the *basic unification argument*:

1. There are live methodological questions about the justification of normative theories.
 2. Metaethical research can help to resolve those methodological questions.
 3. The best way to pursue the project of normative theorizing involves integrating relatively epistemically well-supported research that could help to resolve methodological questions about the justification of normative theories.
 4. Metaethical research is epistemically well-supported relative to normative ethical research.
-
- C. The best way to pursue the project of normative theorizing involves integrating that theorizing with metaethical theorizing (the unity thesis).

As I noted in the introduction, the unity thesis has radical implications for the practice of normative ethics. Many philosophers will thus be inclined to resist the basic unification argument. To do so, however, they need to reject one of the premises of this argument. In the next four sections, I examine the prospects for doing so, addressing important

objections to each of the premises of this argument in turn. As I will show, rejecting any of these premises is no easy task.

3. Methodological controversy and reflective equilibrium

The first premise of the basic unification argument states that there are live methodological questions about the justification of normative ethical theories. My initial case for this premise pointed to the depth and breadth of the methodological controversy facing contemporary normative ethics. The most important worry about this case stems from two facts about John Rawls' method of reflective equilibrium. First, Rawls twice presented this method as (at least part of) a way of permitting normative ethics to proceed without metaethical entanglement. This appears to suggest a direct inconsistency between this method and the unity thesis. Second, Rawls' method is endorsed both widely and strongly,¹⁸ seemingly suggesting that the controversy posited in this premise is exaggerated.

Consider first the challenge to the first premise of the basic unification argument. The idea is that the canonical status of Rawls' method is incompatible with the methodological controversy in normative ethics posited by this premise. This apparent tension is illusory. The best explanation for the remarkably frequent appeals that philosophers make to reflective equilibrium is that they often invoke the idea of reflective equilibrium without intending to wed themselves to the *details* of Rawls' method. As Michael Huemer notes, the term 'reflective equilibrium' is now used extremely broadly, to refer to '...any process of weighing conflicting beliefs against each other and renouncing the less plausible beliefs in order to restore coherence in one's belief system' (2008, 369n3).¹⁹ It is hardly surprising that many are happy to claim this idea as their own, as renouncing it would amount to proclaiming indifference to one's own incoherence. Diluted in this way, however, the 'method' as Huemer describes it here is incapable of challenging or adjudicating the methodological controversies that motivate the first premise of the basic unification argument.

Rawls himself endorsed a considerably more informative method. In developing this method, Rawls twice explicitly aimed to insulate normative ethics from metaethical influence (among other things).²⁰ The unity thesis is thus inconsistent with at least some of Rawls' methodological ambitions. Because of this, it is worth explaining why only Rawls'

early version of the method actually insulates normative ethics. His later (and superior) version does not.

Rawls' early version of the method can be divided into three stages. The first is to assemble what he calls our 'considered judgments' about particular moral cases. These are pre-theoretical judgments that are stably affirmed in conditions which minimize sources of error such as strong emotion or self-interest (1951, 181-3; compare 1999b, 42). The second stage is to find general moral principles that explain these judgments. The final stage of the procedure is to adjust some of the considered judgments and principles, as is necessary in order to produce 'equilibrium' (1951, 184-9).²¹ A person is in reflective equilibrium when the considered judgments and other moral principles that she endorses are coherent, or in Rawls's words, 'coincide' (1999b, 18). This version of the method appears inconsistent with the unity thesis, because there is no obvious way for metaethical considerations to enter into this theory-building process.

Rawls' own later work on the method casts doubt on the significance of this inconsistency. His later conception of the method instructs us to aim at a state that he calls 'wide reflective equilibrium'.²² A reflective equilibrium is *wide* if it is reasoned to from circumstances of ideal information. For Rawls, these ideal circumstances consist in exposure to all plausible conceptions of morality and to all reasonable arguments in favor of (and presumably against) each of these conceptions (1999b, 43; 1999a, 289).²³ Rawls's early version of the method did not include this informational requirement. Because of this, he later described such an approach as a 'narrow' reflective equilibrium methodology (1999a, 289).

Rawls's move to a wide reflective equilibrium methodology was well-motivated. Because it permits only a restricted set of inputs, narrow reflective equilibrium is objectionably conservative, granting our moral judgments a problematically privileged epistemic status.²⁴ For example, moving to wide reflective equilibrium allows us to consider arguments against a moral conception on the basis that it rests on objectionable grounds like bias or ideology.²⁵

This feature of wide reflective equilibrium also entails that the later version of Rawls' method does not conflict with the unity thesis. This is because the informational requirement just mentioned requires that we consider all 'reasonable arguments' for and

against candidate normative ethical theories. The second premise of the unification argument suggests the possibility of constructing reasonable arguments for normative conclusions that rest in part on metaethical premises. The later version of the method instructs us to consider those arguments, if they exist. Because the later version of the method itself says nothing about whether such arguments exist, the content of the method is best understood as *neutral* regarding the truth of the unity thesis.²⁶

To summarize, I have argued here for two claims. First, the apparent wide acceptance of Rawls' method is consistent with the existence of methodological controversy in ethics, because what is actually widely accepted does not adjudicate that controversy. Given the strength of the prima facie case for such controversy, the first premise of the basic unification argument is thus very hard to resist. Second, while the unity thesis is contrary to one of Rawls' central methodological ambitions, it is compatible with the most careful version of Rawls' method.

4. From metaethics to the methodology of normative ethics

The second premise of the basic unification argument claims that plausible metaethical theories can help to resolve the methodological controversy facing normative ethics. In this section, I assess two central objections to this premise. The first is that the idea of metaethics having normative implications violates the neutrality of metaethics. The second is that the contrasting implications of the three metaethics introduced in §2 reflect idiosyncratic phenomena that will not generalize.

It is sometimes claimed that metaethics is (or should be) 'neutral' between substantive normative theories. This idea can be traced back at least to A. J. Ayer, who famously proclaimed that philosophical work on ethics – roughly, metaethics – can 'make no ethical pronouncements' (1946, 103).²⁷ If metaethics must be normatively neutral, it may seem that there must be some sort of error in my claim that various metaethics can have contrasting implications for normative ethics.

This objection faces two serious difficulties. The first is that the clearest alleged support for the normative neutrality thesis tends to itself derive from the content of controversial metaethical theories. Ayer's case is exemplary: his commitment to neutrality

simply falls out of his crude emotivist metaethic. This point undercuts the force of the objection from neutrality, because the fact that one metaethic has such a consequence does not show that all plausible metaethics must.

The second difficulty is that it is crucial to distinguish between a metaethic *entailing* a certain normative thesis, and that metaethic having implications for the *methodology* of normative ethics. Recall the mechanism at the heart of my argument: metaethical theories can have contrasting implications for the correct moral epistemology, which can in turn have implications for the methods that normative ethicists should use in developing and assessing their theories. In order to have force against my argument, a neutrality thesis must therefore ban such methodological implications. However, neutrality theses are typically best understood as claiming only a ban on entailment, not on methodological implications. Ayer is again instructive. He argues that his view entails that legitimate argument that appears to be about values can really only concern the logic of one's moral commitments or relevant matters of empirical fact (1946, 110-112). Far from methodological silence, this is a very strong methodological claim.

One might accept that it is possible for metaethical theories to have implications for the methodology of normative ethics, but deny that the examples I have offered provide adequate reason to show that *in general* they tend to do so.²⁸ At its broadest, the only antidote to this sort of doubt is a comprehensive survey of the implications of leading metaethical theories; a task which considerations of space prevent me from taking up here. A sharper form of the challenge suggests that the features of a metaethic that generate implications for the methodology of normative ethics are idiosyncratic, and can easily be tweaked without affecting the heart of a metaethical view.²⁹ This challenge might seem to be supported by an example due to Robert Audi. Audi observes that while an intuitionist epistemology is typically associated with non-naturalistic metaethics, intuitionism is consistent with either non-cognitivism or naturalistic realism (2004, 2; 21).

It is true that the implications that metaethical views have for moral epistemology do not simply depend upon the broadest metaethical category that they fall into. However, this does not entail that they rest on easily-avoidable idiosyncrasies. To illustrate this point, I will discuss a single metaethic in more detail. I will take Boyd's metaethic as my example,

as the empiricist implications of this view are arguably the most methodologically radical of the examples introduced in §2.

The heart of Boyd's metaethic is a 'direct-reference' style of foundational moral semantics, broadly inspired by the proposals offered by Saul Kripke (1980) and Hilary Putnam (1975) for the reference of names and natural kind terms, and continuous with Boyd's own account of the semantics of scientific vocabulary. On this semantics, moral property terms are natural kind terms that refer directly to the clusters of properties that causally regulate the use of these terms. Boyd counts these clusters of properties as natural kinds in virtue of there being non-accidental causal mechanisms that tend to unify them, and to give us reliable access to them (1997, 116-7). The role of causal links in Boyd's semantics ensures that what's 'in the head' – our moral beliefs or intuitions – will play a quite modest role in our epistemology, roughly analogous to the role that our intuitions about water might play in investigating the chemistry of H₂O. Instead, Boyd's semantics suggests that a successful moral epistemology will provide a systematic understanding of the cluster of natural properties that causally regulate our use of moral terms. Given that the causal processes that unify moral properties and regulate our access to them centrally involve human activity, Boyd plausibly suggests that the methodology made appropriate by his theory is approximately that of the human sciences (1997, 122-4).

As this discussion makes clear, the epistemological consequences of Boyd's view flow from its core commitments: his metaphysical naturalism and semantic externalism. It would be very difficult to avoid these consequences of this theory without radically altering the theory itself.

To see this, it will be helpful to compare a leading naturalistic realist metaethic that does not have such clear empiricist epistemological implications: Frank Jackson and Philip Pettit's moral functionalism. The heart of this view is again an account of foundational moral semantics. Jackson and Pettit argue that the content of a moral term is given holistically, by the 'network of content-relevant connections' that it bears to other vocabulary, where this network is identified by examining our semantic intentions (1995, 22). These connections include input clauses such as paradigm instances of moral wrongness, internal role clauses that fix the relation between normative terms, and output clauses that indicate appropriate motivational responses to moral judgments (Jackson 1998,

130). Jackson and Pettit reasonably suggest that this holistic foundational moral semantics supports an intuitionistic methodology for normative ethics. Thus, the ‘paradigms’ and ‘commonplaces’ just mentioned are ‘candidates for *a priori* truth’ (Jackson and Pettit 1995, 23), because each will play a role in fixing the meanings of moral terms unless the resolution of tensions in the network requires their rejection.

Jackson and Pettit’s metaethic thus supports a very different moral epistemology from Boyd’s. The crucial point is that this difference arises from profound differences between their views in foundational semantics. One can thus predict that naturalistic realist views that share one or other of these semantic pictures will tend to share their epistemological consequences as well. These cases suggest that the implications for the methodology of normative ethics recommended by a metaethic tend to rest on some of its most central elements, not on easily modified idiosyncrasies.

5. Integration, again

The third premise of the basic unification argument claims that:

3. The best way to pursue the project of normative theorizing involves integrating relatively epistemically well-supported research that could help to resolve methodological questions about the justification of normative theories.

The most serious objection to this principle notes that it appears to state an epistemic principle that concerns acquiring evidence, and raises an influential challenge to this conception of epistemic norms.

To begin, recall two of the stages of the integration process briefly sketched in §1. The first stage involves actively acquiring (direct or indirect) evidence concerning a theory, and the second involves the proper use of such evidence in constructing that theory. My defense of premise 3 focused on the second stage of this process, motivating the idea that we should surely not disregard relevant evidence. This might seem to obscure important controversies about the first stage.

Some philosophers, such as Richard Feldman, claim that we can have only prudential or moral reasons (and not epistemic reasons) to seek out evidence. Consider Phoebe, who always perfectly apportions her beliefs to her evidence, but has truly atrocious investigative practices.³⁰ On Feldman’s view, despite this fault, Phoebe is always

fully justified in her beliefs, and is not open to any *epistemic* criticism, because she apportions her beliefs to her evidence (2000, 688-90). It might thus appear that I have illicitly committed myself to denying Feldman's view in this deep epistemological debate.

This objection rests on a mistaken understanding of the unity thesis. As I argued in §1, the unity thesis presupposes a goal of enquiry: to identify and elucidate the correct normative ethical theory. The same is true of the third premise of my argument. The goal just mentioned sounds to me like an epistemic goal (pace Feldman), but nothing crucial depends on this. The point required for my argument is that, because of this constitutive goal, Phoebe would be doing a much better job qua normative ethicist if she improved her investigative practices.

The third premise offers the right explanation of this truth, and of less egregious variants on Phoebe's case. Thus, suppose that Thetis has a minor vice as an enquirer. For example, suppose that (like many of us) Thetis is less diligent in her pursuit of evidence inconsistent with her pet theory than she is in her pursuit of evidence that supports it. Correcting even for such a common human failing would be an improvement in her pursuit of the correct normative theory. It is thus plausible to conclude that the *best* way for someone to pursue a theoretical goal would incorporate all reasonable improvements in her pursuit of evidence, just as the third premise claims.

6. Relative epistemic merit

The final premise of the basic unification argument is that metaethical research is epistemically well-supported relative to normative ethical research. One might challenge this premise by seeking to defend the significant epistemic superiority of normative ethical research, or (more radically) by suggesting that there is *no way* of reasonably adjudicating the relative epistemic status of these two domains.

The first challenge is difficult to motivate in the face of the point made in §2: metaethics and normative ethics are pursued using broadly similar tools in broadly similar contexts. Intuitive attempts to motivate this challenge also appear unpromising. For example, one might support the epistemic superiority of normative ethics by pointing to an example: there is controversy about metaethical theory, whereas it is just *obvious* that it would be wrong to torture an innocent child for fun. The problem is that there is a far

better explanation of this sort of example. While we are very confident in some particular judgments in both domains, theories in both domains are highly controversial. That there is theoretical controversy in normative ethics is too obvious to labor. Examples of theoretically important but relatively uncontroversial metaethical claims include the fact that moral vocabulary functions syntactically in ways identical to vocabulary that is usually interpreted as contributing to the truth conditions of sentences, and the fact that the ethical is supervenient.

The worry that there is *no way* of reasonably adjudicating the relative epistemic status of these two domains might be motivated as follows. Anyone committed to a theory in either metaethics or normative ethics is perhaps disposed to rate the epistemic quality of their theory and its domain as relatively high, and that of its competitors as low. Attempting to adjudicate these competing claims might seem to lead inevitably to dialectical paralysis, or simply amount to begging the question.³¹

Notice first that this worry is dialectically awkward for the opponent of the unity thesis. This is because it suggests that normative ethicists are in no position to legitimately claim that the unity thesis is false. And this in turn raises the question of how they could legitimately proceed as if it were false.

We can do better than raise these sorts of concern about these objections however. We can say something principled about how to adjudicate the controversies that lead to these worries. In outline, the right way to approach this problem is to identify general features, the presence or absence of which in a domain of research would, other things being equal, lead us to raise or lower our estimation of the quality of research in that domain. I call such features *generic indicators of the epistemic quality* of research in a given domain.³² Because these features are generic in this way, they are perfectly suited to provide non-question-begging support for one view over another in a controversy over relative epistemic merit.

Identifying such features and considering how they apply to work in metaethics and in normative ethics provides one promising means of assessing the relative epistemic quality of work in these domains. I argued in §2 that normative ethics and metaethics operate largely in the same intellectual context, and appear to share many methods. Nonetheless, it is possible to identify illustrative contrasts, of which I will briefly sketch two.

First, the fact that one domain of research is more deeply constrained by well-justified theory than another is a non-question-begging reason to believe that the research in the first domain is well-supported relative to research in the second. The central project in metaethics is to examine how ethics fits with our broader theories of how the world works. Because of this, metaethical research typically draws upon, and is constrained by, further philosophical and scientific theory. For example, consider the contrasting views in foundational semantics that I argued in §4 generated the contrasting methodological implications of Boyd's and Jackson and Pettit's metaethics. It would be impossible to responsibly argue the relative merits of these views without drawing on the general debates in foundational semantics between the sorts of view that each side seeks to apply to ethics. Being constrained in this way by well-justified theory should give us some grounds for confidence in metaethics, relative to work in normative ethics that is pursued in isolation from other domains.

Second, the fact that research in one domain is more vulnerable to the distorting influences of emotion, rationalization, and manipulation than another provides a non-question-begging reason to favor the second domain. The sensitivity of our moral judgments to (some of) our emotional reactions is both evident to common sense, and has been suggested experimentally.³³ Our moral judgments also appear more vulnerable to distortion as the result of interpersonal pressure or manipulation.³⁴ Finally, taking oneself to be doing bad things typically makes for an uncomfortable cognitive dissonance. This gives us some reason to worry that our values are simply rationalizations in the face of our own wants or social roles. The psychological mechanisms that underwrite rationalization are (like those that permit manipulation) probably quite general.³⁵ The pressures to rationalization and manipulation are significantly more worrisome in the normative domain than with respect to metaethics, however. This simply because one's role and interests – or the interests of the powerful – are much less likely to be in noticeable tension with one's metaethic than with one's normative commitments.

This is the merest sketch of the use of the generic indicators approach to relative epistemic quality in this context. However, it is worth emphasizing the difficulty of appealing to a differing assessment of such indicators in order to resist the fourth premise of my argument. This is (again) because this premise is so modest: all that it claims is that

metaethical research is not in significantly *worse* epistemic shape than normative ethical research. In order to reject this premise on this basis, the objector would have to mount a convincing case that the generic indicators suggest a *strong* asymmetry in epistemic quality, in favor of normative ethics, which is not a promising hypothesis.

One might object to the generic indicators strategy by noting that the epistemic quality of metaethical research *per se* is irrelevant to my thesis: what really matters is the epistemic quality of metaethical research that will contribute to the justification of claims about the methodology of normative ethics. It would be very bad for my argument if (for example) metaethics was overall quite well supported, but could be neatly divided into the well-supported but irrelevant to the methodology of normative ethics and the relevant but ill-supported.³⁶ The principle underlying this objection is exactly correct. However, it does not undermine the use of the generic indicators strategy. This is because of facts about the structure of actually plausible metaethical theories that have implications for the methodology of normative ethics (such as those that I have discussed in this paper). As we have seen, the methodological implications tend to flow from the central structural claims about the metasemantics, metaphysics, and psychology of ethics contained in those theories. Those structural claims in turn tend to be supported by their fit with wide-ranging bodies of evidence that support theses across the subdomains of metaethics. It is thus quite difficult to imagine that metaethics as a whole is relatively well-supported, but that the parts with implications for the methodology of normative ethics are not.

As I emphasized in §1, I am arguing that the best way of engaging in normative ethics *involves* metaethical investigation. This does not imply the much stronger claim that metaethics has the sort of epistemic superiority over normative ethics that would underwrite pursuing metaethics independently of normative ethics.³⁷ I am inclined to reject this stronger claim on behalf of metaethics. However, I do not argue for this here, and the discussion in this section helps to clarify why. The considerations sketched in this section may support the claim that metaethics is in better epistemic shape than normative ethics. Adequately addressing the question of whether metaethics is best pursued independently of normative ethics would require adjudicating two further questions. First, to what degree is the epistemic quality of metaethics superior to that of normative ethics? Second, what degree of epistemic inequality between two domains is needed to justify the sort of

methodological independence that (e.g.) physics appears to possess from philosophical ontology? These questions would be difficult to answer in a satisfactory way, and this leads me to refrain from arguing explicitly against the metaethicist's stronger priority claim here.

7. The significance of methodological unity

In this paper, I have briefly outlined an argument for the unity thesis, and then examined and rejected a range of the most promising ways that one might hope to resist that argument. In this final section, I begin by reflecting on the positive implications of the unity thesis, before addressing two challenges to the idea that the unity thesis could have such significant implications.

Consider first the implications of the unity thesis for the current practice of normative ethical theorizing. Put aggressively, the point is that, to the extent that they reject the unified approach to moral theorizing defended here, normative ethicists offer a hostage to fortune, in the form of a bald hope that wherever the metaethical chips may fall, their approach will be vindicated.

Consider an example. Frances Kamm is one of the most powerful contemporary normative ethical theorists, gifted with an extraordinary eye for the details and nuances of moral questions. However, her approach to normative theorizing consists almost exclusively in carefully consulting her own intuitive reactions to the sometimes extremely complex cases that she takes to be theoretically crucial. The truth of some leading metaethical views would raise serious questions about the probative significance of this approach. For example, consider Boyd's naturalism, sketched in §3. As noted there, Boyd suggests plausibly that the 'natural kinds' moral metaphysics and externalist semantic account that he defends would have significant methodological implications. These implications seem plainly inconsistent with Kamm's methods. Consider examples that are, on Boyd's theory, relevantly analogous: one could not hope to understand the nature of water or economic growth by appeal to one's intuitive judgments about outlandish cases. If moral goodness were a Boydian natural kind, such a method would be equally unreliable in normative ethics.

For anyone interested in applying a Kamm-style methodology in normative ethics, this sort of example makes two questions salient. First, what alternatives to Boyd's

metaethic *would* vindicate Kamm's method? And second, how plausible are such alternatives, compared to the metaethics that would render such methods dubious? The force of the unity thesis is exactly in the suggestion that the best way to proceed in normative theorizing involves being prepared to seriously address such questions.

The point here is not to criticize Kamm or to defend Boyd. Rather, as a philosopher with strong interests in normative ethics, it seems to me that Kamm's problem is my problem too. We all need principled ways of resolving the methodological controversy that confronts our project, and metaethics is one promising source for discovering such principled resolution.

Consider a striking example. One recent controversy concerns the normative (in)significance of recent empirical research examining the psychology of moral judgment. Thus, some have suggested that our tendencies to judge that there is a moral asymmetry between doing harm and allowing harm to be done, or between intending harm and merely foreseeing it as a side-effect of one's action, can be explained away as results of allegedly morally arbitrary psychological processes.³⁸ There has been much controversy generated by such arguments. The unity thesis suggests that metaethical research could shed light on these debates. This is because accepting different accounts of the nature of moral concepts or moral properties may well lead to very different evaluations of the significance of these psychological processes, and the research that explores them.³⁹

The unity thesis suggests an approach to the methodology for systematic normative ethics that is distinctive in its focus on the relation between metaethics and normative ethics. It thus motivates positive attention to comparatively neglected questions in ethical theorizing. Most notably, it suggests that progress in ethics would be greatly aided by a systematic program of research into the implications of metaethical theories for normative ethics and its methodology. These relations have typically been treated superficially (if at all) by ethical theorists; my own discussion here barely scratches the surface. If one actually wanted to use the results of metaethical research in one's normative theorizing, a much more thorough and careful treatment of such implications would be called for.

It might be objected that taking the unity thesis seriously would make systematic normative ethics *too hard*. After all, it is hard enough to adjudicate between consequentialism, Kantianism, and their various normative competitors. The unity thesis

demands unreasonably that we simultaneously attempt to adjudicate between naturalist realism, expressivism, non-naturalism, fictionalism, and so on. There is something plausible about this worry. However, the right lesson to draw here cannot be that we should have confidence in our ability to make progress in systematic normative ethics, but *only* if we ignore some of the considerations that are likely to have significant import for the justifiability of a normative theory.⁴⁰ The more plausible conclusion is that the unity thesis clarifies the inescapable difficulty of doing systematic normative ethics, in the face of the Gordian knot of entrenched controversy that entangles it.

A related worry is that the unity thesis generalizes in an implausible way. For example, it might seem to demand (absurdly) that philosophers interested in the ethics of global warming do original climate research. Instead, the objector might suggest, such philosophers should simply appeal to the results generated by climate experts and reflect on the ethical implications of those results.

The significance of this objection hinges on the question of whether the ethically important facts about climate change are (a) uncontroversially known within the community of climate scientists, and (b) comprehensible in their ethically relevant respects to a climate non-expert. If both (a) and (b) were true, then the ‘climate ethicist’ should simply responsibly apply the relevant scientific data as background to her argument, and the generalization of the unity thesis would be unimpugned.

If either (a) or (b) were false, then the climate ethicist would evidently do well to collaborate in her research with climate scientists capable of responsibly adjudicating the controversy and interpreting the science. Failing this, a responsible philosopher should carefully couch her results in conditional form: “*if such and such are the facts about the climate, then the following are our obligations...*” However convenient it would be, it could not be responsible, if either (a) or (b) were false, for the non-climate-expert ethicist to make unconditional claims about the ethics of our response to climate change, without either becoming an expert in climate science herself or appealing to expert collaboration.

The division of labor suggested by the conditional reasoning strategy just mentioned makes considerably more sense in the case of climate ethics than it does in the case of the relationship between metaethics and normative ethics. This is because one might reasonably claim that identifying the correct conditional normative principles for

climate ethics exhausts the philosopher's distinctive contribution to the question of climate ethics. One could not say the same thing about 'methodologically conditional' results in normative ethics. Rather, the latter form of conditional theorizing (*if* the correct method for normative ethics is ..., then...) would be a profoundly odd place for a philosopher to end her enquiry.

This conditional form of normative theorizing would entail abandoning the aim of normative theorizing as I have characterized it, perhaps in favor of the goal (mentioned in §1) of making a contribution to the social project of philosophy. Some may take the two objections just canvassed to motivate the replacement of the ambitious goal of identifying the correct normative theory with this more modest 'contribution' goal. It is important to note that the unity thesis is still significant for theorists with this goal. This is because we can think of the unity thesis as applying not merely to a single individual, but to the community of normative ethical enquirers as a whole.

So understood, the unity thesis does not demand that individual philosophers seek to develop well-justified systematic normative ethical theories. However, it does suggest that the community as a whole could be doing better or worse at pursuing that goal. The argument of this paper suggests that *the community* of enquirers is arguably failing to attend adequately to the issue of the potential implications of metaethical theories for normative ethics and its methodology. Even given a 'social division of labor' in ethics, accepting the unity thesis clarifies the way that research on the methodological connections between metaethics and normative ethics could be extremely fruitful, potentially putting the community of normative ethicists in a position to develop their views in a more methodologically sophisticated way.

I do not agree that we are forced to the modesty of ambition just suggested. I am inclined to optimism about what individual philosophers can achieve. I take it that there has been real progress in our understanding of ethics over the past century, and that much of that progress has come (as in other times in the history of philosophy) via the system-building that has been making a recovery in our discipline. Clarifying the distinction between metaethics and normative ethics was an important achievement of Twentieth century ethical theorizing. However, it would be a mistake to take it to force systematic moral philosophers to narrow their focus to one of those domains at a time.⁴¹

References

- Audi, R. (2004). *The Good in the Right*. Princeton: Princeton UP.
- Ayer, A. J. (1946). *Language, Truth and Logic*. 2nd rev. ed. London: Gollancz.
- Berker, S. (2009) "The Normative Insignificance of Neuroscience," *Philosophy and Public Affairs* 37.4, pp. 293-329.
- Boyd, R. (1997). "How to be a Moral Realist," in Darwall, Gibbard, and Railton, eds, pp. 105-136.
- Copp, D. (1985). "Considered Judgment and Moral Justification: Conservatism in Moral Theory," *Morality, Reason, and Truth: New Essays on the Foundations of Ethics*. Ed. David Copp and David Zimmerman. Totowa, NJ: Rowman & Allanheld, pp. 141-68.
- . (2012). "Experiments, Intuitions, and Methodology in Moral and Political Theory," *Oxford Studies in Metaethics* Vol. 7. Ed. Russ Shafer-Landau. Oxford: Oxford UP, pp. 1-36.
- Dancy, J. (1985). "The Role of Imaginary Cases in Ethics," *Pacific Philosophical Quarterly* 66, pp. 141-153.
- Daniels, N. (1979). "Wide Reflective Equilibrium and Theory Acceptance in Ethics," *Journal of Philosophy* 76.5, pp. 256-82.
- Darwall, S., A. Gibbard, and P. Railton, eds. (1997). *Moral Discourse and Practice*. New York: Oxford UP.
- , ---, and ---. (1997a). "Towards *Fin de Siècle* Ethics: Some Trends," in Darwall, Gibbard, and Railton eds, pp. 3-47.
- DePaul, M. and W. Ramsey, eds. (1998). *Rethinking Intuition: The Psychology of Intuition and Its Role in Philosophical Inquiry*. Lanham, MD: Rowman and Littlefield.
- Doris, J. and S. Stich. (2006). "As a Matter of Fact: Empirical Perspectives on Ethics," in Jackson and Smith, pp. 114-152.
- Dworkin, R. (1986). *Law's Empire*. Cambridge MA: Harvard UP.
- . (1996). "Objectivity and Truth: You'd Better Believe It," *Philosophy and Public Affairs* 25.2, pp. 87-139.
- Enoch, D. (2007). "An Outline of an Argument for Robust Metanormative Realism," *Oxford Studies in Metaethics* Vol. 2. Ed. Russ Shafer-Landau. Oxford: Oxford UP, pp. 21-50.
- Feldman, R. (2000). "The Ethics of Belief," *Philosophy and Phenomenological Research* 60.3, 667-695.
- Fodor, J. (1983). *Modularity of Mind*. Cambridge, MA: MIT Press.
- Gazzaniga, M. S., J. E. LeDoux, and J. H. Wilson. (1977). "Language, Praxis, and the Right Hemisphere: Clues to some Mechanisms of Consciousness," *Neurology* 27, pp. 1144-1147.
- Gibbard, A. (1990). *Wise Choices, Apt Feelings*. Cambridge, MA: Harvard UP.
- . (2008). *Reconciling Our Aims*. New York: Oxford UP.
- Greene, J. (2008). "The Secret Joke of Kant's Soul," *Moral Psychology Vol. 3: The Neuroscience of Morality*. Ed. Walter Sinnott-Armstrong. Cambridge, MA: MIT Press, pp. 35-80.

- , F. Cushman, L. Stewart, K. Lowenberg, L. Nystrom, and J. Cohen. (2009). "Pushing moral buttons: The interaction between personal force and intention in moral judgment," *Cognition* 111.3, pp. 364-71.
- Hare, R. M. (1963). *Freedom and Reason*. New York: Oxford UP.
- . (1973). "Critical Study: Rawls' Theory of Justice," *Philosophical Quarterly* 23, pp. 144-155; 241-251.
- . (1981). *Moral Thinking: its Levels, Method, and Point*. Oxford: Clarendon.
- Harman, G. (1977). *Nature of Morality*. New York: Oxford UP.
- Hooker, B. and M. O. Little, eds. (2000). *Moral Particularism*. Oxford: Clarendon.
- Horowitz, T. (1998). "Philosophical Intuitions and Psychological Theory," in DePaul and Ramsey pp. 143-160.
- Huemer, M. (2008). "Revisionary Intuitionism," *Social Philosophy and Policy* 25, pp. 368-92.
- Hussain, N and N. Shah. (ms.) "Metaethics and Its Discontents: A Case Study in Korsgaard," ms. version 2.2.
- Jackson, F. and M. Smith. (2006). *Oxford Handbook of Contemporary Analytic Philosophy*. Oxford: Oxford UP.
- and Philip Pettit. (1995). "Moral Functionalism and Moral Motivation," *Philosophical Quarterly* 45.178, pp. 20-40.
- James, W. (1956). "The Will to Believe," *The Will to Believe and Other Essays in Popular Philosophy*. New York: Dover, 1956.
- Jones, K. (2006). "Moral Epistemology," in Jackson and Smith, pp. 63-85.
- Joyce, R. (2001). *The Myth of Morality*. Cambridge: Cambridge UP.
- Kagan, S. (1998). *Normative Ethics*. Boulder, CO: Westview.
- . (2001). "Thinking About Cases," in *Moral Knowledge*. Ed. E. F. Paul, F. D. Miller, and J. Paul. Cambridge: Cambridge UP, pp. 44-63.
- Kamm, F. (1998). "Moral Intuitions, Cognitive Psychology, and the Harming-Versus-not-Aiding Distinction," *Ethics* 108.3, pp. 463-488.
- Kornblith, H. (1983). "Justified Belief and Epistemically Responsible Action," *Philosophical Review* 92.1, pp. 33-48.
- Korsgaard, C. (2003). "Realism and Constructivism in Twentieth-Century Moral Philosophy," *Journal of Philosophical Research* APA Centennial Supp: Philosophy in America at the Turn of the Century, pp. 99-122.
- Kramer, M. (2009). *Moral Realism as a Moral Doctrine*. Oxford: Blackwell.
- Kripke, S. (1980). *Naming and Necessity*. Cambridge, MA: Harvard UP.
- Levine, P. (1998). *Living Without Philosophy*. Albany, NY: SUNY Press.
- Mackie, J. L. (1977). *Ethics: Inventing Right and Wrong*. Harmondsworth: Penguin.
- McNaughton, D. (2003). "An Unconnected Heap of Duties?" *Ethical Intuitionism: Re-Evaluations*. Ed. Philip Stratton-Lake. Oxford: Oxford UP.
- McPherson, T. (ms.) "From Experimental Philosophy to Normative Ethics?" Available from author.
- . (2008). "Metaethics and the Autonomy of Morality," *Philosophers' Imprint* 8.6, pp. 1-16.
- . (2009). "Moorean Arguments and Moral Revisionism," *Journal of Ethics and Social Philosophy* 3.2, 1-24.
- Miller, A. (2003). *Introduction to Contemporary Metaethics*. Cambridge: Polity, 2003.
- Moore, G.E. (1903). *Principia Ethica*. Cambridge: Cambridge UP.
- Peacocke, C. (1999). *Being Known*. Oxford: Clarendon.

- Putnam, H. (1975). "The Meaning of 'Meaning'," *Mind, Language and Reality: Philosophical Papers, Volume 2*. Cambridge: Cambridge UP, pp. 215-271.
- Rawls, J. (1951). "Outline of a Decision Procedure for Ethics." *Philosophical Review* 60, pp. 177-97.
- . (1999a) "The Independence of Moral Theory," *Collected Papers*. Ed. Samuel Freeman. Cambridge, MA: Harvard UP, pp. 286-302.
- . (1999b). *Theory of Justice* Rev. ed. Cambridge, MA: Belknap.
- Scanlon, T. M. (1998). *What We Owe to Each Other*. Cambridge, MA: Belknap.
- . (2003). "Rawls on Justification." *The Cambridge Companion to Rawls*. Ed. S. Freeman. Cambridge, U.K.; New York: Cambridge University Press, pp. 139-67.
- Shafer-Landau, R. (2003). *Moral Realism: A Defence*. Oxford: Clarendon.
- Singer, P. (1972). "Famine, Affluence, and Morality." *Philosophy and Public Affairs* 1.1, pp. 229-243.
- . (1974). "Sidgwick and Reflective Equilibrium." *Monist* 58.3, pp. 490-517.
- . (2005). "Ethics and Intuitions." *Journal of Ethics* 9.3-4, pp. 331-52.
- Sinnott-Armstrong, W. (2006). *Moral Skepticisms*. New York: Oxford UP.
- Smith, M. (1994). *Moral Problem*. Oxford: Blackwell.
- Street, S. (ms.) "Objectivity and Truth: You'd Better Rethink It." Ms. available at <http://homepages.nyu.edu/~jrs477/streetpapers.html>
- Thomson, J. J. (1990). *The Realm of Rights*. Cambridge, MA: Harvard UP.
- Unger, P. (1996). *Living High and Letting Die: Our Illusion of Innocence*. New York: Oxford UP.
- Walsh, S. D. (2009). "Teleology, Aristotelian Virtue, and Right," *Ethics: The Big Questions*. Ed. J. Sterba. Oxford: Blackwell, pp. 409-418.
- Wheatley, T, and J. Haidt. (2005). "Hypnotic Disgust makes Moral Judgments more Severe," *Psychological Science* 16.10. pp. 780-784.
- Wilkes, K. (1988). *Real People: Personal Identity without Thought Experiments*. Oxford: Clarendon.
- Williams, B. (1993). *Morality (Canto Ed.)* Cambridge: Cambridge UP.
- Zimmerman, D. (1980). "Meta-Ethics Naturalized." *Canadian Journal of Philosophy* X.4, pp. 637-662.

¹ For the claim that such explanatory claims (as opposed to merely extensional claims) form the 'heart' of a normative ethical theory, see e.g. Thomson 1990, 30.

² Many normative theorists share the simple utilitarian's explanatory ambitions, but disagree with her substantive proposal. Note, however, that others think that normative ethicists should have more modest explanatory ambitions. For example, Scanlon offers a contractualist account of wrongness in terms of reasonable rejectability. However, he insists that the notion of the *reasonable* in this account resists any simple non-normative characterization, and hence that application of the contractualist principle will inevitably require good judgment (1998, 245-6). Walsh 2009 argues that most plausible versions of virtue theory and deontology likewise deny that we can give a simple and exhaustive non-normative characterization of rightness. While not as explanatorily ambitious as simple utilitarianism, these theories all fall under the scope of the unity thesis due to their general scope.

³ Thinking of metaethics in terms of this central project sidesteps two potential challenges in thinking about metaethics. First, there is a notable range of characterizations of metaethics. Thus, some worries about metaethics seem to take it to concern only questions of meaning, leading to worries that a distinctive metaethical project disappears, or needs to be redefined, in the face of Quinean worries about the analytic-synthetic distinction (compare Harman 1977, vii, and Zimmerman 1980). By contrast, more recent discussions (e.g. Darwall, Gibbard, and Railton 1997, 7, Miller 2003, 2, and Sinnott-Armstrong 2006, 6), expand the list of questions that count as 'metaethical,' but not in an entirely uniform way. My

characterization of the metaethical project as a kind of location project helps to clarify what is common to these accounts, allowing the differences between these views to be explained by contrasting substantive philosophical commitments.

⁴ Notice that this gloss on metaethics takes it to concern *morality*. Some recent philosophers have helpfully distinguished metaethics which focuses in this way on the moral, from the *metanormative*, which focuses more broadly on practical, or even epistemic normativity (compare especially Enoch 2007). Although I am sympathetic to the idea that a very similar argument could be run in the metanormative case, my argument in this paper will focus on metaethics.

⁵ If metaethics and normative ethics were not distinct, pursuing them separately would simply be confused. For serious doubts about this apparent distinctness, see especially Dworkin 1996 and Korsgaard 2003. For a partial undermining diagnosis of those doubts, see McPherson 2008. Kagan 1998, 5 suggests that there is a ‘continuum’ between metaethics and normative ethics, rather than a clear distinction between them. For a reply to Kagan, see Hussain and Shah ms., 13. Kramer 2009 suggests another interesting view: that these projects may be distinct, but that much of what philosophers typically take to be metaethics is in fact part of normative ethics. I take Kramer’s argument to rest partly on an implicit assumption that metaethical claims must definitionally be neutral between normative theories. For discussion of this issue, see §4 below.

⁶ I am thus not committed, as Smith appears to be, to thinking that philosophers should ‘...give meta-ethical questions a certain priority over questions in normative ethics’ (1994, 2).

⁷ This claim would be uninterestingly modest. Compare Fodor’s analogous general claim that ‘In principle, our botany constrains our astronomy, if only we could think of ways to make them connect’ (1983, 105).

⁸ The nature of intuition is itself a controversial issue. See the papers collected in DePaul and Ramsey 1998 for a variety of views concerning the nature and status of philosophical intuitions.

⁹ Kagan 2001, 44 suggests that at least tacit appeal to intuition may be inescapable.

¹⁰ Representative worries about moral intuitions can be found in Hare 1981, 130ff and Singer 2005.

¹¹ The use of outlandish cases is of course extremely common in ethics. Worries about the philosophical use of imaginary cases are suggested by Dancy 1985 and Wilkes 1988. Distrust in fantastic cases – sometimes dubbed ‘desert island ethics’ – is widespread but rarely receives principled defense.

¹² See Huemer 2008 for a valuable discussion of the relative merits of different types of intuitions. Appeals to the plausibility of moral principles are perhaps as old as ethics. Appeals to the moral (ir)relevance of certain properties and to theoretical desiderata are perhaps less familiar, but are sometimes taken to play a crucial role in normative argument. A striking example of the former is Unger 1996, which follows Singer 1972 in appealing to the intuitive idea that mere physical proximity cannot be a morally relevant property. I take it that the appeal of theoretical desiderata like simplicity and explanatory power helps to explain why we have seen so many efforts to work out normative theories that explain everything about morality in consequentialist or deontological or aretaic terms. Rossians and particularists, of course, argue that these considerations are not decisive. See for example McNaughton 2003.

¹³ Doris and Stich 2006 is a good (if already dated) point of entry into this burgeoning empirical literature.

¹⁴ Another notable contrast is that the pressure under consideration here is not in the first instance rational pressure: because judgments of rationality are made from *within* an attitude-constituted perspective in Gibbard’s account, one could in principle experience rational pressure in almost any direction. These pressures are thus to be understood rather as robust psychological tendencies in ordinary persons.

¹⁵ For two discussions of this pressure, see Gibbard 1990, 275-6 and 2008, 23.

¹⁶ Tom Dougherty suggests an epicycle that may need to be added to the argument. Suppose that one domain of research was as well-supported as another, but it was *much easier* to acquire evidence in the first than in the second. This is arguably a counterexample to the sufficiency of premise 3: it would seemingly be best for researchers interested in the first domain to focus on investigating the first domain and ignoring the second. This may be right, although the details are complex. However, it is not plausibly true of normative ethics and metaethics that research in one is markedly easier than research in the other. Even if required, this amendment will thus not affect the force of my argument.

¹⁷ Compare a simple analogy concerning sources rather than domains. A bit of testimony might suffice to justify my belief that *P*, despite the fact that it would usually be trumped by clear competing visual evidence that not-*P*. Denying that one domain of research or source of evidence is clearly epistemically superior to another thus does not entail skepticism about the latter source or domain.

¹⁸ Two examples will suffice to give the flavor of that endorsement. First, Kagan 1998, 16 suggests that *all* practicing normative theorists are at least implicitly committed to something very similar to the method.

Second, Smith 1994, 40 suggests that Rawls' account of the method successfully systematizes our methodological platitudes.

¹⁹ It is also notable that Scanlon suggests an *interpretation* of Rawls' own account of the method that similarly fails to offer an informative resolution to methodological controversy. On this interpretation, Scanlon suggests that the method '...becomes simply the truism that we should decide what views about justice [or ethics more broadly] to adopt by considering the philosophical arguments for all possible views and assessing them on their merits' (2003, 151). This is a gloss that a proponent of the unity thesis can wholeheartedly endorse.

²⁰ Rawls' central stated goal when first introducing the method in his 1951, §1.3 was to articulate a conception of objectivity in ethics and a method for objective ethical enquiry that each avoid metaethical commitment. Similarly, Rawls 1999a invoked his later version of the method as a central element of his argument for what he called the 'independence of moral theory' (from metaethics, among other things).

²¹ I am abstracting away from a host of fascinating features of Rawls' 1951 proposal here, to focus on those issues most relevant to its relation to the unity thesis.

²² Rawls introduces the term 'wide reflective equilibrium' in 1999a, 289. However, he makes it clear there that he intends this term to have the same sense as what he calls the 'philosophically relevant' equilibrium in 1999b, 43.

²³ This exposition elides an important difficulty in interpreting the *method* of reflective equilibrium. The ideal of reflective equilibrium involves features – exposure to *all* plausible moral conceptions and arguments – that are likely not humanly attainable. This raises an important question to which there is no obviously unproblematic answer: what significance does this methodological ideal have for how *non-ideal* reasoners ought to develop their normative theories?

²⁴ A point made forcefully by Daniels 1979, 258n4. Rawls himself is more cryptic, merely noting that 'Clearly it is the second [i.e. wide] kind of reflective equilibrium that we are concerned with in moral philosophy' (1999b, 43).

²⁵ Not everyone is convinced that the move to wide reflective equilibrium *suffices* to escape the charge of objectionable conservatism. See Hare 1973, Singer 1974, and Copp 1985 for classic challenges to the method that invoke such worries.

²⁶ Rawls 1999a offers a further argument for the 'independence of moral theory.' This argument is complex, and its relevance to normative ethics is made uncertain by Rawls' psychological gloss on what he calls 'moral theory'.

²⁷ Ayer was followed in his defense of normative neutrality by a number of non-cognitivists, including Hare (e.g. 1963, 88-9), although the notion of neutrality in Hare and other later non-cognitivists is arguably weaker than Ayer's. Mackie 1977, 16 makes claims for his error theory that have been (mis-)interpreted as endorsement of a kind of neutrality. Finally, one reading of Dworkin 1996 understands him as treating a form of normative neutrality as a constraint on counting as a metaethical theory (an assumption that may be shared by Kramer 2009). This reading is supported by its continuity with Dworkin's related discussion at 1986, 68. For a discussion and critique of this part of Dworkin, see McPherson 2008, §2 and Street (ms.).

²⁸ One might motivate this denial by noting that metaethicists do not often advertise that their views have such methodological implications. The force of this point is weakened by the fact that it is typically in the dialectical interest of a metaethicist to de-emphasize any controversial implications that her theory might have. As I have argued in §2, implications for the methodology of normative ethics tend to be controversial.

²⁹ Thanks to Nadeem Hussain for pressing a variant of this objection.

³⁰ Phoebe is loosely modeled on Kornblith's example of the physicist (1983, 36), used in arguing for the view that Feldman attacks.

³¹ James sardonically encapsulated the attitude that generates the problem: 'Other philosophies are collections of opinions, mostly false; *my* philosophy gives standing-ground forever' (1956, 12-13; emphasis his).

³² I apply this same strategy for assessing the relative epistemic quality of propositions in my 2009.

³³ Note a single vivid example: elicited disgust reactions in subjects have been shown to have some causal impact on their moral judgments (Wheatley and Haidt 2005). Of course, philosophers can become emotionally invested in their views, no matter the subject matter. The claim suggested by studies like the one just mentioned is that there is a distinctive psychological linkage between emotion and moral judgment in addition to this sort of domain-general worry.

³⁴ Karen Jones suggests plausibly that such worries about ideological influence on our moral beliefs may be ‘...the source of much ordinary skepticism about moral knowledge’ (2006, 66).

³⁵ Compare the vast psychological literature on cognitive dissonance. One amazing non-ethical example is the easy confabulation practiced by split-brain patients called upon to explain the behavior of the half of their bodies over which they do not have control (Gazzaniga et. al. 1977, 1146).

³⁶ I am indebted to an anonymous reviewer for *Pacific Philosophical Quarterly* for raising this worry.

³⁷ Note as well that, in identifying certain worries about the epistemic status of normative theorizing, I do not intend to suggest that we should become skeptics about the possibility of moral theory. My argument has been *comparative*, arguing that there are reasons for taking metaethics to be in at least as good epistemic shape as normative ethics. I have not addressed the different and vexed question of whether normative theorizing can entitle us to reasonable belief in general principles. As I noted above, that one sort of evidence has epistemic priority over a second type does not show that the second type is unable to justify belief. For the attempt to deploy considerations such as these in support of moral skepticism, see Sinnott-Armstrong 2006, 195-210.

³⁸ For the doing-allowing distinction, see for example Horowitz 1998 and Kamm 1998. Greene has been at the forefront of the attempt to understand the psychology of the doctrine of double-effect. He suggests debunking explanations of this doctrine in his 2008, and develops the psychological account in more detail in Greene et al. 2009 (which also references much other relevant psychological literature). For discussion and criticism of the alleged ethical implications of Greene’s work see especially Berker 2009.

³⁹ I argue for this point in detail in my (ms.).

⁴⁰ *Pace* Copp 2012, which argues that we interpret normative ethicists as having a surprisingly unambitious ‘proximate aim’. Copp’s suggestion is seemingly intended to allow normative ethicists to simply ignore the worry that their methods might be leading them away from the *correct* theory. This is not something that I am prepared to do when engaging in normative theorizing.

⁴¹ I am indebted to many people for helping me to develop the ideas set out in this paper. First and foremost are Michael Smith and Tom Kelly. Others who generously offered feedback include Derek Baker, Arudra Burra, Sarah Buss, Ramon Das, Tom Dougherty, Gil Harman, Liz Harman, Nadeem Hussain, Sari Kisilevsky, Colin Klein, Sarah-Jane Leslie, John Maier, Philip Pettit, David Plunkett, Brendan Ritchie, Gideon Rosen, Mark Schroeder, Walter Sinnott-Armstrong, Nick Stang, Sean Walsh, and audiences at Princeton, Cal State LA, Louisiana State, the University of Minnesota Duluth, and Stanford. Last but not least, I am grateful to an anonymous reviewer for *Pacific Philosophical Quarterly*. Almost all of these people disagree with me about something I say herein, and none should be held responsible for any errors that remain despite their efforts.