

INDETERMINACY AND VARIABILITY IN META-ETHICS

Abstract

In the mid-20th century, descriptive meta-ethics addressed a number of central questions, such as whether there is a necessary connection between moral judgment and motivation and whether moral reasons are absolute or relative. I maintain that much of this work in mid-20th century meta-ethics proceeded on an assumption that there is good reason to question. The assumption was that our ordinary discourse is uniform and determinate enough to vindicate one side or the other of these meta-ethical debates. I suggest that ordinary moral discourse may be much less uniform and determinate than 20th century meta-ethics assumed.

INDETERMINACY AND VARIABILITY IN META-ETHICS

“I do not claim to have captured the one precise sense that the word ‘value’ bears in pure speech, uncorrupted by philosophy, that is heard on the Clapham omnibus. So far as this matter goes, I doubt that speakers untouched by philosophy are found in Clapham or anywhere else. And if they were, I doubt they’d have made up their minds exactly what to mean any more than the rest of us have. I take it, rather, that the word ‘value’, like many others, exhibits both semantic variation and semantic indecision” (Lewis, 2000, 86).

“The concept of morality itself bears the accumulated scars of conceptual evolution. Its multiple associations are a bar to summing it up in any one way” (Falk, 1986, 231).

1. Introduction

In the mid-20th century, descriptiveⁱⁱ meta-ethics addressed a number of central questions, such as: Is there a necessary connection between moral judgment and motivation? Are moral reasons absolute or relative? Can we derive an “ought” from an “is”? Do ordinary moral judgments express attitudes or describe states of affairs? Does moral judgment embody metaphysical error? In this paper, I maintain that much of this work in mid-20th century meta-ethics proceeded on an assumption that there is good reason to question. The assumption was that our ordinary discourse is uniform and determinate enough to vindicate one side or the other of these meta-ethical debates. Call this the Uniformity-Determinacy — or UD — assumption.

As an alternative to the UD assumption, I will offer two theses: the Indeterminacy Thesis and the Variability Thesis — the combination of which I will call the “IV Thesis.” The Indeterminacy Thesis holds that some parts of ordinary moral discourse give us no reason to prefer a descriptive analysis that involves one meta-ethical commitment over a descriptive analysis that involves the commitment that has traditionally been taken to be its meta-ethical competitor — e.g., that there are some parts of ordinary discourse that an absolutist account does no better at explaining than a relativist account, and vice versa. The Variability Thesis holds that while some parts of ordinary moral discourse are most accurately analyzed as involving a certain meta-ethical commitment, other parts of ordinary moral discourse are most accurately analyzed as involving the meta-ethical commitment that has traditionally been taken to be its meta-ethical competitor.ⁱⁱⁱ

I do not expect to establish here that the UD assumption is false and the IV Thesis true. But I do hope to make the case for the IV Thesis’ being considered a plausible explanatory rival to the UD assumption — to make the case that descriptive meta-ethics ought to have taken seriously the possibility that the best analysis of ordinary moral discourse would vindicate neither of the traditional sides of many of the 20th century’s central meta-ethical debates.

In section 2, I’ll present a programmatic sketch of how the UD assumption figured in 20th century meta-ethics and of the main aspects of the IV Thesis. In section 3, I’ll situate the differences between UD and IV in the context of the (motive) internalism-externalism debate. In section 4, I’ll briefly discuss some of the implications for descriptive meta-ethics of taking the IV Thesis seriously.

2. The Unitary-Determinate Assumption

How did 20th century philosophers do descriptive meta-ethics? There was, of course, no single method, but one approach was predominate from at least about 1944, when Stevenson published *Ethics and Language*, until at least about 1994, when Smith published *The Moral Problem*. This approach had two stages.

The first stage consisted of gathering examples of moral ideas and terms in everyday language and thought. It consisted of canvassing commonsense moral judgments, linguistic intuitions, and platitudes.

The second stage consisted of trying to show that the best analysis of the evidence gathered in the first stage was that morality had a certain conceptual shape. It consisted of trying to show that what ordinary people say and think is best explained by a moral concept that is internalist or externalist, absolutist or relativist, etc.

As commonly practiced in the 20th century, this two-stage approach assumed UD. It assumed that one can glean from ordinary discourse a single, principled answer to the question of whether morality is internalist or externalist, absolutist or relativist, etc. One indication of the UD assumption is that of the two stages, most of the work in 20th century meta-ethics occurred in the second. Twentieth century meta-ethicists typically presented some examples of ordinary discourse. But they didn't gather data in a comprehensive and systematic way. A handful of illustrative cases were taken to be a sufficient starting point. And if ordinary moral discourse is based on a concept unitary and determinate enough to imply definitive answers to the oft-disputed meta-ethical questions, this procedure is warranted. For if the concept of morality is sharply unitary

and robustly determinate — if the relevant meta-ethical information is encoded in the DNA of every use of moral terms — then one handful of commonsense judgments, intuitions, and platitudes will instantiate the same meta-ethical commitments as any other.

Parties on both sides of each of the 20th century meta-ethical debates realized, however, that there were some bits of data that did not appear at first glance to fit easily into their particular meta-ethical positions. Internalists, for instance, acknowledged that there were some bits of ordinary discourse that an externalist analysis might at least initially look to handle more naturally, and externalists acknowledged that there were some bits of ordinary discourse that an internalist analysis might at least initially look to handle more naturally. So proponents on each side then went on to try to show that while there were some cases that looked to be problematic for their position, those cases could be explained away, but that it was more difficult for their opponents to explain away the cases that looked to be problematic for their position. Each side contended, in other words, that while their position had some awkward elements, the opposing position had *even more* awkward elements, and that this was a powerful reason for thinking that their own position was true. And, once again, if we assume UD, this procedure is warranted. For if commonsense morality is sharply unitary and robustly determinate, then it will be legitimate to hold that, say, internalism and externalism are mutually exclusive and jointly exhaustive meta-ethical options. And if internalism and externalism are mutually exclusive and jointly exhaustive, then establishing that one of those positions fits ordinary discourse less awkwardly than the other will be a powerful reason for believing that the former accurately describes our concept of morality.

If, however, we keep open the question of how unitary and determinate commonsense morality is — if we take the IV Thesis to be a serious meta-ethical contender — then externalism’s being more awkward than internalism will not necessarily push us to the internalist conclusion. For the IV Thesis constitutes a third meta-ethical option, one that may fit the contours of ordinary moral discourse better, with less awkwardness, than either internalism or externalism.^{iv}

According to the Indeterminacy Thesis (which is the “I” of the IV Thesis), parts of our moral thought and language provide no good answers to the questions that were central to much of 20th century meta-ethics, vindicating neither relativism nor absolutism, neither internalism nor externalism, etc. The Indeterminacy Thesis holds that the relationship between some instances of ordinary moral discourse and these meta-ethical debates is analogous to the relationship between ordinary arithmetic and debates in the philosophy of mathematics. There is no fact of the matter as to whether ordinary mathematic usage is better explained by a Platonist or anti-Platonist understanding of number. The way people use numbers in everyday math simply does not contain answers to the questions that animate philosophy of mathematics. That is not to say that the question of what numbers are isn’t philosophically important. But it’s an ontological or substantive question, one on which conceptual analysis of ordinary arithmetic does not gain purchase. Similarly, there may be no fact of the matter as to whether parts of ordinary moral discourse are better explained by, say, absolutism or relativism (even if, as we will find the Variability Thesis holds, there are also some pockets that are best analyzed as absolutist or relativist). That is not to say that the question of whether we ought to hold that moral reasons are absolute or relative isn’t important. Such a question,

however, may be one to which conceptual analysis of ordinary moral discourse may not provide a determinate answer (even if moral metaphysics or prescriptive ethics may.^v

But while meta-ethical indeterminacy might characterize some parts of ordinary discourse, it might not characterize all of them. The best descriptive analysis of some uses of moral terms might involve robust meta-ethical commitments. Those commitments, however, might not all be uniformly consistent with one side of the traditional meta-ethical debates or the other. Some parts of ordinary moral discourse might be best captured by, say, a relativist analysis, while other parts might be best captured by an absolutist analysis, and there may be no reason for holding that one of these parts is any more conceptually genuine or aberrant than the other. Or so the Variability Thesis — which is the second half of the IV Thesis — maintains.

The Variability Thesis has two components: “different person” variability and “different context” variability. “Different person” variability bids descriptive meta-ethics to attend to the fact that we live in a time and place in which some people say they hold certain meta-ethical positions that other people say they do not. There are some, for instance, who say they believe in a divine command theory of the origin of morality, and there are others who say they believe in a socially constructivist view of the origin of morality. “Different person” variability asks us to take seriously the idea that the thought and language of the former really do involve a commitment to moral absolutism and that the thought and language of the latter really do involve a commitment to moral relativism. I do not mean to suggest that the different first-person reports of these two persons alone imply that such differences in meta-ethical commitment exist. Even though two ordinary people may say they hold incompatible meta-ethical beliefs, it could

still be the case that the best descriptive analysis of the moral discourse of both of them is relativist or that the best analysis of both of them is absolutist. The meta-ethical belief that one of these people says she holds may not be involved in her ordinary moral discourse, even if she thinks it is. But then again it may be. It may be that one person's use of moral terms will be best captured (explained with more elegance, less awkwardness) by an account that includes an absolutist component, and that another person's use of moral terms will be best captured by an account that includes a relativistic component. And the difference between first-person reports of meta-ethical commitments gives a variabilist approach (i.e., an approach that employs the variability thesis) at least some prima facie explanatory advantages over approaches based on the UD assumption. For those who accept the UD assumption must shoulder the explanatory burden of explaining away those first-person reports. They must try to show that a significant number of people have false beliefs about how they think about morality or about how they use moral terms. It's possible that a significant number of people are making errors of this kind. But it seems to me that we should start from the presumption that people have fairly accurate beliefs about how they conceive of things and use terms, and that we should abandon that presumption only if there are strong reasons for doing so. An explanation that does not have to disengage someone's first-person reports of her meta-ethical commitments from her actual moral discourse is — at least prima facie — superior to an explanation that does have to disengage them.

Some might object, however, that meta-ethical variability must be rejected because if people conceived of morality differently, they would not be able to engage in the kind of moral disagreement that is a plainly observable fact of our lives. The

variabilist can respond to this objection by constructing a plausible scenario in which two people who hold different meta-ethical commitments can nonetheless engage in moral disagreement. One of them, let us say, is a committed divine command theorist, and the other is a committed cultural relativist. Now there may very well be some first-order moral issues about which these two people can disagree without manifesting their different meta-ethical commitments. They may disagree, for instance, about whether our society is obligated to provide health care to all its members, but their discussion of this matter may not bring their different views of the nature of morality to the surface. Their situation when disagreeing about health care might be similar to that of two mathematicians who disagree about the ontological status of numbers but whose disagreement about the soundness of a particular mathematical proof does not at all implicate the differences in their meta-mathematical views. Or it might be similar to the situation of two lawyers, one of whom is a natural law theorist and the other of whom is a legal positivist, who disagree about how a particular aspect of tort law applies to a specific case.

But while some moral disagreements, such as that over health care, might not bring the divine command theorist's and cultural relativist's different meta-ethical commitments to the surface, other moral disagreements might. When discussing the status of homosexuality, for instance, their different ways of conceiving of morality may very well become manifest in a way that would (or should) register in descriptive meta-ethical analyses. When discussing the status of homosexuality, the divine command theorist and the cultural relativist may engage in a discussion that implicates conflicting views not merely about whether homosexuality is wrong but also about what wrongness

even consists of. In such a case, we should hold that the divine command theorist and the cultural relativist are engaged not merely in a first-order moral disagreement but also in a prescriptive meta-ethical disagreement, in that each holds that his way of thinking about morality (or his set of meta-ethical commitments) is superior to the other's. They disagree not only about what morality requires but about how one ought to conceive of morality in the first place. Now it seems to me that many moral disputes in our culture do involve meta-ethical disagreements of this kind. It seems to me that when people in our culture argue about issues such as homosexuality, abortion, euthanasia, animal rights, and cloning, different views about the nature of morality itself do play a role. And the presence of such meta-ethical disagreement in ordinary moral discourse constitutes strong evidence for the Variability Thesis.

The variabilist explanation of these disagreements can invoke the Rawlsian distinction between the concept of morality and the conception of it (Rawls 1999, 40-46). Adherents of the UD assumption can be taken to hold without sufficient warrant that commitments to one side or the other of the oft-disputed meta-ethical debates are contained in the concept of morality. The variabilist, in contrast, can maintain that the divine command theorist and the relativist can disagree because they hold the same concept of morality (which concept does not include a determinate uniform commitment on the oft-disputed meta-ethical debates) while also holding different conceptions (where the divine command theorist's conception involves different commitments on the oft-disputed meta-ethical debates from the relativist's conception). The variabilist explanation of these disagreements can also invoke what Horgan and Timmons have called the "variantist" character of certain concepts, according to which a concept can be

“subject to certain kinds of identity-preserving differences in correct usage. One and the same concept can be used by two persons (or by one person, at different times) in ways that are governed by somewhat different semantic standards, while still being the *same* concept” (Horgan and Timmons 2002). On the variantist view of concepts (which is suggested by Putnam’s discussion of momentum [Putnam 1988, 9-10] and Lewis’s discussion of semantic indecision and variation [Lewis 1999]), two people can employ the same concept — and disagree with each other — even though the semantic standards that govern one person’s correct employment of the concept may differ somewhat from the semantic standards that govern the other person’s correct employment of the concept.

The variabilist approach, moreover, does not have to attribute to everyone, or even to most people, consistently determinate meta-ethical commitments (so that even if you don’t accept the account of moral disagreement sketched in the previous three paragraphs there is still reason to take the Variability Thesis seriously). For there is another way in which our moral discourse may be variabilist, and this other way may be more common than the first. This other way is “different context” variability, according to which there are some contexts in which moral terms are used in a manner that is best analyzed as involving one commitment and other contexts in which moral terms are used in a manner that is best analyzed as involving the commitment that has traditionally been taken to be the former’s meta-ethical competitor. “Different context” variability implies that the use of moral terms in some contexts is best captured by, say, a relativist analysis, and that the use of moral terms in other contexts is best captured by an absolutist analysis. And, according to “different context” variability, there may be no principled grounds for holding that the relativist use of moral terms in one context is more

conceptually genuine or aberrant than the absolutist use of moral terms in the other context. “Different context” variability emphasizes the differences between the many situations in which we use moral terms, and raises the possibility that these different uses cannot be explained as all flowing from a concept that is uniformly relativist or absolutist, internalist or externalist, etc. The hypothesis here is that meta-ethical analyses that assume UD cannot describe the variety of the different kinds of jobs to which we put our moral terms as accurately as analyses that are variabilist can.

I believe that the Variability Thesis can offer a plausible explanation both of the attractiveness of opposing 20th century meta-ethical positions and of the philosophical stalemate that seemed to come to exist between them. According to this thesis, parts of ordinary moral discourse really are best analyzed as involving a commitment to, say, relativism. But other parts really are best analyzed as involving a commitment to absolutism. Relativists did an excellent job of explaining the former but were doomed to struggle to accommodate the latter. Absolutists did an excellent job of explaining the latter but were doomed to struggle to accommodate the former. A variabilist analysis, on the other hand, can possess the strengths of both positions without being saddled with either of their weaknesses.

Some might object, however, that a variabilist approach is saddled with a significant weakness that afflicts neither the relativist nor the absolutist, nor any other uniformist analysis. The Variability Thesis implies that different people use moral terms in different ways at different times — that moral terms are sometimes used in a way that applies to all of humanity and other times used in a way that applies only to a particular culture, that moral terms are sometimes used in a way that entails motives and other times

used in a way that leaves the question of motive conceptually open, etc. But, so this objection might go, the result of all this would be a moral cacophony, and the variabilist analysis associated with it would be much less coherent than any of the competing uniformist analyses. Coherence is something we want, however; to the extent that an account lacks coherence it is, all things being equal, inferior to an account that has it. So at least on this score, according to the objection, uniformist accounts are superior to the variabilist one.

In response to this objection, the variabilist should emphasize that the Variability Thesis is a descriptive claim, not a prescriptive one.^{vii} Perhaps it would be better if all our uses of moral terms cohered strongly with each other in a certain way, and perhaps uniformist accounts can show us what that way is. But the prescriptive claim that there are good reasons for everyone always to use moral terms in a certain way is consistent with the descriptive claim that we do not all always use moral terms in that or any other uniformly determinate way. If we are engaged in descriptive meta-ethics, we should not start with the methodological assumption that an analysis that renders our use of moral terms more coherent is superior to an analysis that renders it less coherent; to do so would be to beg the question against the variabilist. That is not to say that the Variability Thesis implies rampant moral incoherence. Many of the different uses of moral terms described by “different context” variability may be seen to cohere with each other quite well once the parameters and assumptions that set the different contexts are made explicit and clearly analyzed. It’s also true, however, that some of the differences — especially those described by “different person” variability — may not admit of this kind of reconciliation. Variabilists will probably have to allow that their account implies that our

moral discourse is sometimes cacophonous. But it seems to me that our moral discourse does have its cacophonous moments and thus that an accurate descriptive meta-ethical analysis should accommodate that fact (even if it would also be good for someone engaged in prescriptive meta-ethics to show us that we would be better off — that we could harmonize what is now cacophonous — if we all always used moral terms in a certain uniform way).

But perhaps the objection mentioned in the preceding paragraph is better expressed in terms of simplicity. The uniformist analyses, so this way of putting the objection might go, explain more phenomena with fewer principles, while the variabilist approach has to make use of both a relativistic conception of moral reasons and an absolutist one, of both an internalist conception of moral judgment and an externalist one, etc. In responding to this form of the objection, variabilists can first point out that the fact that one account has fewer explanans than another does not alone make the former preferable to the latter. The account also has to handle better the explananda. To fully defend their position, however, variabilists would then have to show that their complicated meta-ethical picture does indeed fit the phenomena better — that the more simple, uniformist accounts simply cannot accommodate how people use moral terms as plausibly as a variabilist account can. Of course showing this is not something that can be done by the kind of general, programmatic discussion I've offered up to this point. To show this, one must dive into the details of what the observable phenomena are and present specific arguments for why the variabilist explanation handles them better. I try to make some steps in that direction in Gill 1999, Gill 2008, and in section 4 of this paper, but obviously a great deal of work will remain.

So according to the IV Thesis, there are pockets of ordinary moral discourse that are most accurately analyzed as involving certain meta-ethical commitments. There are other pockets, however, that are most accurately analyzed as involving the commitments that have traditionally been taken to be the former's meta-ethical competitors. And there are wide swaths of ordinary discourse that give us no reason for preferring a descriptive analysis that includes one of the competing commitments over the other. A clean meta-ethical theory, one that takes a singular stand on each of the oft-disputed meta-ethical questions and claims that that stand applies to all the observable moral phenomena, may seem to be more elegant. But the clean meta-ethical theory may do a worse job of describing how we actually use moral terms. Our use of moral terms may itself be very messy.

I should note that I am putting forward the IV Thesis as one that might pertain to *our* moral discourse — to the moral thought and language of Anglo-Americans of the last hundred years or so.^{viii} That is to say, the IV Thesis is supposed to be consistent with the idea that the variability and indeterminacy of our moral discourse is an historically contingent state of affairs. It is not intended to imply that moral discourse could never have included determinate commitments that would vindicate absolutism or relativism, internalism or externalism, etc. Indeed, it seems to me likely that other groups of people in other times and places have talked and thought in ways that are best captured by one or the other of the traditional meta-ethical positions. But our moral discourse may not be not like that, and if it isn't, meta-ethics should accommodate and try to explain this historically contingent fact.^{ix}

This historically contingent variability and indeterminacy of our moral discourse, as implied by the IV Thesis, bears similarity to our discourse about death. When we say someone is dead, do we mean to say that her respiration and heartbeat have permanently ceased (traditional death), or that although her respiration and heartbeat may be maintained by machines her brain has irreversibly ceased functioning (brain death)? Undoubtedly, there have been communities in which virtually everyone almost always had the same determinate understanding of death. But our thought and talk about death is not like that. In some contexts — e.g., when discussing cadaveric organ procurement — people mean by “death” brain death. In other contexts — e.g., when certain religious groups are discussing the end of life — people mean by “death” traditional death. And in still other contexts — e.g., when discussing what our grandchildren will remember about us after we’re gone — we may not be committed one way or the other, and our conversation may fit equally well with either a “brain death” or “traditional death” analysis. The analogy between morality and death should not be pushed too far; there are significant differences between how the two concepts operate. But the idea that the ways we talk and think about death are not uniformly better explained by a “brain death” analysis or a “traditional death” analysis is crucially similar to the idea that the ways we talk and think about morality are not uniformly better explained by a relativist or absolutist analysis, nor by an internalist or externalist analysis, etc.

4. The Internalist-Externalist Debate

Traditionally, (motive) internalists held that there was a necessary connection between the moral judgment that X was right and a motive to do X — a conceptual link between

moral judgment and the motive of the judger. Internalists allowed that at times people thought and said things that looked like they were morally judging that X was right without possessing any motive to do X. Internalists claimed, however, that these seeming moral judgments were conceptually aberrant, secondary, or somehow parasitic on moral judgments that were necessarily connected to motive. (Motive) externalists, in contrast, denied that moral judgments were necessarily connected to motive. They held that a person's morally judging that X was right always left *conceptually open* the question of whether the person had a motive to do X. Externalists acknowledged that there were cases in which the connection between a person's moral judgment and her motive looked to be very intimate indeed. They claimed, however, that this connection was merely a very strong contingent one.^x

The IV Thesis holds that traditional internalists and externalists both got something importantly right about ordinary moral thought and language — and that they both got something importantly wrong. Ordinary people make some moral judgments that internalists correctly analyzed as being conceptually connected to motive, and externalists were wrong to deny that moral terms can be used in ways that involve such a necessary connection. Ordinary people also, however, make some moral judgments that leave conceptually open the question of motive, as the externalists realized, and internalists were wrong to claim that such judgments are in any way conceptually inferior to moral judgments that are necessarily connected to motive.

In this section, I will try to elucidate some ways in which the UD assumption played a role in the 20th century (motive) internalism-externalism debate, and to make the case that a position based on the IV Thesis ought to be considered a serious rival to the

more traditional internalist and externalist positions. I will start by describing the internalism of R.M. Hare, move on to the externalism of David Brink, and then examine the internalism of Michael Smith. Obviously, I will not say enough in this single section to justify the claim that all twentieth century internalist and externalist positions were inadequate. But I do hope that my brief treatment will raise doubts about how the internalism-externalism debate was conducted — doubts that will open the door for consideration of the IV Thesis.

Hare contends that paradigmatic value-judgments are “action-guiding” and “entail imperatives,” and this makes his view a classic example of internalist meta-ethics. He acknowledges, however, that in some contexts, the “standards of values” in play are so “stable” or “ossified” that it is possible to say “You ought to x” or “x is good” and mean by it a purely descriptive, non-action-guiding, non-imperative-entailing statement. Thus, while Hare at times argues that “good” cannot mean “possessing property x” because we would then be unable to use “good” to commend something for possessing x, he also acknowledges that we sometimes do use “good” in a completely descriptive way (Hare, 1964, 112; cf. 124, 147-8). Or as Hare writes with regard to “ought,” “Thus it is possible to say ‘You ought to go and call on the So-and-sos’ meaning by it no value-judgment at all, but simply the descriptive judgment that such an action is required in order to conform to a standard which people in general, or a certain kind of people not specified but well understood, accept. And certainly, if this is the way in which an ‘ought’-sentence is being used, it does not entail an imperative; we can certainly say without contradiction ‘You ought to go and call on the So-and-sos, but don’t’” (Hare, 1964, 164).

Hare also says, however, that the completely descriptive uses of “ought” and “good” are “secondary” while the evaluative uses are “primary” — that the descriptive uses “are not genuine value judgments” (Hare 1964, 164, 118-122, 127, 148, 169-171). The question we have to ask is what warrants Hare in maintaining that “ought” and “good” judgments that are action-guiding and imperative-entailing are conceptually primary or genuine while “ought” and “good” judgments that are not action-guiding and not imperative-entailing are conceptually aberrant or ungentle.

Hare gives four reasons for holding that in order to be “genuine,” a value judgment must be action-guiding or imperative-entailing. He claims, first, that we learn the descriptive meanings of “good” only after we have grasped its action-guiding meaning. As he puts it, “We have knowledge of the evaluative meaning of ‘good’ from our earliest years; but we are constantly learning to use it in new descriptive meanings, as the classes of objects whose virtues we learn to distinguish grow more numerous” (Hare, 1964, 118). Now whether children’s “earliest” understanding of “good” is action-guiding or purely descriptive is an empirical question that I do not think we can satisfactorily answer by relying simply on the kinds of intuition Hare seems to use. (For what it’s worth, it seems quite plausible to me that children first come to think of “good” as denoting a certain class of actions.) But even if Hare is correct in holding that our earliest understanding of certain words is action-guiding and not purely descriptive, that would still not establish that our descriptive uses are conceptually (as opposed to simply chronologically) secondary. Even if (as seems doubtful to me) our descriptive uses come later, they still may be just as conceptually central and genuine as the action-guiding uses that came earlier.^{xii}

Hare claims, secondly, that the action-guiding meaning of “good” is primary because “we can use the evaluative force of the word to *change* the descriptive meaning” (Hare, 1964, 119). Hare is certainly correct that this sort of change does occur: the standards of what counts as a “good motor-car,” for instance, may change from one decade to the next, in which case the evaluative meaning will have stayed the same while the descriptive meaning will have changed. But Hare can rely on these sorts of phenomena to establish the primacy of the action-guiding meaning of “good” only if it is never the case that descriptive meanings stay the same while evaluative meanings change. There has to be a conceptual asymmetry between the two meanings in order to establish Hare’s point. Sometimes, however, it is the descriptive meaning that stays the same while the evaluative meaning changes. Evaluative meaning can sometimes be “emptied” out of a term so that it comes to “be used as a purely descriptive word for designating certain characteristics of the object” (Hare, 1964, 120). So sometimes evaluative meanings remain fixed while descriptive meanings change, but other times descriptive meanings remain fixed while evaluative meanings change. And the symmetry of these phenomena lends more support to a variabilist approach than to any uniformist internalism.

Hare claims, thirdly, that when people use terms such as “ought” and “good” in a way that does not entail imperatives or guide action, they are not even trying to make a genuine value-judgment. In such cases, when only the descriptive meaning is in use, people are using the terms in “an inverted-commas or an ironic” way (see Hare, 1964, 165 and 125). They are consciously insincere, intending merely to parrot or mock the values of other people while signaling their own distance from them. Now, once again,

Hare is correct in pointing out that people do sometimes use “ought” and “good” in ways that are insincere, ironic, or placed in inverted commas. As Hare himself acknowledges, however, (and as I will address in more detail in my discussion of Brink and Smith) noting this fact cannot dispose of all the “awkward cases” (Hare 165). For Hare cannot plausibly maintain that every time people use the terms “ought” and “good” in ways that are not action-guiding and not imperative-entailing, they are being insincere or signaling ironic intentions. In some contexts, people use “ought” and “good” to sincerely and non-ironically describe a “sociological” or “psychological fact,” without their statement’s having any necessary implications for action (Hare, 1964, 167).

So what is Hare’s fourth line of defense against the “awkward cases” that the inverted commas reply cannot explain away? How does he advance his internalist position in the face of sincere, non-ironic “ought” and “good” judgments that do not guide action or entail imperatives? He does so by stipulation. In the end, he seems to allow that he has no argument for the claim that every genuine, sincere, conceptually respectable value-judgment is action-guiding and imperative-entailing. So all he can do is *define* value-judgments as those that guide action and entail imperatives. As he puts it, “My answer to the objection then is, that cases which are alleged to be value-judgements not entailing imperatives will always on examination be found to be cases where what is meant is not [a genuine value-judgment], but [a statement of sociological or psychological fact]. This contention is, of course, impossible to prove or even to render plausible, unless we know when we are to count a judgement [a genuine value-judgment]; but I propose to get over this difficulty in the only possible way, by making it a matter of definition. I propose to say that the test, whether someone is using the

judgement ‘I ought to do X’ as a value-judgement or not is, ‘Does he or does he not recognize that if he assents to the judgement, he must also assent to the command ‘Let me do X’? Thus I am not here claiming to prove anything substantial about the way in which we use language; I am merely suggesting a terminology which, if applied to the study of moral language, will, I am satisfied, prove illuminating” (Hare, 1964, 168-169; cf. 164). This statement of Hare’s more or less concedes the truth of what I am calling the Variability Thesis. For Hare acknowledges here that within ordinary discourse we can find some “ought”-judgments that are best given an internalist analysis and other “ought”-judgments that are best given an externalist analysis, and that there is no principled, non-stipulative reason for holding that one of these kinds of judgments is conceptually more genuine than the other. Hare chooses to focus his attention on the first kind of judgment because he thinks it will “prove illuminating.” But he gives no reason to deny that someone else could focus attention on the second kind of judgment and also produce illuminating results. He gives no reason to deny that it would be philosophically worthwhile to determine the kinds of psychological and sociological facts some “ought”-judgments try to describe.

Moreover, Hare also more or less concedes the truth of what I am calling the Indeterminacy Thesis. For he says that “it is usually impossible for an ordinary person, untrained in logical subtleties, to ask or to answer the question” that would determine whether his “ought” or “good” judgment is a real value-judgment or not (Hare, 1964, 167). Or as he goes on to put it, “The situation is very similar to that of the scientist who is asked by the logician ‘Is your statement that phosphorus melts at 44 degrees Celsius analytic or synthetic; if you found a substance which was in other respects just like

phosphorus, but which melted at another temperature, would you say ‘It isn’t really phosphorus’ or would you say ‘Then after all some phosphorus melts at other temperatures’? The scientist might well ... answer ‘I don’t know; I haven’t yet come across the case which would make me decide this question; I have got better things to worry about’. Similarly, the ordinary person, making moral decisions on the basis of his accepted principles, very rarely has to ask himself the question that we have just asked. So long as his value-judgments correspond with the accepted standards and with his own feelings, he does not have to decide which he is saying, because, as we might put it, all three are as yet for him materially equivalent” (Hare, 1964, 167-8; cf. 172). Here Hare seems to be allowing that in certain contexts, there is no more reason to analyze a person’s judgment as guiding action and entailing an imperative or as not guiding action and not entailing an imperative.

So there are passages in Hare that seem to concede both parts of the IV Thesis. Hare seems to acknowledge that his internalist-prescriptivist account of moral language captures part of our ordinary discourse, but that an externalist-descriptivist account captures other parts, and that still other parts are no better captured either by an internalist-prescriptivist or by an externalist-descriptivist analysis. But the limitations of Hare’s position are not always readily apparent on the surface of his text; his official position seems to be that his prescriptivist account captures the essence of moral language, and that any uses of moral terms that do not fit his account are so peculiar or conceptually aberrant that they are hardly of any importance. And the meta-ethicists who have followed Hare have also suggested that moral terms essentially have a single, determinate conceptual shape.

Consider, for instance, Brink's arguments against Hare and for externalism. Brink says that we can conceive of "amoralists," or people who make genuine moral judgments without having any motive to act in accord with them. He gives as examples sociopaths, "Plato's Thrasymachus, Hobbes's Fool, and Mike, the conman in David Mamet's film *House of Games*" (Brink, 1989, 47). These figures are *global* amoralists, in that they seem to lack motives to act in accord with morality in general. We can also add to Brink's argument *local* amoralists, people who lack motives to perform particular acts they judge to be required by morality. Local amoralists are motivated to perform many of the actions they judge to be morally required, but do not possess motives for others. I am thinking here of people like a student I knew who seemed to care about doing many of the things he judged moral (he seemed to care about not causing unnecessary pain and about being loyal to his friends) but who said about cheating on an exam, "Sure it's wrong, but so what? If I know I can get away with it, why should I care?" and of an identity-thief on a television ad who said, "I know it's wrong, but I don't care." An internalist impressed by Hare's "inverted commas" response might try to explain away cases of global and local amoralism by claiming that the people in question are being insincere or ironic. Brink argues, however, that this response does not take these cases "seriously enough" (Brink, 1989, 47). For it does not seem that when amoralists use moral terms they are all always being insincere or ironic. Nor does there seem to be a principled, non-question-begging reason for claiming that they are confused about how to use moral terms or "mistaken about what morality requires" (Brink, 1989, 47-48).

I think Brink is right in claiming that we can conceive of a person who makes a genuine moral judgment without possessing any motive to act in accord with it. Indeed, I think there are plenty of actual cases of this. And such cases do undermine the internalist view that every moral judgment is necessarily connected to a corresponding motive. If internalism is the claim that necessarily whenever a person makes a genuine moral judgment she possesses a corresponding motive, then amoralist counterexamples do defeat it.

Brink goes on to draw a stronger externalist conclusion from the amoralist cases, however. He maintains that the connection between moral judgment and motive is *always* contingent — that moral judgments on their own are never necessarily connected to motive.^{xiii} But the case of the amoralist does not establish this conclusion. Or rather, the case of the amoralist establishes the externalist conclusion only if we assume UD, which implies that internalism and externalism are mutually exclusive and jointly exhaustive options for an account of the relationship between moral judgment and motive. But the UD assumption is something we should question here. For it could be that while in certain contexts a moral judgment is not necessarily connected to a corresponding motive, in other contexts a moral judgment is necessarily connected to a corresponding motive. That amoralists use moral terms a certain way does not mean all people always use moral terms in that way. Our uses of moral terms might vary in a way that the Variability Thesis captures better than either traditional internalism or traditional externalism.

To gain a clearer picture of the variabilist possibility, let's consider two cases. The first case is the student who believes that cheating on an exam is wrong but does not

care about its wrongness. We can imagine having a conversation with this student in which we try to arouse in him a motive not to cheat. But this conversation could be held against the backdrop of agreement between us on the immorality of cheating. If we do succeed in getting him not to cheat even when he could have got away with it, the best account of what has happened may be not that we have gotten him to change his moral judgment but that we have aroused in him a motive to do something that he had thought all along was morally correct. An externalist analysis might fit this case better than an internalist one.

But now imagine a different case — a person who has found out that a friend in the hospital is feeling depressed and who has decided that he should visit. In this case, we might find it unimaginable, or at least very puzzling, for the person to be completely lacking the motive to visit his friend, or for him to say all in one breath, “Sure, I ought to visit him, but I just don’t care to.” The context in which we are speaking might make it appropriate to take the person’s judgment about what he ought to do to be necessarily connected to a motive to do it. An internalist analysis might fit this case better than an externalist one.

To bring out the difference between these two cases, we can use the kind of questions Hare used to distinguish different types of “ought”-judgments. Imagine that after our discussion had aroused in the student a motive not to cheat, we had asked him, “If another person in the same situation as yours didn’t care at all not to cheat, would you take that person to be in disagreement with you about what is morally correct in this situation?” Given what the student said at the beginning of the discussion, it seems that he might well answer No to that question. The student seems to believe that one’s

judgment of the immorality of cheating on a test can be separated from one's motives about how to behave during the test. Now imagine, however, we asked the person of the second case, "If another person in the same situation as yours didn't care at all about visiting his friend in the hospital, would you take that person to be in disagreement with you about the right thing to do in this situation?" To this question it's entirely possible that the person would answer Yes. The person might think that one's judgment about whether one ought to visit a sick friend cannot be separated from one's motives about how to behave toward a sick friend.^{xiv}

Now if we take the student who doesn't care about the wrongness of cheating on an exam and the friend who is motivated to visit the hospital to be one and the same person, then we have a pair of cases that exemplify "different context" variability. For these cases would then describe a single person whose discourse is best given an externalist analysis when he is talking about how morality applies to his role as a student, and whose discourse is best given an internalist analysis when he is talking about how morality applies to his role as a friend. But we can also imagine cases that support "different person" variability. We can, that is, imagine a global internalist, someone who always uses moral language in a way that is best analyzed as involving a necessary link between his judgment of what is moral and his motives; perhaps Hare himself was this type of person. And we can also imagine a global externalist, someone who always uses moral language in a way that is best analyzed as leaving open the question of whether he has a motive to do what he judges to be moral; maybe Thrasymachus and sociopaths are examples of this type.

In addition to the Variability Thesis, we can also apply the Indeterminacy Thesis to this debate. Imagine, for instance, the members of a medical ethics committee are engaged in conversation about whether the hospital should continue or discontinue aggressive treatment on an elderly patient with dementia. The professional context will play a role in how this conversation proceeds. None of the committee members will, for example, raise the possibility of pursuing a course of action that is in manifest violation of explicit standards of medical codes of ethics. The context will not allow for that kind of suggestion. Indeed, if one of the members did raise the possibility of doing something manifestly in conflict with codes of medical ethics, it's likely she would be seen as changing the subject. Now let us say that the committee eventually comes to the unanimous conclusion that they ought to recommend discontinuation of treatment. Our question is whether each committee member's judgment that the committee ought to recommend discontinuation of treatment is conceptually linked to a motive to make that recommendation. We can imagine an externalist committee member, someone who takes her contribution to the committee's deliberations to concern only the working out of the implications of professional standards and not to necessarily involve her personal motives at all. And we can also imagine an internalist committee member, someone who is deeply personally committed to the principles at the center of the professional standards and thus whose contributions to the committee's deliberations are all necessarily connected to his personal practical commitments. There may, however, also be members whose participation on the committee does not give us any reason to prefer an internalist or externalist analysis of their moral judgment. Such members, we can suppose, possess motives to recommend the discontinuation of treatment. But there may

be nothing in their discourse that gives us any reason to prefer either the view that their possessing these motives is conceptually linked to their judgment that discontinuation is the right course of action or the view that the question of whether they possess these motives is left conceptually open by their judgment that discontinuation is the right course of action. And if we asked them which view better captures their moral judgment, they might say the same thing as Hare's scientist: "I don't know; I haven't yet come across the case which would make me decide this question; I have got better things to worry about."

Let us now turn Michael Smith's internalist meta-ethics. Do his arguments push us away from the meta-ethical variability and indeterminacy I have been proposing?

In response to Brink's amoralist challenge, Smith maintains that amoralists "do not really make moral judgments" after all (Smith, 1994, 68). They may "try to make moral judgments," but if they possess no corresponding motive they "fail" to do so. According to Smith, anyone who uses moral language in a way that has no practical implications is placing her moral judgments in inverted-commas, whether she realizes it or not. So Smith's position implies that a person may think she's making a genuine moral judgment, that the people to whom she's speaking may think she's making a genuine moral judgment, but that nonetheless she is not making a genuine moral judgment. This is an explanatorily expensive implication. All else being equal, a position that does not have to attribute these kinds of errors to ordinary users of moral terms is preferable to a position that does. Unless there are powerful opposing considerations, we should suppose that people know what they're saying, and that they succeed in making the kinds of judgments they are trying to make. Of course, Smith

believes that an externalist position is saddled with implications that are even more costly than having to attribute to some people unconscious inverted-commas uses of moral terms. He may be right about that. And if internalism and externalism were mutually exclusive and jointly exhaustive options — i.e., if we accepted the UD assumption — then externalism’s greater cost would be enough to establish that we ought to opt for internalism. But if we question the UD assumption, then a third option becomes available: the variabilist alternative. And by distinguishing the meta-ethical implications of one pocket of moral discourse from those in another, the variabilist may be able to avoid the unpalatable implications of both internalism and externalism.

Let us now examine the details of Smith’s argument for his internalist position. He begins his argument by noting what he takes to be a plainly observable fact: “a *change in motivation* follows reliably in the wake of a *change in moral judgement*, at least in the good and strong-willed person” (Smith, 1994, 71). He shows that internalism explains this fact by attributing to “good and strong-willed” people a “non-derivative” concern to do what they think they ought to do, “where this is read *de re* and not *de dicto*,” while externalism explains it by attributing to most people a “derivative” or “self-consciously *moral* motive” to do what they think they ought to do, where this is read *de dicto* and not *de re* (Smith, 1994, 73-74). Smith then contends that the internalist explanation of the reliability of a change in motivation following a change in moral judgment is clearly superior to the externalist one. For “commonsense tells us that if good people judge it right to be honest, or right to care for their children and friends and fellows, or right for people to get what they deserve, then they care non-derivatively about these things. Good people care non-derivatively about honesty, the weal and woe

of their children and friends, the well-being of their fellows, people getting what they deserve, justice, equality, and the like, not just one thing: doing what they believe to be right, where this is read *de dicto* and not *de re*. Indeed, commonsense tells us that being so motivated is a fetish or moral vice, not the one and only moral virtue” (Smith, 1994, 75).

Now there are certainly cases in which the attribution of a non-derivative, *de re* concern better explains a change in moral motivation than the attribution of a self-consciously moral concern to do what is right *de dicto*. Suppose, for instance, that I form the judgment that I ought to visit a friend in the hospital, or that I ought to spend more time with my children, or that I ought to do more to combat racism. And suppose that in the wake of that judgment I develop a motive to visit my sick friend, or spend more time with my children, or combat racism. It is very plausible to hold that, if I am a good and strong-willed person, my motive will be a non-derivative, *de re* concern to visit my friend, spend more time with my children, or combat racism, and not simply the self-consciously moral motive to do what is right *de dicto*.^{xv}

There are other cases, however, that are better explained by attributing to the agent a self-consciously moral motive to do what is right *de dicto*. Suppose, for instance, that a car has needlessly parked exceedingly close to mine, and that as I pull out I nick it, leaving a very small dent. I consider driving on but after some reflection come to the conclusion that the right thing to do is to leave a note. Or suppose that after waiting at the Wal-Mart checkout line for an unforgivable amount of time I am given ten dollars too much in change. I consider keeping the extra change but on reflection decide that the right thing to do is give it back. Now suppose that in the wake of each of these

judgments I develop the motive to do what I have judged to be morally required. Which is the better account: that I possess a non-derivative, *de re* concern to leave the note and return the ten dollars, or that my motive to leave the note and return the ten dollars is derived from a motive to do the right thing *de dicto*? The second option seems to me to be at least as plausible as the first, and probably more so.

Perhaps Smith would point out that his “non-derivative *de re* concern” explanation applies only to “good and strong-willed” people. I don’t think, however, that commonsense implies that someone is bad or weak-willed if her motive to leave the note or return the extra change is derived from a self-consciously moral concern to do the right thing *de dicto*. Nor does commonsense imply that such a person’s moral motives are peculiar or unhealthy, or that she is morally fetishistic. If she visited her friend in the hospital or spent more time with her children only because of a self-consciously moral motive to do the right thing *de dicto*, we probably would think something like that was wrong with her. But leaving the note and returning the extra change are different from visiting a friend. The connection between motive and moral judgment in the one case may differ from the others. The most accurate analysis might be one that can explain that difference without having to hold that in the nicked car and Wal-Mart cases the person involved is bad, weak-willed, or morally fetishistic.

Furthermore, because such an analysis would be variabilist, it would not force us to attribute to the person in the nicked car and Wal-Mart cases a use of moral terms that is ironic, “merely conventional,” unconsciously placed in inverted commas, or otherwise conceptually aberrant. This is all to the good, as the moral judgments in these cases seem to be about as straightforward and ordinary as moral judgments get. That is not to say

that the moral judgment that I ought to visit a sick friend in the hospital or spend more time with my children is non-straightforward or out-of-the-ordinary. It is to say, rather, that a variabilist account — one that can count all those judgments as genuine and ordinary while also providing an explanation of the connection between motive and moral judgment in one set of cases that is different from the connection in the other set — seems to capture commonsense better than an uniformly internalist or externalist one.

There are, moreover, some ordinary cases in which the most accurate analysis of a change in motivation following in the wake of a change in moral judgment might be the indeterminist one, according to which we have no reason to prefer the internalist or the externalist account. Consider, for instance, a group of introductory philosophy students who, after participating in a class on Singer’s “Famine, Affluence, and Morality,” come to accept the judgment that the morally right thing to do is to give more money to famine relief. Some of the students may make this judgment in a way that must be interpreted as leaving open the possibility that they possess no motive to give, and some of them may make this judgment in a way that makes it legitimate to infer that they possess a non-derivative motive to give. Other students, however, may make the judgment in a way that gives us no reason to prefer either an internalist analysis or an externalist one.

There are, of course, many more things that can be said for both the internalist and externalist positions. I am not claiming that my brief treatment conclusively settles any issues. I hope, though, that I’ve shown that a number of influential arguments in the internalism-externalism debate have assumed UD, and that alternative views based on the

IV Thesis deserve to be considered serious alternatives to the more uniformist positions that dominated the 20th century debate.

The IV Thesis would gain even more plausibility if I could show that other 20th century meta-ethical debates — such as the debates between absolutists and relativists, between expressivists and descriptivists, and between error theorists and non-error theorists — have also been explanatorily hamstrung by the UD assumption, and that approaches based on the IV Thesis look to have at least some *prima facie* explanatory advantages. I have begun this work in Gill 1999 and Gill 2008. But for now, I can only point to that projected work and say that the IV Thesis will gain plausibility to the extent that such work proves fruitful.

5. Implications for Meta-ethics

I want to conclude by briefly discussing the implications for descriptive meta-ethics — i.e., for investigations of how ordinary people use moral terms — of taking the IV Thesis to be a serious explanatory rival to the UD assumption.

If we leave open the possibility that our concept of morality is variable and indeterminate, descriptive meta-ethics will have to involve much more empirical investigation than it typically did in the 20th century. As I noted in section 2, many of the 20th century's meta-ethical analyses were based on a mere handful of examples of ordinary usage. A few illustrative anecdotes were often taken to be a sufficient empirical starting point for meta-ethical inquiry. And if the UD assumption is true this procedure is warranted. For if ordinary moral discourse is based on a concept unitary and determinate enough to imply definitive answers to the oft-disputed meta-ethical questions, then the

correct analysis of any one handful of examples will also be correct for every other handful. Indeed, the fact that the UD assumption legitimizes basing a meta-ethical analysis on an empirically slender base constitutes further evidence of the assumption's role in 20th century meta-ethics. If, however, we keep open the question of how unitary and determinate our concept of morality is when doing descriptive meta-ethics — if we take the IV Thesis to be a serious rival to the UD assumption — we will have to give much more time and effort than most 20th century meta-ethicists did to the gathering of data. For if we take seriously the possibility that our concept of morality is multifarious, then one handful of examples may have meta-ethical implications that another handful will not.

One way of gathering data is to pay close attention to moral language and thought as it occurs in specific real-world contexts. The practice of professions, such as medicine and business, is one kind of context to examine. Discussion of particular topics, such as public discourse about the moral grounding and implications of the U.S. Constitution, might be another. Such studies, which would require tools of sociological and ethnological research as well as philosophical analysis, could reveal a great deal about how we talk and speak in certain morally significant situations, and such results could be of great assistance in improving our moral talk and thought in those situations. The mistake to guard against would be to apply conclusions based on data from one context to other contexts which may set different constraints on the use of moral terms.

Another way of gathering descriptive meta-ethical data is to test persons' moral reactions under experimentally responsible conditions. As Stich and Doris put it, "A credible philosophical methodology of *thought* experiments must be supplemented by ...

systematic investigation with *actual* experiments” (Stich and Doris, forthcoming).^{xvi} In addition to the kind of empirical work Stich and Doris recommend, meta-ethicists can also examine the social and linguistic mechanisms that effect moves from one conversational context to another, as well as the moral principles that operate in the different kinds of contexts we inhabit. Particular contexts may very well impose robust restrictions on our moral language and thought, and meta-ethicists can use the tools of philosophy of language to elucidate such things. The mistake to guard against, once again, would be to generalize from the restrictions in one context to blanket claims about the concept of morality as a whole.^{xviii}

Let me close by noting that although I have maintained that descriptive meta-ethics may not give us good reason to adopt one or the other of the traditional meta-ethical positions, I do not mean to imply that there are no good reasons at all for adopting one meta-ethical position or another. There may be no conceptual bar to using moral terms in a wide variety of ways, but there may nonetheless be good reasons for preferring certain meta-ethical commitments to others. For there may be good reasons for thinking that certain meta-ethical commitments fit better with the best ontological or scientific world-view, or that a certain set of meta-ethical commitments constitutes a more rationally coherent whole than other sets, or that certain meta-ethical commitments better promote first-order moral values. These sorts of reasons for preferring certain meta-ethical commitments to others are *prescriptive*. They are reasons — based on metaphysical, epistemological, or first-order normative grounds — for thinking that it would be *better* if people used moral terms in a certain way. But of course the fact that it would be better if people used moral terms in a certain way does not imply that they *do*

use moral terms in that way. Some of us at least some of the time may say and think things that conflict with the best ontology and science. We may be guilty of internal incoherence. We may commit ourselves to morally disastrous views. And if some of us do this some of the time, it should be no surprise to find evidence of it in ordinary moral discourse. Indeed, it would be surprising to find that the best description of ordinary moral discourse uncovers no trace of these kinds of mistakes.

I myself think there are good reasons to adopt certain meta-ethical commitments and to reject others. To make that case, however, I would have to present prescriptive arguments (based on metaphysics, epistemology, or first-order normative views). If my task is descriptive meta-ethics, I must remain neutral toward the wide array of uses of moral terms that is present in our miscellaneous moral discourse.

Anscombe, G.E.M. (1958) "Modern Moral Philosophy." *Philosophy* 33: 1-19.

Brink, David O. (1997) "Moral Motivation." *Ethics* 107: 4-32.

Brink, David O. (1989) *Moral Realism and the Foundations of Ethics*. Cambridge: Cambridge University Press.

Darwall, Stephen. (1995) *The British Moralists and the Internal 'Ought.'* Cambridge: Cambridge University Press.

Doris, John and Stich, Stephen. (Forthcoming) "As a Matter of Fact: Empirical Perspectives of Ethics."

Falk, W.D. (1986) "Morality, self, and others" in *Ought, Reasons, and Morality*. Ithaca: Cornell University Press.

Gibbard, Allan. (2003) *Thinking How to Live*. Cambridge, MA: Harvard University Press.

Gill, Michael B. (1999). "Relativism and the Concept of Morality." *Journal of Value Inquiry*.

Gill, Michael B. (2008). "Meta-Ethical Variability, Incoherence, and Error," in *Moral Psychology: The Cognitive Science of Morality*. Edited by Walter Sinnott-Armstrong. MIT Press.

Hare, R.M. (1964) *The Language of Morals*. New York: Oxford University Press.

Horgan, Terry and Timmons, Mark. (2002) "Conceptual Relativity and Metaphysical Realism." *Philosophical Issues* 12: 219-250.

Hume, David. (2000) *A Treatise of Human Nature*. Edited by Norton and Norton. Oxford: Oxford University Press.

Lewis, David. (2000) "Dispositional theories of value" in *Papers in Ethics and Social Philosophy*: 68-94. Cambridge: Cambridge University Press.

Lewis, David. (1999). "Reduction of mind" in *Papers in Metaphysics and Epistemology*: 291-324. Cambridge: Cambridge University Press.

Loeb, Don (forthcoming). "Moral Incoherentism: How to Pull a Metaphysical Rabbit out of a Semantic Hat."

MacIntyre, Alasdair. (1981). *After Virtue*. Notre Dame: University of Notre Dame Press.

Mackie, J.L. (1977). *Ethics: Inventing Right and Wrong*. New York: Penguin.

Nagel, Thomas. (1979). "The fragmentation of value" in *Mortal Questions*: 128-141.
Cambridge: Cambridge University Press.

Nichols, Shaun. (2004) *Sentimental Rules: On the Natural Foundations of Moral Judgment*. New York: Oxford University Press.

Putnam, Hilary. (1988) *Representation and Reality*. Cambridge, MA: MIT Press.

Rawls, John. (1999) *A Theory of Justice*. Cambridge, MA: Harvard University Press.

Scanlon, T.M. (1998) *What We Owe to Each Other*. Cambridge, MA: Harvard University Press.

Scanlon, T.M. (1982) "Contractualism and Utilitarianism" in *Utilitarianism and Beyond*: 103-128. Edited by Sen, Amartya and Williams, Bernard. Cambridge: Cambridge University Press.

Schneewind, J. B. (1998) *The Invention of Autonomy*. Cambridge: Cambridge University Press.

Smith, Michael. (1994) *The Moral Problem*. Oxford UK & Cambridge USA: Blackwell.

Stich, Stephen and Weinberg, Jonathan M. (2002) "Jackson's empirical assumptions."

Phenomenology and Philosophical Research 62: 637-643.

Stich, Stephen and Ravenscroft, Ian. (1994) "What is Folk Psychology?" *Cognition* 50:

447-468.

Stich, Stephen. (1993). "Moral Philosophy and Mental Representation," in *The Origin*

of Values (eds. Hechter, Nadel, Michod): 215-228. New York: Aldine de Gruyter.

Timmons, Mark. (1998). *Morality without Foundations*. Oxford: Oxford University

Press.

Williams, Bernard. (1985) *Ethics and the Limits of Philosophy*. Cambridge, MA:

Harvard University Press.

Wittgenstein, Ludwig. (1958) *Philosophical Investigations*. Oxford: Basil Blackwell.

ⁱⁱ There are a number of different topics that can be called “meta-ethical.” I will focus here on only one of them, a topic that can be called “descriptive meta-ethics.” The task of someone engaged in descriptive meta-ethics is to provide the best analysis of the ordinary uses of moral terms. The descriptive task is different from the ontological task of determining what if any moral properties actually exist. The descriptive task is also different from the prescriptive task of showing that we ought to use moral terms in a certain way. The mid-twentieth century meta-ethicists I am concerned with here were concerned to answer the descriptive question — even if the early 20th century theorists such as Moore, Prichard, and Ross were more concerned with the ontological meta-ethical questions, and early 21st century theorists have become more concerned with the ontological questions as well. Stevenson, for instance, says that his work is concerned to analyze “the judgments of the ordinary man as he finishes reading the morning’s newspaper” (Stevenson, *Facts and Values*, v) and that his meta-ethical “conclusions are based upon observation of ethical discussions in daily life, and can be clarified and tested only by turning to that source” (Stevenson, *Ethics and Language*, 13). Hare compares his task to that of a descriptive grammarian (see Hare, 1964, i and 4) and he says explicitly that he is giving an account of moral terms as they *are* used, not as they *might* be used (see Hare, 1964, 92). Mackie makes it clear that he is trying to explain “ordinary thought” or what the “ordinary user of moral language means to say” (Mackie, 1977, 31-33). And Smith says that he intends to give an account of the features of morality “that

are manifest in ordinary moral practice as it is engaged in by ordinary folk” (Smith, 1994, 5). Brink is not always as clearly descriptivist as Hare, Mackie, or Smith, but he too says that his account fits better than its rivals with “commonsense moral thinking” (Brink, 1989, 37; although Brink also says that this way of putting it is “perhaps a little misleading”). Meta-ethicists who believe that there is a blanket Yes or No answer to the question of whether commonsense morality is in error (see Mackie, 1977, 48-49; Smith, 1994, 4-13; Brink, 1989, 34) will have to be particularly concerned to keep their conceptual task separate from the prescriptivist one. For the only way to answer the question of whether commonsense morality as a whole is in error is to first give an account of commonsense morality that does not prejudge whether or not it is erroneous.

ⁱⁱⁱ I do not claim that what I am calling the IV thesis is original. The quotations from Lewis 2000 and Falk 1986 at the beginning of this paper show that they had the same ideas. More recently, Nichols has raised the “possibility that there might ... be considerable variation in intuitions about moral requirements. Thus ... it is a substantive empirical assumption that there is a stable and cross-culturally uniform set of intuitions or platitudes that comprise *the* folk concept of morality” (Nichols, 2004, 75; cf. 93, 169-170). Loeb makes many similar points in forthcoming work. The two major historical antecedents of the IV Thesis are Wittgenstein and Hume. Just as Wittgenstein maintains that while all words might look alike they may nonetheless have different functions (Wittgenstein, 1958, 6-7), so I want to maintain that while moral terms might look alike they nonetheless have different functions. And just as Wittgenstein maintains that our concept of a game is “blurred” and best captured only by “family resemblance”

(Wittgenstein, 1958, 32-36), so I want to maintain that our concept of morality may be blurred and best captured only by family resemblance. I discuss Hume in footnotes below.

^{iv} In this section I will speak of internalism, externalism, absolutism, relativism and the like in very general ways in order to paint in broad brushstrokes the meta-ethical picture I want to advance. I acknowledge, however, that each of these meta-ethical positions admits of many distinctions that I do not take note of. When speaking of internalism, for instance, we can distinguish between motive-internalism and reason-internalism (and of course there are different kinds of motive-internalism and reason-internalism). When speaking of relativism, we can distinguish between cultural or constructivist relativism and individualist or personal relativism (and there are different kinds of each of these views as well). And so on. My suggestion is that if we were to examine in detail any of the micro-debates among these more refined meta-ethical positions, we would still find that each side is afflicted by awkwardnesses that would make a position based on the IV Thesis a legitimate explanatory rival. But of course the details do matter. And without having done the fine grain work on the micro-debates, I cannot claim to have shown that none of the more refined meta-ethical positions can succeed at giving an elegant uniformist-determinist account.

^v We can clarify the Indeterminacy Thesis by applying to meta-ethical debates the distinction between methodologically external and methodologically internal accounts, which has been helpfully deployed in discussions of grammatical theory and folk psychology (see Stich and Ravenscroft 1994). A methodologically external meta-ethical account tries to capture ordinary moral discourse but does not take a stand on the ideas

ordinary people have when they use moral terms. A methodologically external account claims that certain meta-ethical positions do a better job of explaining our use of moral terms than other meta-ethical positions, but it does not claim to have described the internal mental representations of ordinary users of moral terms. The things that make up a methodologically external account aren't necessarily "in the head." A methodologically internal account, in contrast, tries not only to capture the observable phenomena of how moral terms are used but also to describe the ideas of ordinary users of moral terms. A methodologically internal account does try to capture internally represented moral commitments. It does try to say something about what's "in the head." I intend the Indeterminacy Thesis to be a rival to both methodologically external and methodologically internal meta-ethical accounts that involve the UD assumption. If we allow that the Indeterminacy Thesis is a legitimate explanatory rival to methodologically external accounts, we should keep open the possibility that some parts of ordinary moral discourse seriously underdetermine the question that separates, say, meta-ethical relativists from meta-ethical absolutists. Some parts of ordinary moral discourse may be no better explained by a relativist analysis than by an absolutist one; the relativist and absolutist analyses might fit certain parts of ordinary moral discourse equally well. If we allow that the Indeterminacy Thesis is a legitimate explanatory rival to methodologically internal accounts, we should keep open the possibility that the mental representations of some ordinary users of moral terms are neither relativistic nor absolutist. It may be the case that what's "in the head" of some ordinary users of moral terms commits them neither to relativism nor to absolutism. (Note that methodologically internal indeterminacy is not exactly the same as methodologically external indeterminacy. In the

methodologically internal case, the suggestion is what is “in the head” of some people may not commit them to, say, either relativism or to absolutism. In the methodologically external case, the suggestion is that there may be no fact as to whether some parts of ordinary moral discourse are better explained by a relativist or absolutist analysis. In each case, however, the indeterminacy in question implies that we have no good reason on purely descriptive or analytic grounds to favor one side of the traditional 20th century relativist-absolutist debate over the other.)

^{vii} In section 4, I will say more about the relationship between the descriptive and the prescriptive meta-ethical tasks. The important point to keep in mind is that the meta-ethical approaches I am trying to raise doubt about here are only those that purport to be descriptive of our ordinary moral discourse. See footnote 2 for more on this matter.

^{viii} Note that I am maintaining that meta-ethical variability and indeterminacy exist *within* our culture. I think it’s likely that even more variation exists between different cultures. That is to say, I suspect the UD assumption is even less plausible when applied to the moral discourse of humanity as a whole than when applied to the moral discourse of a single culture.

^{ix} To expand on this idea, we might call attention to cultural and historical factors that lend plausibility to the idea that the phenomena of our ordinary moral discourse are not unitary and determinate enough to settle the oft-disputed meta-ethical debates. Along these lines, one should acknowledge that there may exist elsewhere relatively homogenous and isolated groups of people whose moral discourse does imply a very robustly determinate moral concept. But twentieth century Britons and Americans constitute a very different kind of group. Our moral language and thought has roots in

numerous disparate sources, such as Hellenistic thought, the Judeo-Christian tradition (which is of course not a single position but a constellation of ideas, some of which are incompatible with each other), various branches of the sciences and social sciences, and some non-western cultures. The result of this interaction of multiply diverse sources is an unruly conglomeration. We should not be surprised, consequently, to find in our ordinary discourse some things that suggest one answer to a meta-ethical question, and other things that suggest the contrary answer. For the former may have their roots in ideas developed within one way of thinking about morality, while the latter may have their roots in ideas developed within a very different way of thinking about morality.

That our historical and cultural heritage has left us with a moral hodgepodge is an idea raised in Anscombe 1958 and explored in detail in McIntyre 1977 and Williams 1985. (Although I agree with much of the historical pictures of these works, I disagree with the dire conclusions MacIntyre [and perhaps Anscombe] draw. I think our moral discourse is not in such terrible shape — that the miscellaneous character of it is not cause for the great alarm that MacIntyre sounds.) Other works in the history of philosophy, such as Darwall 1995 and Schneewind 1998, reveal how features that may now appear to us to be fundamental to morality evolved over time. It seems to me that the more attention we pay to these sorts of accounts of the historical development of our moral ideas, the more plausible a rejection of the UD assumption will begin to look. (For an ahistorical treatment of this issue, see Nagal's "Fragmentation of Virtue").

One historical source of meta-ethical variation that I've tried to describe in *The British Moralists on Human Nature and the Birth of Secular Ethics* is differing attitudes toward original sin or the Fall of Man. A conspicuous strand of Christian Protestantism

emphasizes original sin in a manner that seems to be incompatible with the internalist idea that there is a necessary connection between a person's motivations and what she ought to do. A conspicuous strand of Greek eudemonism avoids original sin, emphasizing instead the possibility of a unified life in a way that seems to imply the internalist idea. But both the Protestant strand and the Greek strand — which together imply incompatible positions on the internalist debate — make their presence felt in our moral discourse.

^x If internalism is defined as the claim that necessarily whenever A judges that x is moral, A possesses a motive to x; and if externalism is defined as simply the negation of internalism; then, of course, a single case in which a person can judge that x is moral without possessing a motive to x will mean victory for externalism. Now the Variability Thesis does hold that there can be cases (such as those of the amoralist) in which a person judges that x is moral without possessing a motive to x. So if externalism is defined as simply the negation of internalism, then the Variability Thesis is in the externalist camp. But it seems to me that externalists such as Brink advance a stronger claim. It seems to me that Brink holds that the state of mind that is a moral judgment *never* motivates on its own, not just that motivation is *sometimes* external to a moral judgment but that it is *always* so, that moral judgment (i.e., *every* moral judgment) is the kind of thing that cannot motivate unless combined with some other mental state, that the connection between moral judgment and motivation is always mediated. An old-fashioned way of expressing this kind of strong externalism is to say that morality stands in need of a *sanction*. This kind of strong externalist view seems to me to be implied by Brink's claim that externalism "makes the motivational force of moral judgment and moral belief

a matter of contingent psychological fact, depending on both the content of people's moral views and their attitudes and desires (Brink, 1989, 46) and "[M]oral motivation depends upon our beliefs about the rational authority of morality" (Brink, 1997, 30). And what I am claiming is that this stronger kind of externalism is just as unwarranted as internalism. The Variability Thesis holds that some of the states of mind that can legitimately be called moral judgments do on their own motivate, and that some of the states of mind that can legitimately be called moral judgments do not on their own motivate.

^{xii} For discussion of how different parts of our moral thought and language could have different origins, see Hume, 2000, 307-322 and 367-378.

^{xiii} See footnote 11.

^{xiv} The difference between the case of visiting a friend in the hospital and the case of leaving a note on a nicked car is something that Hume was well aware of. Hume distinguishes clearly between natural virtue (which includes things such as kindness to friends and care for children) and artificial virtue (which includes matters of justice, such as keeping contracts and paying taxes). And the connection between motivation and judgments of natural virtue, according to Hume's view, may differ from the connection between motivation and judgments of artificial virtue. See Hume, 2000, 307-322 and 367-378.

^{xv} Some have objected to this suggested internalist analysis of the case of the person who judges that he ought to visit his friend in the hospital, maintaining that while there is some essentially motivating feature in this case, this is due to the nature of friendship, and not to the nature of the moral judgment. The reason we would be puzzled if the

person said he didn't care to visit his friend in the hospital, according this objection, is because of how we conceive of friendship, not because of how we conceive of moral judgment. In response, I would first point out that there are other cases in which there moral judgments seems to be essentially motivating even though friendship is not involved. In the main text, for instance, I mention the judgment that I ought to spend more time with my children, or do more to combat racism. Secondly, I would maintain that the friend-in-the-hospital case could be recast so that it doesn't directly involve the concept of friendship. The relationship between the person in the hospital and the potential visitor could be described by ethically neutral accounts of the actions the two have performed in the past. But recasting their relationship in this way may not make it any less puzzling for the potential visitor to say, "I know I ought to visit him, but I just don't care to." Thirdly, I would point out that someone's being my friend is not always connected to my having a motive to act in a particular way. What makes us expect a motive is my judging that I ought to do something vis-à-vis my friend. For instance, my judging that my friend's mortgage payment is difficult for him to meet will not necessarily lead you to expect that I possess a motive to help him meet it. On the other hand, my judging that I ought to help him meet it will probably lead you to expect that I possess a motive to do so. And fourthly, I would suggest that this objection seems to assume that we can distinguish between different parts of a moral judgment (between, for instance, the 'ought' part and the 'friendship' part of the judgment that I ought to help my friend) that in fact may form a seamless whole in our minds. And the seamless whole that is the judgment that I ought to visit my friend may have an essentially motivating

character that the seamless whole that I ought to return ten dollars to Wal-Mart may not have.

^{xvi} See also Nichols 2004 and Stich and Weinberg, 2002, 640-642.

^{xviii} A more recent development relevant to the issues discussed here is exemplified by Scanlon 1998 and Gibbard 2003, who have advanced analyses of certain restricted aspects of our moral discourse. Scanlon, for instance, says that he is presenting an account not of morality as a whole “but rather of a narrower domain of morality having to do with our duties to other people” (Scanlon, 1998, 6). And Gibbard says that he is presenting an account only of the “concept of ought” (Gibbard, 2003, x). Scanlon acknowledges that the restricted nature of his most recent account contrasts with earlier work of his (Scanlon 1982) in which he did say he was attempting to explain morality as a whole. And I think Gibbard’s focus is also narrower than morality as a whole, which was the topic of Hare, Brink, and Smith. This recent work of Scanlon and Gibbard seems to me to be evidence of a moving away from the uniformist approach of 20th century meta-ethics toward a more variabilist approach in that Scanlon and Gibbard are contending that while a certain kind of internalist or externalist analysis captures some parts of our moral discourse it may not capture others. (That said, I think that even Scanlon and Gibbard’s recent works still are more uniformist than the phenomena dictate — that the analyses they present accurately capture even a more restricted domain that they contend. But of course backing that claim up would require carefully working through the details of their rich and sophisticated accounts.)