

Meta-ethical Variability, Incoherence, and Error

Michael B. Gill

Introduction

Moral cognitivists hold that in ordinary thought and language moral terms are used to make factual claims and express propositions. Moral non-cognitivists hold that in ordinary thought and language moral terms are not used to make factual claims or express propositions. What cognitivists and non-cognitivists seem to agree about, however, is that there is something in ordinary thought and language that can vindicate one side of their debate or the other.

Don Loeb raises the possibility — which I will call “the variability thesis” — that ordinary moral thought and language contains both cognitivist and non-cognitivist elements, and that there is no principled reason for thinking that either the cognitivist or non-cognitivist elements are conceptually more primary or aberrant than the other. According to the variability thesis, cognitivists accurately capture some aspects of what we think and say when we use moral terms and non-cognitivists capture other aspects, but neither side provides a correct analysis of ordinary moral thought and language as a whole.

Loeb also contends that if ordinary moral thought and language is variable in this way, we may be forced to conclude that our use of moral terms is fundamentally confused. The variability thesis might imply a new kind of irrationalism, which Loeb calls “moral incoherentism.” Moral incoherentism is new in that it holds not that we use moral terms non-cognitively (as Blackburn and Gibbard hold) nor that what we take to be moral properties are never instantiated in the world (as cognitivist error theorists such as Mackie hold) but rather that our moral thought and language are so meta-ethically self-contradictory that it is impossible to coherently apply moral terms at all.

I think the variability thesis is eminently plausible and of great significance for meta-ethical inquiry; I discuss this point in Part 1. I do not think, however, that the variability thesis gives as much support to moral incoherentism as Loeb thinks it does; I discuss this point in Part 2. In Part 3, I briefly sketch my own view of the relationship between moral variability and error.

Part 1

The variability thesis can explain well the state of debate between cognitivists and non-cognitivists. Very impressive, highly sophisticated positions have been developed on both the cognitivist and non-cognitivist

sides. It's possible, of course, that one side is completely wrong, the impressiveness and sophistication of its advocates' arguments notwithstanding. But it's also possible that both sides have gotten a lot right. The variability thesis fits with this latter possibility, as it implies that there really are cognitivist aspects to our moral discourse, which the cognitivists have accurately analyzed, and that there really are non-cognitivist aspects, which the non-cognitivists have accurately analyzed. The one big mistake both sides have made, the variability thesis implies, is to assume that the fundamentals of our moral discourse are entirely cognitivist or entirely non-cognitivist.¹

Loeb initially applies the variability thesis to the debate between cognitivists and non-cognitivists, but I take it that he thinks it can be applied to other contemporary meta-ethical debates as well. Is there a necessary connection between moral judgment and motivation or only a very strong contingent connection? Are moral reasons objective or relative? According to the variability thesis, such questions pose a false dilemma. For within our moral thought and language can be found both internalist and externalist, and objectivist and relativist, aspects — and none of those aspect can be shown to be conceptually more central or aberrant than the others. Both sides of each debate have correctly analyzed some aspects of our moral thought and language. But both sides have also made the mistake of assuming that our moral thought and

language must at some fundamental level admit of only an internalist or externalist, or only an objectivist or relativist, analysis (p. XXX, p. XXX).

Such is the possibility raised by the variability thesis. But what can we point to besides the persistence of meta-ethical disagreement to support it?

Loeb contends that in order to test the variability thesis — or, for that matter, any meta-ethical position — we must engage in serious empirical investigation. Meta-ethicists should not rely simply on their own armchair intuitions, nor on haphazardly collected impressions and anecdotes. Meta-ethical inquiry should be based, rather, on a responsible method of data-collection, one that attends to large, representative samples of the uses of moral terms and does its best to control for observer bias.² Loeb notes that there may be especially vexing difficulties in collecting data that will bear in direct and uncontroversial ways on certain meta-ethical debates (such as the debate between cognitivists and non-cognitivists) (see pp. XXX-XXX). But I think he is absolutely right in his contention that meta-ethicists ought to make a much more determined and thorough attempt than they have traditionally done to ground their positions in systematic observations of the phenomena they are trying to give an account of.

Part 2

Of course we cannot say in advance that empirical investigation will confirm or disconfirm the variability thesis, just as we cannot say in advance that it will confirm or disconfirm cognitivism or non-cognitivism, internalism or externalism, objectivism or relativism, or any other meta-ethical position. If we heed Loeb's call for a new, more serious empirical grounding for meta-ethics, all of these positions have to be taken, at this stage, to be merely hypotheses in need of testing. But let us consider nonetheless what the implications would be if the variability thesis were vindicated. If we were to find that ordinary moral thought and language contain both cognitivist and non-cognitivist aspects — and if there could be found no principled way of granting conceptual priority to, or of dismissing as conceptually aberrant, one of these aspects — what conclusions would we be compelled to draw about our concept of morality?

According to Loeb, the variability thesis implies or at least strongly suggests that our moral thought and language embody incompatible commitments, that participants in moral discourse are engaged in an activity that harbors internal contradictions, that incoherence infects the very semantics of moral terms. If the variability thesis is true, according to Loeb, it would not simply be the case that people who use moral terms are trying to refer to something that does not exist, which is the kind of

error made by 19th century adherents to the theory of phlogiston, an entity that scientific investigation revealed did not exist but at least could have existed. Participants in moral discourse would, rather, be like people discussing the characteristics of a round square, an activity we know to be mistaken without having to do any scientific investigation at all.

I do not think, however, that the variability thesis leads as directly to the conclusion that moral thought and language is as pervasively incoherent as Loeb suggests. That we can find in our uses of moral terms both cognitivist and non-cognitivist (as well as internalist and externalist, and objectivist and relativist) aspects might signal ineluctable incoherence, but then again it might not. It all depends on where those different aspects are located, on how they're distributed in our thought and language. If the different aspects are implicated by each and every use of moral terms, then Loeb's diagnosis of incoherence will be apt. But if one of the aspects is implicated within one pocket of moral thought and language, and the other aspect is implicated within a different pocket of moral thought and language, then Loeb's moral incoherentism might be a misdiagnosis. For it might be perfectly sensible to use a moral term in a way that involves one commitment, and also to use a moral term in a way that eschews that commitment, just so long as the first use occurs in a situation that is semantically³ insulated from the situation in which the second use occurs.

Two things that cannot both be coherently asserted at the same time might each be coherently asserted at different times.

To illustrate this possibility, we can point to our uses of “happy” and its cognates. Sometimes “happy” is used non-cognitively. If a person says, “I’m happy,” she may be expressing an attitude rather than trying to describe anything about herself. But sometimes “happy” is used cognitively. If a person says of someone else, “He’s happy,” she may be trying to describe something about the other person. And among cognitive uses of “happy,” there are also variations. Sometimes “happy” is used to refer exclusively to occurrent feelings. Sometimes it is used to refer to dispositions. Sometimes it is used to refer to a condition that is necessarily connected to objective (non-affective) features of a person’s situation. Now if we were to gather up all the commitments implicated by all our different uses of “happy” and take each of them to be implicated by every use of “happy,” then our thought and language about happiness would look to be inexorably confused, akin to discussion of a round square. There would be something ineluctably incoherent about a mode of discourse that commits us to holding that one and the same person, at a single time and in a uniform way, is both happy and not happy. But our happiness discourse is not that confused or incoherent. It is perfectly sensible, when describing a person’s life as a whole, to use “happy” in a way that involves a commitment that, were one seeking to express

gratitude to another for a recent act of kindness, one would eschew. What would be erroneous would be an analysis that fails to take into account the differences between the various contexts in which “happy” is used.

It seems to me that just as we can give a reconciling pluralist account of “happy,” so too might we be able to give a reconciling pluralist account of moral terms. And the possibility of such an account opens up the conceptual space to accept the variabilist thesis while rejecting Loeb’s moral incoherentism. For if such a pluralist account were at hand, we could hold that some pockets of ordinary moral discourse really are best analyzed as thoroughly cognitivist (or objectivist, or externalist), and some pockets really are best analyzed as thoroughly non-cognitivist (or relativist, or internalist). But the existence of such variability would not necessarily reveal that ordinary moral thought and language are incoherent or confused. It may, reveal, rather, that our moral terms are flexible enough to be put to numerous different kinds of uses. Moral terms, according to this possibility, can be used cognitively and non-cognitively, relativistically and objectively, externally and internally — even if they cannot coherently be used in all these different ways at the same time. So ordinary folk do not necessarily make any mistake when they put moral terms to these different uses, so long as they do not try to use them in too many different ways simultaneously. If there is a mistake that needs diagnosing here it is that of meta-ethicists who have assumed

that if a certain feature (such as a commitment to cognitivism, or to objectivism, or to a necessary connection between moral judgment and motivation) is implicated by the use of moral terms in one pocket of discourse it must also be implicated by the use of moral terms in every other pocket.

Let me fill in a bit this notion of pockets of moral discourse. What I have in mind are two types of cases. The first type of case is that in which some people use a moral term in one way while other people always use the moral term in a different way. The second type of case is that in which some people use a moral term one way in certain situations and the same people use the moral term in another way in other situations. There may, for instance, be some people — say, some college sophomores — who always use a moral term in a relativistic way, while they may be others — say, some priests or rabbis — who always use the moral term in an objectivist way. And there also may be some people who use moral terms relativistically in certain situations — say, when discussing the moral status of individuals in distant times or places, or when conversing with other people who themselves use moral terms in a predominantly relativistic way — and who use moral terms objectively in other situations — say, when assessing the laws, policies, or customs of their own country, or when conversing with other people who themselves use moral terms in a predominantly objectivist way. Meta-ethical objectivists typically

begin their accounts by focusing on the pockets of moral discourse that truly are objectivist. They then try to show that, despite the appearance of relativistic commitments in some other pockets, everyone's uses of moral terms are fundamentally objectivist — that the seemingly relativistic uses are parasitic on objectivist uses, or insincere, or implicitly placed in inverted commas, or otherwise conceptually aberrant. Meta-ethical relativists typically do the same sort of thing, although they of course begin with uses of moral terms in relativistic pockets and then try somehow to dismiss or accommodate the seemingly objectivist uses. Meta-ethical objectivists and meta-ethical relativists, in other words, both begin by pointing to uses of moral terms that really do fit very well with their own analyses, and then both have to fight like hell to explain away the uses that look not to fit.

Loeb rightly raises the possibility that when we pay sufficiently close attention to how moral terms are actually used in ordinary discourse, we will find that there truly are both objectivist and relativistic commitments — and that neither sort of commitment can be convincingly explained away. But it seems to me that Loeb does not give sufficient consideration to the possibility that the objectivist and relativistic uses may be semantically insulated from each other. He seems to share with meta-ethical objectivists and relativists what I call the “uniformity assumption” — namely, that if a meta-ethical commitment is implicated in

one pocket of ordinary moral discourse then it must also be implicated in every other pocket. This is the assumption I want to question. The possibility I want to raise is that some pockets of moral discourse are consistently objectivist, that other pockets are consistently relativistic, and that the uses of moral terms in each of these types of pockets is completely coherent. That we can find in ordinary moral discourse as a whole both objectivist and relativistic commitments does not imply that both of these different commitments are implicated in every pocket of ordinary moral discourse — just as that we can find in ordinary “happiness” discourse as a whole a good reason to say that one and the same person is happy (in one sense) and not happy (in another sense) does not imply that our thought and language about happiness suffers from any ineluctable incoherence.⁴

But let us now consider how this idea that ordinary discourse involves variable meta-ethical commitments that are semantically insulated from each other might apply to the debate between cognitivist and non-cognitivists. According to the variabilist-insulationist hypothesis, there may be some people who use moral terms in a way that is best analyzed as non-cognitivist and other people who use moral terms in a way that is best analyzed as cognitivist. Examples of the first sort might be Beavis and Butthead, who use value terms in ways that seem to be most accurately analyzed as non-factual, as expressing attitudes rather than

propositions. Examples of the second sort might be certain evangelicals, who claim that they use value terms to represent God's will and whose use of value terms turns out to track perfectly factual claims they hold about God's will; or hard-core utilitarians, who claim that they use value terms to represent certain facts about the production of happiness and whose use of moral terms turns out to track perfectly factual claims they hold about the production of happiness. Non-cognitivists have developed impressive, sophisticated ways of accounting for the seemingly cognitivist character of the uses to which certain evangelicals and hard-core utilitarians put moral terms. And cognitivists would surely have ways of accommodating or dismissing the thought and language of people like Beavis and Butthead. But why think that either of these uniformist, one-size-fits-all positions captures all the phenomena of the uses of moral terms better than a variabilist-insulationist account? Why think that the moral discourses of priests, rabbis, evangelicals, utilitarians, Beavis and Butthead all share the same meta-ethical commitments? The possibility I want to raise is that the phenomena can be better captured — accounted for in a way that is explanatorily more virtuous than either cognitivism or non-cognitivism — by the view that Beavis and Butthead really do use moral terms thoroughly non-cognitively and that certain evangelicals and utilitarians really do use moral terms thoroughly cognitively. Loeb criticizes cognitivists and non-cognitivists for failing to consider the

possibility that ordinary moral thought and language contains both cognitivist and non-cognitivist features. With this criticism I completely agree. It seems to me, however, that Loeb's moral incoherentism is based on the idea that cognitivist and non-cognitivist features are both implicated by all, or at least by most, ordinary uses of moral terms. But I think we should at least take to be a live option the possibility that these features are usually insulated from each other — that our moral discourse rarely implicates both kinds of features simultaneously.

Indeed, as I noted in my discussion of objectivism and relativism, it seems to me plausible that even the uses to which a single person puts moral terms might be variable and insulated in this way. Some of us might spend time in both the non-cognitivist camp (in which Beavis and Butthead can be found) and in the cognitivist camp (in which certain evangelicals and utilitarians can be found) but not reside permanently in either. Some person may, for instance, use moral terms in a way that is best analyzed cognitively when she is discussing policy-choices in a professional setting (let us say, when, in her capacity as a physician, she is serving on an ethics committee that is trying to decide whether to alter the hospital's DNR policies) and yet also use moral terms in a way that is best analyzed non-cognitively when she is discussing personal issues in a non-professional setting (let us say, when she is talking with a close friend about how a mutual acquaintance of theirs went about ending a romantic

relationship). But such a person would not necessarily be guilty of any incoherence or confusion. She could very well be applying moral terms effectively and appropriately in both the professional and personal situations. The meta-ethical commitments it would be fair to attribute to her in one situation would be different from the meta-ethical commitments it would be fair to attribute to her in the other situation, but because the two situations are insulated from each other she may have no problem using moral terms in completely sensible ways on both occasions. Moreover, there may be no principled reason to hold that her use of moral terms in one of the situations is conceptually more central or aberrant than her use of moral terms in the other situation. Even though there are meta-ethical differences between them, both kinds of uses may be orderly, sincere, clear — and completely ordinary — cases of moral thought and language.

Some might continue to insist, however, that certain uses of moral terms just must be conceptually superior to other kinds. All the different uses, some might hold, cannot be equally copasetic. In responding to this way of thinking, variabilist-insulationists should start by drawing a distinction between the attempt to analyze ordinary uses of moral terms — which we can call descriptivist meta-ethics — and the attempt to articulate the way of using moral terms that is normatively the best — which we can call prescriptivist meta-ethics. The variability thesis, as Loeb

points out very clearly, is entirely on the descriptivist side (p. XXX, p. XXX), and I intend the variabilist-insulationist hypothesis to be entirely on that side as well. But even if this hypothesis turns out to be an accurate account of how people actually use moral terms, there may still be excellent normative reasons to prefer one, uniform way of using moral terms to any other. One can be a variabilist in descriptive meta-ethics (holding that cognitivist and non-cognitivist, and relativist and objectivist, and internalist and externalist, commitments are all present in ordinary moral thought and language) and still contend that one set of meta-ethical commitments is normatively superior to any other set (holding that everyone should always use moral terms only in ways that involve, say, cognitivist, objectivist, and externalist commitments). Indeed, it seems to me rather unlikely that the particular set of meta-ethical commitments that any of us comes to think are normatively the best will also turn out to be exactly the same ones that characterize all ordinary uses of moral terms.

Some might hold, however, that even within descriptive meta-ethics there are principled reasons to reject a variabilist account even before delving into the details of lots of particular examples of ordinary usage. One of the reasons that might be offered for this position is that there must be some uniform commitments that characterize all ordinary uses of moral terms because we do in fact classify all these uses as *moral*.

When we classify some bit of thought or language as moral, we do not mean to say merely that words like “moral,” “right,” or “virtuous” have been used. We mean to say, rather, that a certain concept is in play, and that concept must have some more or less determinate shape in order to serve the function it does in fact serve. But morality could not serve this function if it were taken to cover bits of thought and language with as widely disparate features as a variabilist account implies. In response, variabilists should, I think, accept that their account implies that our concept of morality is something of a hodgepodge — more of a miscellany than a system. But the hodgepodge-miscellany picture may very well fit the phenomena better than the uniformist accounts that meta-ethicists have traditionally offered. The many bits of thought and language we classify as moral may bear only a family resemblance to each other. But that we can give nothing more precise than a family resemblance account of morality does not imply that the concept is unusable. Our uses of “game” bear only a family resemblance to each other. But the concept of a game is perfectly usable nonetheless. Perhaps it’s not the most orderly concept we possess. Some of our “game” talk may be imprecise, and that imprecision may at times produce confusion and miscommunication. But most of the time our “game” talk works quite well. Most of the time, when people use “game,” all the participants in the conversation know what is being talked about to an extent clearly

sufficient for a coherent, sensible, useful application of the term. Similarly, morality may not be our most orderly concept. Some uses of moral terms may produce confusion and miscommunication. Consider the potential pitfalls of a conversation between an evangelical and an undergraduate cultural relativist, or between Beavis and a rabbi. But such confusion and miscommunication does not occur in all our moral conversations. Often enough, we manage to use moral terms coherently and effectively. And we manage this even though the way we use moral terms in one situation may differ from the way we use them in another situation. Ordinary users of moral terms usually find it relatively easy to pick up on the conversational clues that signal the meta-ethical parameters of any particular moral discussion. It's only traditional meta-ethics that, in this regard, may have been a bit clueless.

Part 3

Let me now briefly discuss the relationship between the variabilist-insulationist view I've sketched and several kinds of meta-ethical error.

The Uniformist Error

The possibility of a variabilist-insulationist account implies that it is a mistake to assume that just because a meta-ethical commitment is

implicated by one bit of moral thought and language it is also implicated by every bit of moral thought and language. I attribute this uniformist error to many of the twentieth century descriptive meta-ethicists who argued that our moral discourse as a whole is thoroughly absolutist or relativist, externalist or internalist, cognitivist or non-cognitivist.

I do not think, however, that the uniformist error infects most of ordinary moral thought and language. When ordinary people think or talk about specific moral matters, they are usually engaged in an activity that does not commit them to views on whether moral discourse as a whole is uniformly absolutist, externalist, cognitivist, or whatever. When a physician in a hospital ethics committee meeting is considering changes to DNR policy, for instance, it may be fair to attribute to her meta-ethical commitments about how to think and talk about matters of professional ethics in official settings. But that does not mean it is fair to attribute those same meta-ethical commitments to her when she is talking with a close friend about how a fraught personal issue ought to be dealt with. The meta-ethical commitments this second, personal conversation involves her in may be specific to matters and settings quite different from the matters and settings of conversations about professional ethics in official committees. Thought and language directed at specific, local moral matters need not — and typically do not — implicate global positions on moral discourse as a whole.

That is not to say that all ordinary users of moral terms consciously reject the uniformity assumption and embrace the variabilist-insulationist hypothesis. I suspect that if asked some people would say that all moral thought and language does involve certain uniform meta-ethical commitments, and others would give a more variabilist answer. But however that may be, I don't think there is good reason to attribute to most instances of ordinary moral thought and language a commitment to the uniformist assumption. The analogy to our concept of a game is once again instructive. Most ordinary people have probably given little or no thought to the question of whether all the things we call a "game" share some robust conceptual characteristics or bear only a family resemblance to each other. Probably, if asked, some people would give one answer and other people would give the other. But now consider those people who say that all the things we call a "game" do share some robust conceptual characteristics (meta-game uniformists). Do we have any reason to think that a commitment to the conceptual uniformity of game-discourse as a whole has infected all of their particular first-order discussions of football, poker, duck-duck-goose, and so on? I think it's clear we do not. That a person gives a uniformist answer to the meta-level question about the concept of a game does not give us reason to think that her ordinary thought and language about particular games has involved any uniformist commitments. Similarly, that a person gives a

uniformist answer to a meta-level question about the concept of morality does not constitute strong evidence that all of her ordinary moral thought and language has involved a commitment to robust meta-ethical uniformity.

Loeb's moral incoherentism

According to Loeb's moral incoherentism, ordinary people "use the moral words *both* to make factual assertions *and* to do something incompatible with the making of such assertions" (p. XXX). This is Loeb's "attenuated" error theory, according to which internal inconsistency ineluctably afflicts moral discourse (p. XXX).

While I agree with Loeb that we can find in ordinary discourse cases in which moral words are used to make factual assertions and cases in which moral words are used to do something incompatible with the making of such assertions, I do not think that there are good grounds for holding that the error of internal inconsistency is as pervasive as his moral incoherentism implies. This is because I think it likely that the cases in which moral terms are used one way are usually semantically insulated from the cases in which moral terms are used the other way. My hypothesis is that within some particular conversations, moral terms are used in a coherently cognitivist (or objectivist or externalist) way; within other conversations, moral terms are used in a coherently non-cognitivist

(or relativist or internalist) way; and most ordinary people are fairly adept at semantically navigating between and within the two different types of conversations.

At the same time, I also think some people fail at this navigational task, and perhaps most of us fail at least some of the time. At least some of the time, the semantic variability of our moral terms can cause our moral discourse to flounder or founder. I suspect these kinds of problems occur at the interpersonal level more often than at the intrapersonal level. This kind of interpersonal problem arises when one person in a conversation uses moral terms in a way that is semantically different from the way the other person uses them. It might be the case that the first person always uses moral terms in a certain way and the second person always uses them in another, incompatible way, as, for instance, in conversations between die-hard relativists and evangelicals, or between Beavis and a rabbi. Or it might be the case that for various reasons the first person, when discussing a particular moral issue, happens to use moral terms in a way that is incompatible with how the second person happens to use moral terms when discussing that issue, even if the two of them may use moral terms consonantly when discussing many other issues. Two people may be able to discuss the ethics of famine relief in a perfectly coherent manner but then encounter formidable semantic obstacles when the conversation turns to gay marriage. In such cases,

moral communication can become very difficult, like trying to hold a serious conversation over a bad cell phone connection, or it may break down altogether, reducing the parties to hurling verbal clods at each other. I imagine we can all recall some situations, personal and societal, in which moral conversation seems to have degenerated in this way.

But even if these interpersonal cases are more common, I think Loeb is right that profound semantic confusion can also occur intrapersonally. A single person's thought and language might involve meta-ethical commitments that really do conflict with each other and this might make it impossible for her to use moral terms coherently. And perhaps most of us have been afflicted by this kind of confusion at one time or another. Perhaps most of us have at one time or another been in a state not merely of moral perplexity (unclear about the right thing to do in a particular situation) but of meta-moral perplexity (unclear about what would even constitute being right in a particular situation). I don't think ordinary people end up in this state every time they use moral terms, but it seems reasonable to me to hold that some of us are in this state at least some of the time. In certain situations or in the face of certain issues, the semantic insulation that separates one set of our meta-ethical commitments from another set may break down.

Now what I've said in the preceding two paragraphs may make my disagreement with Loeb seem to be slighter than it might have first

appeared. I think Loeb is right that some uses of moral terms suffer from incoherence. And perhaps there is no obstacle to his holding, as I do, that some uses of moral terms are coherent (p. XXX). Our disagreement may thus seem to be only about the frequency of these types of cases: I think coherence is more common than he does. And my challenge to Loeb may, then, simply be to give us reasons to think that breakdowns in our uses of moral terms are typical rather than atypical.

It seems to me, however, that a number of Loeb's claims suggest not merely that *some* uses of moral terms are incoherent but that incoherence clings to *every* use of moral terms — that the error of holding incompatible meta-ethical commitments infects the semantics of moral thought and language as whole (p. XXX-XXX, XXX, XXX-XXX, and XXX.) And if Loeb holds that incoherence pervades all our moral thought and language, then the difference between us is not simply about the frequency of breakdowns. For I think that some uses of moral terms are free of incoherence altogether. I think that, even though serious meta-ethical mistakes of the kind Loeb describes do sometimes occur, there is no insurmountable semantic barrier to our using our moral terms in a completely coherent way.

Mackie's cognitivist error theory

In addition to thinking that the kind of incoherence on which Loeb focuses afflicts some (but not all) of our moral thought and language, I also believe that some instances of ordinary moral thought and language commit the kind of error Mackie focused on. Mackie claimed that our moral discourse, while not internally logically inconsistent, commits us to the existence of entities that in fact do not exist — that our moral thought and language presuppose false views of what the world is like. And it seems to me likely that at least some of the time some people do think and talk in ways that make it fair to attribute to them this kind of error. There may, for instance, be people who hold certain theological views about the origin of morality, and these views may permeate their moral thought and language to such an extent that an accurate analysis of their moral thought and language would have to include them. Their theological views may be so entrenched in their moral thought and language that the latter could not be maintained without the former; if they were ever to abandon their theological views, their moral discourse would undergo a drastic change or they would abandon thinking and talking in moral terms altogether (p. XXX). But these persons' theological views, which an accurate analysis will have to include as integral to their moral discourse, may be false. And if views inextricably built into persons' moral thought and language are false, then it seems to me just to convict their moral discourse as a whole of the kind of error on which Mackie

focused. Or, to take another example, consider self-professed relativists who hold certain views about the causal role culture plays in the development of morals, and whose moral thought and language are so embedded in those views that an accurate analysis of their moral discourse will have to include them. Such people, we can imagine, could not give up their views about the cultural origins of morality and still engage in the kind of moral thinking and talking that they currently do engage in. But these people's views about the causal role culture plays in the development of morals may be false. And in such a case, once again, it seems to me just to hold that their moral discourse as a whole is erroneous, and erroneous in roughly the way Mackie describes. In some cases, a person's meta-theory may so saturate her practice that we will be able to discern simply by observing her practice what meta-theory she holds; in some cases, persons' meta-theories about morality may infect their first-order moral practices. And if in such cases the meta-theory is erroneous, then the first-order practice will be erroneous as well.

But I do not think there are good grounds to hold, as Mackie did, that everyone's first-order moral practice is inextricably linked to an erroneous meta-theory. For one thing, some people may hold meta-theories that are not erroneous, and their first-order moral practice may reflect their non-erroneous understanding. And for another thing, many people much of the time may use moral terms in ways that do not

embody any very specific meta-theory at all, erroneous or otherwise.⁵ Many people much of the time may use moral terms in ways that do not contain observable evidence of the robust kinds of meta-ethical commitment that characterize the uses of moral terms of the theologians and relativists I described in the previous paragraph. Indeed, even some people who say they hold to certain meta-ethical claims (say, that morality is determined by God's will, or that morality is nothing but cultural norms) may use moral terms in ways that indicate that they would continue to think and talk about particular moral matters in pretty much the same way even if they were to abandon those meta-ethical claims, and this gives us good reason to think that their first-order moral practices are not infected by their meta-ethical views. As Blackburn has put it, their uses of moral terms may be capable of being "clipped on to" a non-erroneous metaphysic as easily as they can be clipped on to an erroneous one (1993, p. 151). And so it's just to hold, in such cases, that even if the meta-ethical claims they say they hold are false, their ethical thought and language does not embody error.

Parts of ordinary moral thought and language embody error. But the kinds of error some parts embody differ from the kinds of error other parts embody. And still other parts of ordinary moral thought and language do not embody error at all. A blanket error theory will not capture all the phenomena any more than a blanket non-error theory will.

A successful theory will have to distinguish what is erroneous in our moral discourse from what is not, and between different kinds of error. Such a theory will be richly multifarious, complicated, messy. But that is just what we should expect of any account that is true to what human beings actually think, say, and do.

Notes

1. See Loeb, p. XXX. I discuss the variability thesis in more detail in my “Meta-ethical Variability and our Miscellaneous Moral Discourse” (forthcoming).
2. For a full treatment of this empirical approach to meta-ethics, see Nichols (2004).
3. Throughout my discussion, I follow Loeb in taking the meta-ethical issues to involve the semantics of moral terms. It’s not clear to me, however, that these issues really are about semantics rather than pragmatics. I am not sure, that is, whether all the meta-ethical commitments Loeb and I discuss are part of the semantics of moral terms or some of the commitments are merely pragmatically implicated by moral terms. So while I say throughout my discussion that the meta-ethical commitments in one situation can be *semantically* insulated from the meta-ethical commitments in another situation, it could be the case that what I should be saying is that the meta-ethical commitments in one

situation can be *pragmatically* insulated from the meta-ethical commitments in another situation.

4. But the following objection might be raised to my attempt to analogize our moral terms to “happiness.” If I say that a person is happy, and mean by it that she is experiencing a pleasurable occurrent emotion, and you say that the same person is not happy, and mean by it that her life as a whole is not going well, we will probably be able to quickly disambiguate our uses of the term “happy” and thus dissolve any disagreement there might have seemed to be between us. But if I say that some action is wrong and you say that the same action is not wrong, then even if we are using the term “wrong” in different ways (say, I am using it objectively and you are using it relativistically), it seems likely that there will still be some kind of conflict between us that will not be as easily dissolvable as was the apparent disagreement between us about whether a person is happy. I cannot address this objection fully here, but I would respond to it, first of all, by claiming that a lot of apparent moral disagreements *can* be dissolved by disambiguating between the different senses in which moral terms are being used, and, secondly, by claiming that people using moral terms in meta-ethically different ways can nonetheless disagree (or at least come into conflict) with each other about some first-order moral issues (see Loeb, pp. XXX-XXX).

5. The possibility raised but not developed here is that moral terms are often used in ways that are indeterminate with regard to the questions separating cognitivists and non-cognitivists, internalists and externalists, objectivists and relativists. There may be nothing in some uses of moral terms that gives us any principled reason for taking them to implicate one kind of meta-ethical commitment rather than the other. I discuss this possibility — which I call the Indeterminacy Thesis — in some detail in “Meta-ethical Variability and our Miscellaneous Moral Discourse.”