

*Midwest Studies in Philosophy, XXXI (2007)*

1 Can Moral Obligations Be Empirically  
2  
3 Discovered?  
4

5 JESSE J. PRINZ

6  
  
7 **T**hese are good days for empirically minded philosophy. It is now common to  
8 see philosophers cite laboratory studies in support of their theories, and a  
9 growing number of philosophers are conducting their own psychological experi-  
10 ments. From a worm's eye perspective, it looks a little bit like there is a method-  
11 ological revolution taking place. Gone are the days when philosophers had to read  
12 psychology journals covered in discrete bindings so as to hide a perverse interest in  
13 the latest sexy results. At least that is true in most areas of philosophy. Lately, the  
14 methodological revolution has been gaining ground in ethics. Philosophers now  
15 routinely cite the burgeoning empirical literature on moral judgment. The majority  
16 of this work has been descriptive in nature. Psychology and neuroscience are being  
17 used to determine what goes on in people's heads when they decide that something  
18 is morally right or morally wrong. Traditionally minded ethicists continue to feel  
19 insulated from the onslaught of empirical moral psychology. Ethics has long been  
20 in the business of telling us how we ought to behave, rather than merely describing  
21 how we do behave. Such normative claims appear to have a kind of immunity to  
22 empirical investigation. Psychology can reveal what we morally detest, but it  
23 cannot reveal what we should detest. Or so it would seem. In this chapter, I want  
24 to explore this supposition and argue that empirical methods can actually shed  
25 light on what we ought to do. Along the way, I will also make some other contro-  
26 versial claims. I will say that morality has an emotional basis, and I will endorse a  
27 form of relativism. Of course, such conjectures require much more discussion than  
28 I can offer here. My arguments will be preliminary, and I will content myself with  
29 having shown how, in principle, one might go about trying to discover obligations

© 2007 Copyright The Authors  
Journal compilation © 2007 Blackwell Publishing, Inc.

1 empirically. Along the way, I will indulge in periodic reflections on methodology,  
2 which can be accepted quite independently of the contentious philosophical  
3 conclusion I draw. If my philosophical conclusions are rejected, I hope they are  
4 rejected on empirical grounds.

### 5 6 **STEP 1: THE EMOTIONAL BASIS OF MORAL CONCEPTS**

7  
8 Moral judgments, such as the judgment that it is wrong to sleep with your best  
9 friend's lover, contain moral concepts, such as the concept *WRONG*. There is no  
10 shortage on controversy about the nature of concepts, but according to the major-  
11 ity view in cognitive science (and increasingly in philosophy) concepts are mental  
12 representations: concrete particulars in the head that aim to represent things. The  
13 concept *WRONG*, on this approach, is a mental representation that aims to represent  
14 the property of wrongness, whatever that turns out to be. Concepts are also pre-  
15 sumed to be the tools by which we grasp words and the mechanisms by which  
16 words come to refer. Thus, the concept *WRONG* is something in my head that allows  
17 me to understand the word "wrong" and it is that in virtue of which the word refers,  
18 in fact it is not an empty expression. I will not defend this basic approach to what  
19 concepts are here (see Prinz 2002), but I want to draw out an obvious implication.  
20 If concepts are in the head, then we can study moral concepts using whatever  
21 means are available for determining what goes on psychologically when people use  
22 moral vocabulary.

23 This is where neuroscience and psychology come in. There are now numer-  
24 ous studies that are designed to discover what goes on in the head during moral  
25 reflection. Consider, for example, the barrage of recent neuroimaging studies in  
26 which researchers measure brain activity during moral reflection and compare it  
27 with brain activity during nonmoral tasks. Some studies ask subjects to deliberate  
28 on moral dilemmas, some ask them to decide whether something is morally right or  
29 wrong, some show subjects morally significant images, some have them perform  
30 morally significant actions in virtual worlds, and so on (Greene et al. 2001; Heek-  
31 eren et al. 2003; King et al. 2006; Moll et al. 2002). These studies tend to show a  
32 similar pattern. In addition to brain areas associated with attribution of mental  
33 states (did the agent act intentionally?), the major players in moral cognition are  
34 brain areas associated with emotion. Common hotspots include areas such as the  
35 orbital frontal cortex and the temporal pole, which are involved in assigned emo-  
36 tional significance to events, and areas such as cingulate cortex, which are associ-  
37 ated with emotional experience. This has led researchers to conclude that emotions  
38 are centrally involved in moral judgment.

39 Neuroscience cannot reveal which emotions are active when people make  
40 moral judgments, because we cannot yet distinguish the neural circuits underlying  
41 distinct discrete emotions. But behavioral studies can shed some light here. Experi-  
42 ments have been designed to determine what emotions arise in reflecting on  
43 morally significant events. The most direct way to measure this is present partici-  
44 pants with an event and ask them what they would feel. For example, Rozin et al.  
45 (1999) gave subjects a list of moral transgressions that they had divided into three  
46 categories: crimes against persons (such as stealing or harming), crimes against

1 community (such as violations of status hierarchies), and crimes against nature  
2 (such as violations of sexual mores). In a forced-choice design, they found that such  
3 events elicit anger, contempt, and disgust respectively. In an unpublished follow-up  
4 study, I present subjects with a similar list of transgressions, but rather than asking  
5 how they would feel if someone else performed the immoral act, I asked how they  
6 would feel if they themselves behaved immorally. I found that crimes against  
7 persons elicit guilt, crimes against nature elicit shame, and crimes against com-  
8 munity can elicit either. The main lesson, for present purposes, is that there are a  
9 number of different emotions that arise in the context of moral judgment. What  
10 they share in common is that they are all emotions of blame. Each feeling can be  
11 described as negative appraisal of the person of action being judged.

12 The studies that I have been describing so far establish that emotions arise  
13 when we make moral judgments, but they do not establish how those emotions  
14 relate to judgments. Are they mere sequelae of moral reflection, or do they play a  
15 more integral role? To address this question, experiments have been designed to  
16 show that emotions can influence moral judgments. For example, Wheatley and  
17 Haidt (2005) showed that hypnotically induced disgust amplifies judgments of  
18 moral wrongness even in response to scenarios that subjects in a control group  
19 would regard as morally benign. Similarly, Schnall et al. (forthcoming) showed that 3  
20 wrongness judgments sometimes increase when participants assess scenarios at a  
21 filthy desk, and compared with a clean desk.

22 It has also been shown that moral judgments diminish when negative emo-  
23 tions are reduced. Schnall et al. (submitted) showed that subjects are offer lower  
24 wrongness judgments if they are allowed to wash their hands, and Valdesolo and  
25 DeSteno (2006) have shown that watching a comedy routine dramatically reduces  
26 the judgment that it is wrong to kill one person in order to save five. Similar  
27 conclusions can be drawn from studies of psychopaths. Psychopaths are known to  
28 have a profound deficit in negative emotions such as fear and sadness, and they also  
29 have a profound deficit to understand moral judgments. Blair (1995) has evidence  
30 that psychopaths treat moral rules as if they are mere social conventions.

31 Putting these pieces together, we might venture that moral judgments  
32 contain an emotional component. The neuroimaging studies suggest that emotions  
33 arise when we make moral judgments. The hypnosis and dirty desk studies suggest  
34 that those emotions are not mere epiphenomena but are related to what people  
35 report. The hand washing, comedy, and psychopathy studies suggest that these  
36 emotions do not merely influence moral attitudes but may be necessary for such  
37 attitudes. Remove emotions, and people no longer judge that things are morally  
38 wrong. The evidence does not deductively entail that moral judgments contain  
39 emotions, but that conclusion explains and predicts the empirical evidence.

40 The studies mentioned so far provide *a posteriori* support for the conclusion  
41 that moral judgments contain emotions. By that I mean it is the kind of support that  
42 we might not discover by merely reflecting on what we mean by moral terms.  
43 Philosophers have long assumed that the structure of concepts is available to  
44 introspection. This need not be the case. For example, psychologists have used  
45 empirical evidence to support the claim that most concepts are associated with  
46 prototypes: representations of typically category instances. Evidence for this claim

1 includes the fact that people are faster at classifying prototypical instances and  
2 people learn to classify prototypical instances earlier in development. We might not  
3 have discovered the existence and prototypes by introspection, but that does not  
4 mean that we should deny that concepts are prototypically structured. The ubiquity  
5 of prototype effects has been used to support the hypothesis that concepts just are  
6 prototypes, or at least contain prototypes as proper parts. The argument for such a  
7 conclusion is simple: concepts are what we use to grasp words; empirical evidence  
8 suggests that we use prototypes to grasp words; therefore, concepts are prototypes.  
9 Forget what you think about the second premise and, hence, the conclusion. My  
10 point here is that empirical methods can be used to discover the structure of  
11 concepts. Thus, we can empirically discover that moral concepts have an emotional  
12 basis.

13 That said, the emotional basis of moral concepts might be knowable *a priori*  
14 too. We may have conscious access to this fact. To me it seems perfectly obvious  
15 that moral judgments contain emotions. The problem with this bald assertion is that  
16 philosophers tend to have different intuitions. What seems obvious to me may  
17 seem obviously false to you. This is another reason why empirical methods can be  
18 useful. When philosophers disagree about the intuitions that direct their theories,  
19 they may be victims of theoretical bias. Theories can influence our intuitions. To  
20 escape such bias, we can try to consult the intuitions of people who are not  
21 professional philosophers, such as college undergraduates.

22 Suppose we want to know whether people overtly recognize that emotions  
23 have an intimate link to morality. We can design a study to see whether people  
24 regard emotions as epiphenomena that are merely contingently associated with  
25 moral views, or whether instead they recognize what the empirical findings tend  
26 to support: that morals have an emotional basis. Toward this end, I asked a group  
27 of North Carolina college students to consider two scenarios (Prinz, unpublished  
28 data). One describes a person who verbally condemns marijuana smoking but has  
29 no negative emotional response to it. The other describes a person who verbally  
30 insists that marijuana smoking is morally acceptable, but feels disgusted at those  
31 who smoke and ashamed when he himself smokes. For each scenario, subjects were  
32 asked whether the person's moral values were reflected in what they said or in how  
33 they felt. In both cases, a significant and sizable majority said that the moral value  
34 is reflected in the emotion (ninety percent in the first case and seventy percent in  
35 the second case).

36 Recall that the studies mentioned earlier already provide empirical support  
37 for the claim that moral judgments have an emotional component. The marijuana  
38 study shows that this empirical hypothesis may even readily accessible as a con-  
39 ceptual truth. The marijuana study suggests that people tend to attribute moral  
40 judgments when and only when emotions are in place, and the other studies suggest  
41 that people engage in what seem to be moral judgments when and only when  
42 emotions are in place.

43 The marijuana study helps to adjudicate between some competing interpre-  
44 tations of the other results. On one interpretation, emotions exhaust or our moral  
45 concepts. The concept *WRONG*, on this view, is just a negative feeling (or perhaps 4  
46 a disposition to have a negative feeling). But some philosophers might resist a

1 reduction of WRONG to a mere feeling; they might suppose that moral concepts have  
2 two parts: an emotional part and some kind of descriptive part. If that were the  
3 case, it should be possible to have these two parts independently. But then, in the  
4 scenario where someone feels disgusted about marijuana smoking but insists that  
5 it is okay, subjects should not be so quick to judge that the agent is opposed to  
6 marijuana. They should take his verbal behavior as evidence for the assumption  
7 that he lacks the nonemotional component of moral judgments. Thus, subjects  
8 seem to think that emotions are sufficient for moral judgments. We can also learn  
9 something from the other scenario, in which the agent feels no negative emotions  
10 but says that marijuana smoking is wrong. When presented with the handwashing,  
11 comedy, or psychopathy studies, one might be tempted to draw a modest conclu-  
12 sion that emotions are not really necessary for moral judgments, but merely facili-  
13 tate moral judgments. The fact that most subjects think the dispassionate marijuana  
14 critic is not really morally opposed to the drug suggests that they regard emotions  
15 as necessary to moral judgment. These observations suggest that emotions are not  
16 merely components of the concept WRONG, much less contingent components, but  
17 may actually constitute that concept in some sense. If one feels disgust and other  
18 emotions of blame toward an action, one thereby judges it to be wrong, and if one  
19 fails to have these emotions, there is nothing left to call the moral judgment.  
20 Further studies can help confirm that the constitution thesis is part of our folk  
21 theory of moral concepts.

22 One interesting feature of the marijuana study is that it complicates the  
23 standard philosophical distinction between *a priori* and *a posteriori* methods. When  
24 subjects fill out their responses on a questionnaire, they are using their *a priori*  
25 intuitions. But, no single individual's response provides adequate evidence for the  
26 claim that there is a conceptual link between morality and emotion. At best, we  
27 could say that the individual draws such a link. It is only when participants are  
28 averaged together that evidence for a general (that is to say, widespread) link  
29 emerges. So the link between moral judgments and emotions, though available to  
30 many through armchair reflection, is in this sense *a posteriori*. It would be helpful to  
31 introduce a third term for empirically oriented philosophy. We can say that concep-  
32 tual claims made on the basis of analyzing multiple *a priori* intuitions are *a*  
33 *posterior priori*—made after the prior. Many experimental philosophers have con-  
34 ducted *a posteriori priori* studies, and in this respect their methods continue to use  
35 armchair philosophical reflection as a window into the structure of concepts. I think  
36 these methods are strongest when coupled with less direct methods for observing  
37 conceptual structure like those described above. Without studies of how we actu-  
38 ally moralize, studies of folk intuitions about moralizing could be accused of  
39 getting morality badly wrong. Without studies of folk intuitions, studies of what we  
40 do when asked to moralized could be dismissed as showing that people are not  
41 really moralizing in those cases, but are instead doing something else (e.g., Sinnott-  
42 Armstrong [2006] argues that emotions are merely used as heuristics that stand in  
43 place of real moral deliberation).

44 Collectively, the studies that I have been discussing support the conclusion  
45 that moral judgments have an emotional basis, and, more specifically, that the  
46 concept WRONG is constituted by negative emotions. Elsewhere, I have elaborated

1 on this hypothesis considerably (Prinz 2007). Here I want to focus on a possible  
2 implication.

### 3 4 **STEP 2: FROM CONCEPTS TO REFERENCE**

5  
6 It is widely assumed in philosophy that if moral concepts are constituted by  
7 emotions, then they do not refer to anything. This is the kind of view that was  
8 defended by emotivists. They claimed that sentences expressing moral judgments,  
9 such as “stealing is wrong,” do not aim to assert anything that can be evaluated as  
10 true or false. “Wrong” expresses an emotion, and thus “Stealing is wrong” is just an  
11 expression of a negative feeling toward stealing. Some emotivists suggest that when  
12 we say stealing is wrong we are also enjoining listeners to feel as we do about  
13 stealing, but we are not attributing any property to stealing. Emotivists seem to  
14 assume that if a word expresses an emotion it cannot also refer to a property. That  
15 assumption is confused. Emotions can refer to properties. After all, we use emo-  
16 tions to classify things. We can divide the world into disgusting things and nondis-  
17 gusting things, scary things and nonscary things, irritating things and nonirritating  
18 things, and so on. The very fact that we use emotions to sort things suggests that  
19 emotions function as representations. Indeed, sorting can be characterized as a  
20 primitive form of reference. All things being equal, if a mental state is used to sort  
21 things into a category that can be distinguished from other categories, then that  
22 mental state refers to that category.

23 This is the starting place for a naturalistic theory of reference. If we think of  
24 referring a kind of activity, we can easily identify behaviors that indicate that  
25 reference is going on. If I put some items in bin A and others in bin B in a  
26 predictable way, then I am clearly distinguishing these items, and this can be  
27 characterized in semantic terms (I am representing A items and B items differ-  
28 ently). Emotions are good classifiers. I could fill a bin with disgusting things and  
29 another bin with nondisgusting things. That suggests that disgust is in the business  
30 of representing. If I say that something is disgusting, I imply that it has the property  
31 that disgust represents.

32 To figure out what a mental representation represents, we need to apply our  
33 best theory of reference. Philosophers have typically developed theories of refer-  
34 ence from the armchair. Even card-carrying naturalists about reference (those who  
35 want to characterize reference in nonsemantic terms) tend to carry on as if their  
36 project is immune to empirical evidence. The standard strategy is to assume that  
37 we have internal representations that refer to certain categories (such as natural  
38 kinds) and then see whether there could exist a relation between mind and world  
39 that uniquely maps internal states onto those categories. They try to show that  
40 certain natural relations could in principle get the semantics of our categories right.  
41 This project has its place, but I think it suffers from two limitations. First, practi-  
42 tioners never take the obvious next step of seeing whether the concepts we use  
43 actually do bear the specified relation to any category in the world. Suppose you  
44 think that concepts refer to what they are nomically related to (cf. Fodor 1990). It  
45 would be worthwhile to test this theory by seeing whether the concept that lights  
46 up when we use the word “dog” also bears a nomic relationship to the property of



1 doghood. If it does not, that would be an embarrassment to the theory. But as soon  
2 as we pose this test, we immediately get into trouble. First, it is far from clear how  
3 one can test for a law of this kind. The mental state type that is active when I use  
4 the word “dog” will probably be activated when I see some dogs and not others.  
5 Does that show that it refers to just some dogs? One should say it refers to all dogs,  
6 because the dogs that do not cause the mental state to activate belong to the same  
7 category as the dogs that do cause it to activate: they are all dogs. The semantic  
8 theorist cannot simply stipulate that the concept is nomically related to this larger  
9 category. Rather, the decision to treat the mental state as a representation of dogs  
10 seems to be a kind of decision that we make, not just as semantic theorists, but also  
11 as users of the concept. Of course, we do not have the mental resources to identify  
12 all dogs (we do not know the dog essence). But we can, as a matter of policy, treat  
13 the mental state that is activated by paradigm instances of dogs as a natural kind  
14 concept. When we do that, we are, in effect, adding another component to our  
15 theory of reference. We are saying that reference requires both nomic relations and  
16 semantic policies about how to select from the various possible categories that  
17 might be nomically related. I do not want to suggest that these policies cannot be  
18 naturalized. They probably can. I want only to say that the failure to think seriously  
19 about how to test semantic theories has led some theories to ignore the possibility  
20 that reference involves this further component.

21 This brings me to a second limitation of semantic theories as they are devel-  
22 oped by naturalistic philosophers of mind. Very often, reference is treated as a  
23 purely relational property that has no bearing on the behavior or psychology. It is  
24 a fact about me, as I am situated in the world, that my concepts refer to certain  
25 things and not others, but this fact is behaviorally inert. If the concept I express by  
26 “water” referred to XYZ rather than H<sub>2</sub>O, it would make no difference to me. By  
27 focusing on such examples, philosophers have implied that the theory of reference  
28 will have little to do with the theory of psychology (construed narrowly as a the  
29 science of behavior and processes that take place in our heads). But this is a  
30 mistake. If reference depends on semantic policies, then it depends on psychologi-  
31 cal states. Indeed, the reason why we say that my water concept refers to the local  
32 watery stuff is that I have a semantic policy to this effect. If I were totally neutral  
33 about the question of whether my water concept refers to a natural kind or to any  
34 clear liquid, then it would not refer to a natural kind. The study of reference should  
35 centrally involve an investigation of what we take our concept to refer to. Semantic  
36 policies play a role in determining reference, and such policies lead themselves to  
37 empirical investigation. We can ask people, “What kind of things are dogs?” or  
38 “Could something be a dog if it didn’t look like one?” Again, one can think of  
39 referring as a kind of activity. It is a complex activity that involves both sorting  
40 things and forming beliefs about the nature of the things we sort. Theories of  
41 reference should be tested and they should not be seen as psychologically inert.

42 In semantic theory, the term “naturalism” refers to attempts to provide  
43 reductive explanations. In epistemology and moral psychology, it refers to attempts  
44 to ground philosophical theories in empirical evidence. What I am suggesting is  
45 that semantic theory should be naturalized in this second sense. A theory of  
46 reference should be seen as a psychological theory—a theory of mental processes

1 and behaviors. What our concepts refer to is up to us, and if we want to figure out  
2 what reference is, we should see what we say about reference and how we engage  
3 in the activity of referring.

4 Taking these lessons on board, we can see how one would go about to  
5 determine what, if anything, our moral concepts refer to. First, we need to see what  
6 things we sort under our moral concepts. If, as I have suggested, moral concepts  
7 have an emotional basis, this amounts to the question of what things induce the  
8 relevant emotions (anger, contempt, guilt, etc.). Second, we need to see what are  
9 semantic policies that are regarding moral concepts in order to see how to extend  
10 and delimit the class of things that elicit these emotions.

11 There has been a considerable amount of empirical research on what elicits  
12 each of our moral emotions (see Haidt 2003, for a review). One set of findings was  
13 already mentioned above: anger and guilt are elicited by crimes against persons;  
14 disgust and shame are elicited by crimes against nature; and contempt, guilt, and  
15 shame are elicited by crimes against community. Many other studies add details  
16 and refinements to this big picture. Guilt is most readily induced when the person  
17 harmed is someone we care about (Baumeister et al. 1994). Disgust can be induced  
18 by crimes against persons when those crimes are especially senseless, or violent, or  
19 motivated by demeaning attitudes such as racism (Haidt et al. 1997). Shame often  
20 arises when someone imagines being judged by an audience (e.g., Ferguson et al.  
21 1991).

22 As these observations already make clear, the class of things we call “morally  
23 wrong” is something of a hodgepodge. It includes such varia as theft, rudeness,  
24 and bestiality. In this respect, WRONG is quite unlike DOG, SPOON, or POTATO. These  
25 concepts refer to categories whose members vary, to some extent, but it is fairly  
26 easy to identify intracategorical commonalities. Moreover, there is considerable  
27 cross-cultural variation in what elicits our moral emotions and gets labeled morally  
28 wrong. We morally condemn infanticide, cannibalism, slavery, honor killing,  
29 torture, polygamy, first-cousin marriage, and male dominance, but these practices  
30 are or have been widely accepted by other cultures. Even within our society, there  
31 is considerable variation: Southerners are more likely to favor corporal punish-  
32 ment (Nisbett and Cohen 1996), liberals and conservatives differ on whether 5  
33 distribution should aim for equality or equity, and members of various religious  
34 groups diverge on questions about abortion and birth control. By comparison,  
35 there is little cultural disagreement, I suspect, about what counts as a potato.

36 Moral variation poses a challenge for anyone aiming to explain what our  
37 moral concepts refer to. With natural kind concepts and many artifact concepts, a  
38 large portion of the things that we classify under a single label share some unifying  
39 properties. We can discover what category members have in common. It sometimes  
40 takes years of scientific research to do so, and there is often enough variation in the  
41 category that scientists are forced, at certain junctures, to simply stipulate how a  
42 term should be used (are whales fish?). But it is, nevertheless, fairly easy to see how  
43 unifying features might be identified, and there is plenty of reason for optimism  
44 about unifying projects before the search for essences begins. Members of natural  
45 kind concepts typically *seem* similar on the face of it, and there is considerable  
46 consensus about what objects fall into what categories.



1           When we encounter a category that is extremely heterogeneous on the face  
2 of it, there are four possibilities. One possibility is that there is a unifying essence  
3 that just is not obvious, but could be discovered through hard work. In the case of  
4 morality, that seems highly unlikely. It is not even clear how we would go about  
5 finding a unifying essence. Should we put stealing and bestiality under a micro-  
6 scope and see what they share in common? If different people classify different  
7 things as wrong, whose class of wrongs should we consider in our hunt for  
8 essences? One might think that the search for the essence of wrongness has been  
9 a central task in moral philosophy. Perhaps normative ethicists are like the chem-  
10 ists of the deontic domain. I am skeptical of this interpretation. I think normative  
11 ethicists are rarely looking for an essence that would unify moral judgments.  
12 Rather, they tend to abandon many ordinary intuitions about what is wrong and  
13 build general theories based on a handful of cases. In the West, moral philosophers  
14 have typically focused on harm norms, ignoring the fact that ordinary people often  
15 condemn acts in which no one is harmed (such as consensual incest). Some moral  
16 philosophers focus on moral rules that involve character, deliberately ignoring or  
17 postponing consideration of rules that pertain to action. Others do the opposite.  
18 The search for common essence seems as fruitless as using a washcloth for a  
19 blanket: pull it up to cover one part of your body, and the rest of your body will be  
20 exposed. If normative ethicists are trying to unify the class of things we moralize,  
21 they are not succeeding. More often, I think they are trying to develop new systems  
22 of values: abandon your folk moral attitudes, they say, and follow a more coherent  
23 set of action guiding principles. That is a noble project, but it will not help with the  
24 question at hand: What do our ordinary moral concepts refer to? My present point  
25 is that ordinary moral concepts probably do not refer to a unifying essence.

26           This brings us to a second possibility. Perhaps moral concepts simply fail to  
27 refer. If there is no intrinsic essence unifying the majority of things we classify as  
28 *wrong*, then maybe that concept is empty. Suppose I introduce the word “*proggle*”  
29 to refer to the category containing potatoes, spoons, and dogs. Suppose that you  
30 pick up the term but use it for potatoes, spoons, and spiders. Someone else uses  
31 “*proggle*” when talking about eggplants, spoons, and gorillas. At a certain point one  
32 might ask what *proggle* refers to. There is certainly no intrinsic essence, even as any  
33 one individual uses the word. That might raise the specter of elimination: there  
34 are no *proggles*. Whether this consequence follows depends the semantic policies  
35 underlying the use of the *PROGGLE* concept. Suppose I intend the concept to refer  
36 to a natural kind. If so, the fact that there is no intrinsic essence in my category  
37 suggests that my concept is incoherent. My classificatory behaviors cannot be  
38 reconciled with my semantic policies. From this it would follow that my concept  
39 does not refer.

40           Can a similar conclusion be drawn in the case of moral concepts? That  
41 depends on our semantic policies toward moral terms. Figuring out what those  
42 policies are is an empirical project. We need to investigate whether moral concepts  
43 are treated as natural kind concepts or in some other way. Do people assume that  
44 there is a unifying essence to the moral? If so, is this just a dispensable belief or is  
45 it integral to how people use moral concepts? Undoubtedly some people think  
46 there is an essence to morality. For example, some think that all moral rules issue

1 from divine decree. But three things are noteworthy about that particular belief.  
2 First, it is not universal. Second, those who have it would concede that something  
3 could still count as morally wrong if God had not commanded against it (Nucci  
4 2001). Third, divine decree is not an intrinsic essence, so those who take this to be  
5 the unifying feature of moral rules do not thereby assume that morals constitute a  
6 natural kind. Indeed, there is good reason to suppose that people do not ordinarily  
7 think of wrongness the way they think of tigers or gold. Unlike those categories, we  
8 do not establish sciences for discovering the essence of morality (this article not  
9 withstanding), and many of us do not feel compelled to consult an expert when  
10 deciding whether something is really good. With natural kinds, every alleged  
11 instance of the category can potentially be questioned: Are Chihuahuas *really*  
12 dogs? With morals, some cases seem to be non-negotiable, though these will differ  
13 from person to person. If I read a headline reporting that experts had discovered  
14 that rape is not really wrong, I would assume it was a very sick joke. Basic values  
15 just do not seem contestable. And when they are contested, we do not respond by  
16 calling in experts. Rather we reflect deeply to make sure that we really have the  
17 values that we think we have. Is my moral abhorrence of rape really basic? And if  
18 it is not, might it be up for debate after all? I sincerely doubt the answer is yes, but  
19 to find out, I need to look inward into my values, not merely outward at the world.  
20 In sum, I think we should empirically investigate whether people treat moral  
21 concepts as natural kind concepts, and I suspect the answer will be negative.

22 If moral concepts are not natural kind concepts, then variation within the  
23 range of things that people moralize will not necessarily entail conceptual inco-  
24 herence and vacuity. A third possibility (beyond intrinsic essences and vacuity) is  
25 that WRONG is a disjunctive concept, revering to some fixed list of basic wrong, and  
26 other wrong derived from those. Many traditions have lists of basic wrongs (the Ten  
27 Commandments), and we certainly possess some disjunctive concepts (recall the  
28 philosophical literature on JADE). What troubles me about this proposal is that  
29 disjunctive concepts are rare, and they usually arise in cases where we were aiming  
30 for a natural kind and ended up pointing to objects that lack intrinsic unity. JADE is  
31 a disjunctive concept because people thought that nephrite and jadeite were intrin-  
32 sically alike, and that assumption proved false. When that was discovered, we  
33 probably should have become eliminativists about jade (the concept turned out to  
34 be incoherent), but it proved useful to drop the natural kind policy, and preserve a  
35 concept that referred to two stones rather than one. In the case of moral categories,  
36 the disunity is quite apparent, unlike jade. It hardly seems that we made a mistake  
37 and thought that bestiality and promise breaking were intrinsically alike. More-  
38 over, we need a story about why we co-classify such heterogeneous behaviors in the  
39 first place. We have such a story with JADE. If we lack such a story with WRONG, it  
40 would seem that we arbitrarily grouped together a motley disjunction and smacked  
41 on a single label (recall the progle case). But if there is some feature in virtue of  
42 which we co-classify wrongs, then perhaps we should say that feature, whatever it  
43 is, provides the sought after categorical unity. The point can be made in terms of  
44 semantic policies: absent any evidence that we simply think of moral wrongs as an  
45 arbitrary list, we should look for an alternate account of how we think about this  
46 category.

1           If bestiality and promise breaking do not share intrinsic properties, how is it  
2 that we co-classify them. This question has both a proximal answer, a distal answer,  
3 and a really distal answer. The proximal answer is already anticipated by the earlier  
4 discussion. These wrongs, and all others, are united by the fact that they cause a  
5 certain class of emotions in us, emotions of blame. We would feel an other-blaming  
6 emotion if we caught someone sleeping with a sheep or breaking a promise, and  
7 we would feel a self-blaming emotion if we were caught performing these act  
8 ourselves. In this respect, very different wrongs feel alike. By comparison, consider  
9 the range of things we find funny or the range of foods we find delicious. A subtle  
10 pun and a slapstick routine may have no interesting intrinsic properties in  
11 common, but they both amuse us. A serving of tiramisu tastes nothing like the squid  
12 ink pasta in the first course, but both cause gustatory pleasure. Very different things  
13 can elicit similar feelings, and when that happens, we group those different things  
14 together. In the moral case, we can ask why such different things cause such similar  
15 emotions. This brings us to the distal story. Different actions elicit similar feelings  
16 because, during the course of moral education, we get punished for very different  
17 things. Punishment takes a number of different forms, from physical discipline to  
18 social ostracism, but all these different forms have a common effect: they condition  
19 emotional responses. When children are corrected for misdeeds, the methods of  
20 correction are emotionally evocative. Caregivers also instill emotions by express-  
21 ing outrage when morally questionable behavior is discussed. If the topic of bes-  
22 tiality come up, you can be sure that strong emotions will be expressed. Moral  
23 learners pick up on these emotions and acquire emotional dispositions. Why do  
24 caregivers condemn some actions and not others? The short answer is that they  
25 were emotionally conditioned to condemn those things, and the long answer (the  
26 really distal story) is that every value in a given culture has a genealogy. Many were  
27 introduced to promote social cohesion, some came about through happenstance,  
28 others reflect bygone superstitions, and so on. The details do not matter here (see  
29 Prinz 2007, for some examples). What matters is that there is always an explanation  
30 for why something gets moralized, and moralization consists in the cultivation of  
31 emotional dispositions. When we classify two different things as wrong, we do so on  
32 the basis of the emotions they elicit.

33           This suggests a fourth proposal about the semantic consequences of moral  
34 diversity. We could say that *WRONG* refers to a response-dependent property: the  
35 property of causing emotions of self-blame and other blame. Other philosophers  
36 have made similar proposals, most notably John McDowell (1985) and David  
37 Wiggins (1987). They liken wrongness to funniness: the funny is, arguably, that 6  
38 which causes feelings of amusement. In order to defend a response-dependent  
39 account of what moral concepts designate, it would be good to show that this  
40 conclusion follows from an independently plausible account of how concepts  
41 refer. I suggested that reference often involves two factors: a pattern of nomic  
42 co-occurrence and a semantic policy. It is trivially true that the property of causing  
43 emotions of blame causes emotions of blame in a law-like way. The property can be  
44 defined in that way. But in order for moral concepts to refer to this response-  
45 dependent property, we also have to have a semantic policy that has that implica-  
46 tion. This is harder to establish.

1 I have argued that we do not treat moral concepts as natural kind concepts  
2 or as disjunctive concepts. But is it really plausible that people think the wrong is  
3 just that which causes feelings of disapproval? I think it is more likely that we have  
4 a flexible semantic policy when it comes to moral concept. If it turned out that  
5 there was a mind-external source of morality (such as a divine decree), we would  
6 happily say that morality is an objective feature of the world. But if there is no  
7 objective feature, we are happy to treat moral facts as subjective. Response-  
8 dependence is a kind of fallback plan. One could garner empirical support for this  
9 claim by asking people (in jargon-free language) whether failure to find a mind-  
10 independent source of morality would lead them to conclude that moral claims are  
11 vacuous. I do not know for sure how such a study would come out in the moral case,  
12 but I take the outcome as quite predictable in the case of concepts such as FUNNY,  
13 DELICIOUS, or BEAUTIFUL. Suppose there is no intrinsic essence to things that amuse  
14 us; should we say that nothing is funny? Suppose there is no essence to things we  
15 find delicious or beautiful; should we say these properties do not exist? No. Instead  
16 we people have different senses of humor, we say that food preferences are a  
17 matter of taste, and beauty is in the eye of the beholder. These platitudes reveal a  
18 tendency to treat emotionally based concepts as response-dependent, when we find  
19 that they do not correspond to mind-independent features of the world. The same  
20 fallback strategy may be in place with color concepts. Locke argued that these refer  
21 to secondary qualities—powers to cause certain experiences in us.

22 There is a plausible diagnosis of why we seem to default to response-  
23 dependence in all these cases. Even if we think surfaces are really colored and good  
24 jokes are intrinsically funny, we also know that color and funniness are manifest  
25 features of the world. They are not hidden posits that we can reveal only through  
26 scientific instruments. In this they differ markedly from concepts such as WITCH,  
27 FLOGISTON, and other concepts introduced to refer to (especially unobservable)  
28 natural kinds. We may have a semantic policy that says something like “save  
29 appearances!” If some apparent property turns out not to have a mind-  
30 independent essence, we do not say the apparent property does not exist. Instead,  
31 we say it is merely apparent; in other words, response-dependent. This explains why  
32 people continue to believe in the existence of colors, beauty, and funniness.

33 So the trick to determine whether moral properties are response-dependent  
34 is not to see whether people antecedently believe that. Some people may think  
35 morality has a mind-independent source. The real question is whether people think  
36 of morality as a manifest feature of the world. Do they think that any empirical  
37 discovery could show that morality does not exist? Would the discovery that moral  
38 values come from us entail that there are no moral truths? To my knowledge,  
39 no one has directly investigated lay opinions about such questions. One piece  
40 of evidence comes from the fact that there seems to be considerable diversity of  
41 opinion about where morality comes from. Some claim to believe in divine  
42 command, others do not. Some people are moral absolutists, and some are relativ-  
43 ists. In a study of moral objectivist, Nichols (2004) showed that people are divided  
44 on the question of whether there is a fact of the matter who is right in moral  
45 disputes. Kelly et al. (2007) found similar results, with over half of their subjects  
46 saying that it is okay to hit someone if the local authorities allow it. The individual

1 differences suggest that moral concepts are semantically flexible. Some people  
2 have strong convictions about whether morality is mind-independent or mind-  
3 dependent property, but these convictions are not semantic policies. They are just  
4 collateral beliefs about the nature of moral properties. If someone is wrong about  
5 whether morality depends on human responses, that does not entail eliminativism.  
6 In other words, the variation is that beliefs about the nature of morality can be  
7 taken as evidence for the view that we regard morality as a manifest kind. And,  
8 if that is right, we will try to save appearances if we discovery that there are no  
9 mind-independent moral facts.

10 More evidence is needed before any confident conclusions can be drawn  
11 about moral reference, but let me sum up with the conclusions that seem most  
12 plausible to me. First, the diversity of moral values problematizes efforts to iden-  
13 tify a mind-independent essence of moral wrongs. So I do don't think WRONG  
14 refers to a mind-independent property. Second, that does not entail eliminativ-  
15 ism, because the assumption of mind-independence is not a semantic policy when  
16 it comes to moral concepts. Third, the best account of what WRONG refers to,  
17 given our semantic policies and the range of things we call "wrong," is a  
18 response-dependent account: WRONG refers to the property of causing emotions  
19 of self-blame and other-blame. Such an account would need to be refined in  
20 various ways (whose response matter? under what conditions?), but it will serve  
21 for present purposes (see Prinz 2007).

### 22 **STEP 3: FROM MORAL TRUTH TO OBLIGATIONS**

23  
24  
25 Suppose I am right, and WRONG refers to a response-dependent property. It then  
26 follows that there are two ways to figure out what is wrong. First, there is a  
27 first-person method. If something seems to cause emotions of blame in me that is  
28 direct evidence that it is wrong. It is not decisive evidence. I may misinterpret my  
29 feelings, or those feelings could derive from an extraneous source, such as a dirty  
30 desk. But, all things being equal, if something seems to make me feel angry and  
31 guilty, it probably has the disposition to cause these emotions in me, and that  
32 entails that it is wrong. From a first-person perspective, we use our emotions as a  
33 guide to morality, and they are a pretty reliable guide because moral properties are  
34 constituted by powers to cause emotions. Phenomenologically, we do not experi-  
35 ence this as an inference. We do not typically say to ourselves, "that made me mad,  
36 so it must be wrong." Rather, the anger constitutes our belief that something is  
37 wrong. So, from a first-person perspective, there is often an immediate, causal  
38 transition from perception of an event to representation of that event as wrong, by  
39 means of an emotional response. Thus, moral judgments appear to us as intuitions,  
40 rather than inferences. Put another way, the state transition from perceiving an  
41 event to condemning it is not a deductively valid inference; it is a causal chain. So,  
42 from the first-person perspective, we do not *deduce* moral facts from descriptive  
43 facts.

44 But there is also a second way to figure out whether something is wrong.  
45 From a third-person perspective, we can investigate what things in the world  
46 are disposed to cause emotions of blame. One could adopt this third-person

1 perspective with regard to oneself. For example, a therapist might help me see that  
2 I get mad when my spouse stays out late at night, and this insight might lead me to  
3 conclude that I must be regarding that behavior as wrong. If I also believe that  
4 WRONG refers to that which is disposed to cause emotions of blame, I may conclude  
5 that my spouse really has done something wrong.

6 This kind of inference can also take place when we are investigating other  
7 people. Anthropologists studying the Gisu of Uganda have discovered that  
8 members of this group show moral outrage when someone does not avoid his or  
9 her mother-in-law (Heald 1990). If wrongness is the property of being disposed to  
10 cause such responses, then the anthropologist can conclude that it is wrong for  
11 members of the Gisu to visit their spouses' mothers.

12 When we make inference of this kind, we are discovering wrongs empirically.  
13 Purely descriptive premises go into the discovery process, and these deductively  
14 entail that something is wrong. We observe that someone has an emotional  
15 response to something; we know that wrongness consists in causing such a  
16 response; and we infer that the thing causing the response is wrong. The argument  
17 makes use of the semantic theory I argued for in the last section, but a similar  
18 argument could be constructed for different theories of reference. For any theory  
19 of reference and for any theory of concepts, there is some relation R that obtains  
20 between a concept and what it represents. If it is a naturalistic theory (in either of  
21 the senses mentioned above), one can empirically investigate what bears R to the  
22 concept WRONG. Thus, one can empirically discover wrongs.

23 It follows trivially that one can empirically discover obligations. I do not want  
24 to offer a complete analysis of the concept OUGHT here, but one entailment relation  
25 strikes me as relatively uncontroversial: if something is wrong, then one ought not  
26 do it. If one can discover wrongs empirically, and one can derive oughts from  
27 wrongs, then one can discover oughts empirically. One can derive an ought from  
28 an is.

29 The argument that I have been developing is easy to summarize. It goes like  
30 this:

- 31 P1. The concept WRONG is constituted by emotions.  
32 P2. Emotions refer to response-dependent properties.  
33 P3. There is some person P, such that some type of action A, disposes P to feel  
34 the emotion constituting the concept WRONG.  
35 C1. Therefore, doing A is wrong for P.  
36 P4. "It is wrong for *x* to do *F*" entails "*x* ought to refrain from doing *F*."  
37 C2. Therefore, P ought to refrain from doing A.

38 I suspect many readers will not be convinced by this argument. I have given only  
39 preliminary arguments for the core premises, but, as I just noted, a similar argu-  
40 ment could be constructed for different theories of moral concepts and moral  
41 semantics. For that reason, I will not spend time reviewing objections that focus too  
42 heavily on specific premises. But I must consider two objections. According to the  
43 first, the conclusion of this argument may seem so obviously wrong that no argu-  
44 ment could possibly support it. According to the second, the conclusion can only be  
45 defending on a reading that robs it of normative force.



**OBJECTION 1: OBLIGATING EVIL**

1  
2  
3 The first objection I want to consider concerns the relativist implications of the  
4 view that I have been defending. If the truth of wrongness judgments depends  
5 on emotional dispositions, and different people are disposed to have different  
6 emotional responses, then the truth moral judgments will vary from individual to  
7 individual. I think this is an unavoidable consequence of the theory, and one that  
8 we must be prepared to accept if we are open to the idea that the semantics of  
9 moral concepts can be empirically studied. Here, as elsewhere, we are hostage to  
10 empirical fortune, and I think relativism is supported by the evidence. Relativism  
11 is a hugely controversial thesis, which I cannot defend here (see Prinz 2007). But,  
12 there is a worry that I must address here. On the face of it, the relativism entailed  
13 by my theory of moral concepts seems to undercut my argument for deriving  
14 obligations from descriptive premises. That would be a serious embarrassment.

15 Here is the worry. We can empirically discover that members of some other  
16 group take themselves to be obligated to do something. For example, we can  
17 discover that the Gisu take themselves to be obligated to avoid their mother-in-  
18 laws. But from this it does not follow that they ought to avoid their mother-in-laws.  
19 The point becomes all the more obvious when we consider cultural groups that  
20 have values that we find abhorrent. Can we say of the extremist Rwandan Hutus  
21 that they ought to have massacred the Tutsis? Can we say of those who engage in  
22 honor killings, such as some traditional members of the Yedizi religious group,  
23 that they ought to stone their sister's and daughter's to death when they become  
24 romantically involved with members of other groups? Can the mere fact that  
25 someone sincerely believes that they are obligated to do something horrific entail  
26 that the person really ought to? If not, then something has gone wrong with the  
27 argument above.

28 I think this worry is more moral than philosophical. We do not like the idea  
29 that people are obligated to do things that we find morally repellent. If the theory  
30 has that consequence, we will regard it as very unfortunate, but that does not make  
31 the theory wrong. To make the objection stick, one would have to convert disdain  
32 for the conclusion into a substantive objection. One strategy for doing that is to  
33 argue that the conclusion of the argument can generate contradictions. The argu-  
34 ment suggests that people ought to engage in honor killings, after all failure to  
35 engage in that practice causes strong emotions of blame in people who abide by  
36 that custom. But it also follows from the argument that one ought not engage in  
37 honor killing. Some people (myself included) find the practice reprehensible, and  
38 that implies that it is wrong, and if it is wrong, one ought not do it. This seems like  
39 a contradictory pair of entailments.

40 This objection is not decisive, because the apparent contradiction can be  
41 eliminated in three ways. First, one can follow Gilbert Harman (1975) and say that  
42 the word "ought" is semantically restricted so as to apply only when the person to  
43 whom it is applied possesses the values in question. On this interpretation, we  
44 cannot say that the Yedizi ought to refrain from honor killings, so we cannot  
45 generate a contradiction. Second, one can relativize obligations to value systems.  
46 We can say that the Yedizi are obligated to engage in honor killings on their value

1 system and not ours. Third, we can eliminate the contradiction by placing the  
2 negation inside the obligation. In other words, we can say that the Yedizi are  
3 obligated to engage in honor killing and not to engage in honor killing. This  
4 formulation escapes the *logical* contradiction, even if it places the Yedizi under  
5 *practically* inconsistent requirements. The narrow scope of “not” is justified on the  
6 grounds that saying “not obligated” implies that there is no obligation, which is, if  
7 the semantic theory is right, false.

8 But there is another semantic worry that is harder to avoid. Arguable, the  
9 concept OUGHT can only be used to express one’s own values. Like WRONG, the  
10 concept OUGHT may have an emotional basis. When we say that someone ought to  
11 do something, we may be expressing how we feel. I certainly do not feel good about  
12 honor killing, and, consequently, it is not just odd for me to say that people ought  
13 to engage in this practice; such a statement would actually semantically anomalous.  
14 It would be equivalent to me saying that fine Bordeaux wines are yucky. A child  
15 might call a Bordeaux “yucky” but I cannot apply this word without feeling the  
16 disgust that the word conventionally expresses. Likewise, I cannot say someone  
17 ought to engage in honor killing if I do not have the feelings that conventionally  
18 constitute ought thoughts. (And, perhaps I cannot say that something is wrong  
19 unless it is wrong according to my values.) If this intuition about the concept of  
20 OUGHT is right, then there cannot be any deductive argument that uses “ought” as  
21 the main verb in a conclusion without having some other evaluative term used in  
22 the premises. There can be no derivation of an ought from an is.

23 This is a powerful objection, but it can be answered. I offer three lines of  
24 reply, which range from very concessive to not concessive at all. First, I might  
25 concede that no ought can be derived from an is, while maintaining something  
26 weaker: we can derive an *obligation* from an is. Perhaps we can say that people who  
27 support honor killings are under an obligation to kill, while denying that they ought  
28 to kill. To my ear, this does not sound inconsistent. On this reply, I can continue to  
29 claim that obligations are empirically discoverable, which is an interesting result.

30 The second reply is slightly less concessive. I might concede that we cannot  
31 derive conclusions about what other people ought to do, but I could derive conclusions  
32 about what I ought to do. Recall that we can take a third-person stance  
33 toward ourselves. Suppose that I am undergoing a severe depression, and as a  
34 result I am not feeling strong emotions of any kind, whether negative or positive  
35 (see Mele 1996). I can still discover, say by reflecting on my past behavior, that  
36 certain actions are disposed, when I am healthy, to cause emotions of blame.  
37 Knowing this fact about myself, I can infer that those actions are wrong, and that  
38 I ought to avoid them. Here, the “ought” is semantically licensed because I am  
39 disposed, under better conditions, to feel the emotions that normally accompany  
40 the word. If this reply is right, the one can derive an ought from an is, as long as the  
41 argument is stated in the first-person.

42 The third reply is least concessive. Perhaps there are two uses of “ought,” one  
43 of which is expressive, the other of which is not. By analogy, there may be two uses  
44 of delicious. I can say that chicha (a fermented beverage made from corn and  
45 saliva) is delicious to Amazonian people while denying that it is delicious to me.  
46 This use of the word does not express any feelings. But if I drink some Bordeaux

1 as say “that’s delicious!” I have thereby expressed how I feel. Notice that this  
2 expressive use would be the default reading, for pragmatic reasons, if I did not  
3 specify a context. If I say, “Chichi is delicious,” I imply that it is delicious to me. But  
4 I can cancel this implication by saying that it is delicious to Amazonian people.  
5 Likewise, there may be a nonexpressive use of ought. I may be able to cancel the  
6 implicature that I am deploying the expressive use by saying, “You ought to avoid  
7 your mother-in-law if you are a member of the Gisu.” If “ought” has a nonexpres-  
8 sive use, then I can derive an ought from an is. It is an empirical question whether  
9 the word “ought” ever allows such a nonexpressive use. One might investigate by  
10 polling intuitions about sentences like the one I just formulated about the Gisu.

11 I am confident about the first two replies, and less sure about the third. At a  
12 minimum, then, I would say we can derive an ought from an is when talking about  
13 our own values, and we can derive an obligation from an is when talking about  
14 people whose values differ from our own.

## 15 **OBLIGATION 2: ANEMIC OUGHTS**

16  
17  
18 I have been arguing that we can use empirical methods to discover obligations.  
19 Purely descriptive premises reveal the demands that moral values place on us.  
20 Critics may find my derivation of norms somewhat unsatisfying. In my introduc-  
21 tion, I hinted that empirical findings could be used to draw normative conclu-  
22 sions. A normative conclusion is a claim about what we ought to do. The  
23 argument that I have put forward seems both too powerful and too weak for that  
24 purpose. It is too powerful because it seems to prove too much. It can be used to  
25 show that all of us are obligated to do whatever anyone thinks we should do. If  
26 someone thinks that I am morally obligated to indulge in cannibal feasts every  
27 February, then the argument that I put forward seems to entail that there is such  
28 an obligation. But it would make a joke of normative theory to say that we are  
29 all obligated to do everything that any moral system required. In this sense, the  
30 argument is too strong to do any work in normative ethics. On the other hand,  
31 it is also too weak. I conceded in the last section that the argument I have put  
32 forward cannot be used to adduce an other-directed ought judgment stated with  
33 emotional force. This is the kernel of truth in Hume’s law: you cannot derive an  
34 oomph from an is. If the oughts that one can derive have no oomph, then they  
35 cannot be taken as recommendations for action. When I say that the Gisu ought  
36 to avoid the mother-in-laws, I am not recommending that they do that. I am, in  
37 this sense, not really making a normative claim. So perhaps my introduction was  
38 an exercise in false advertising. Perhaps I have not shown how to draw normative  
39 conclusions after all. The oughts that can be deduced from empirical premises  
40 are too anemic.

41 This objection reveals a genuine limitation of the argument that I have been  
42 defended, but there is also an available reply. I just implied that my argument for  
43 deriving oughts entails that everyone is obligated to do everything that anyone  
44 happens to think they should do. But this is actually a mistake. As formulated  
45 above, the argument is actually restricted to norms that are possessed by the person  
46 upon whom the derived obligation is placed. The argument entails that if a person,

1 P, is disposed to have a certain emotional response to an action, A, then P ought not  
2 do A. It is silent on the question of whether P falls under obligations that she or he  
3 would not endorse. I have not said that every agent ought to do whatever is  
4 required on other agents' moral theories. The argument does not require that I  
5 ought to avoid my mother-in-law, for example. It entails only that each person  
6 ought to do what his or her own values require.

7 This restriction provides a response to the normativity objection. Notice, that  
8 from P's point of view, P ought to do what P's norms require. So the ought in the  
9 conclusion has normative force in one important sense: it is an ought that would be  
10 psychologically compelling to P. Trivially, we feel obligated to follow our own  
11 values. Indeed, we are bound by our values even when we do not realize what our  
12 values require of us. If P does not realize she ought to perform some action by her  
13 own system of values, it still follows that she ought to perform that action; she  
14 is committed to norms that carry that requirement. Even if the "ought" in the  
15 conclusion is not, semantically speaking, a normative ought (because the person  
16 drawing the inference may not have the requisite feelings), it is an ought that has  
17 normative jurisdiction over the person to whom it applies. So whenever the argu-  
18 ment is applied there is a corresponding ought with oomph.

19 This kind of normativity is weaker than the kind that many philosophers  
20 seek. Philosophers often want norms that are binding *regardless* of our desires  
21 and inclinations. I am skeptical that there are such norms in the moral domain.  
22 My skepticism stems from my conviction that moral concepts refer to response-  
23 dependent properties. If the arguments for that claim succeed, then the idea of  
24 moral norms that transcend our inclinations will be difficult to sustain. Thus, I think  
25 the revised formulation of the argument may yield conclusions that track the only  
26 kind of normativity we should expect to find in the moral domain. The conclusion  
27 of the argument might not be stated in an emotionally expressive way by the  
28 person making the argument, but the ought would be binding.

29 This suggestion may seem disturbing at first. I have suggested that there is a  
30 kind of normativity such that everyone is bound by the values they possess. The  
31 Gisu are required not to visit their mother-in-laws, and members of the Yedizi are  
32 required to kill those who dishonor them. That is an unsettling consequence of a  
33 theory. But notice that when I observe the existence of these requirements, I am  
34 not expressing my own values. I think honor killing is a ghastly institution. Indeed,  
35 I have been arguing that all of us are bound by our own values, so I am bound to  
36 condemn honor killing.

37 Against this suggestion, one might object that people can question whether  
38 they are really bound by the values that they happen to endorse. A member of the  
39 Yedizi can say, "I know I morally value honor killing, but should I? Is it really  
40 something I am obligated to do?" It is a consequence of my theory that this kind  
41 of question is really incoherent. That may look like a fatal objection to the theory,  
42 but it is not an embarrassment if there is a reasonable explanation of why such  
43 questions seem intelligible to the people who formulate them. Fortunately, a rea-  
44 sonable explanation is available. Recall that I am not committed to the view that  
45 people realize morality is a response-dependent property. People can (and often do)  
46 believe that morality is mind-independent. Someone who endorses a divine

1 command theory, for example, can wonder, "Are the values I endorse really God's  
2 requirements or were they planted by some kind of devil?" Someone who wonders  
3 whether her or his values are binding may be in the grip of an erroneous theory of  
4 where values come from.

5 There is also another explanation of how we can question of values. Suppose  
6 I come to recognize that my values are response-dependent. After I make this  
7 discovery, I can still wonder whether the values I have are the best possible values  
8 that I could have. For example, I could wonder whether my moral values lead me  
9 to flourish. I can then ask, in a prudential voice: "Should I do what my morals tell  
10 me to do, or should I pursue actions that will bring me greater well-being?" I may  
11 also recognize that I have inconsistent values, and that will lead me to think that I  
12 am morally obligated to refrain from doing something that I am also obligated to  
13 do. A member of the Yedizi might believe in the basic sanctity of human life, and  
14 this could lead to the thought: "I morally value honor killing, but I ought not  
15 engage in that practice because I am also morally obligated to protect human life."  
16 In sum, I think we are morally obligated (by our own standards!) to do whatever  
17 our moral values require of us, but we sometimes fall under competing obligations  
18 because we are morally inconsistent or because we possess nonmoral values that  
19 are difficult to reconcile with morality. Thus, the fact that people question whether  
20 they should do what their moral values demand is easily explained. It can reflect an  
21 understandable mistake about moral ontology, or it can reflect the recognition that  
22 values compete.

23 The process of reviewing our values, searching for inconsistencies, and decid-  
24 ing what to do is a central part of normative ethics. The crucial thing to notice about  
25 this process is that we do it from within a system of values. We typically begin with  
26 something that we already value and try to draw out its implications. We cannot  
27 adopt a transcendental stance, and if we could, we would not feel compelled to  
28 endorse the values that emerged from such a stance. I think philosophical work in  
29 normative ethics is an internal, value-expressive project, and normative theories  
30 are usually systematic extrapolations of values that the intended readers of such  
31 theories already endorse.

32 When building normative theories, philosophers sometimes select a small set  
33 of antecedently accepted values, and those values are usually discovered through  
34 introspection. Introspection is essentially an empirical method. We access our  
35 mental states through a process of internal observation. But introspection is a  
36 limited method; we do not always observe accurately, we are prone to theory-laden  
37 interpretations; we do not have conscious access to all of our thoughts, and the  
38 values revealed by introspection may be unique to the introspector. Moreover, the  
39 practice of cherry picking from antecedently accepted values results in normative  
40 theories that are parochial: they capture some of our moral commitments at the  
41 expense of others.

42 I think philosophical approaches to normative theory construction can be  
43 supplemented by the methods used in the empirical sciences. Experimental psy-  
44 chology, for example, can help us discover our normative commitments. Psycholo-  
45 gists do this in a variety of ways: sometimes they ask people to rate morally  
46 significant scenarios, sometimes they measure emotional responses, sometimes

1 they observe behavior, and so on. We can discover some of our own values in this  
2 way, and in some cases we are surprised. We may find that our values are incon-  
3 sistent; we may find that some values are more fundamental than others; we may  
4 discover the comparative weightings of values; and we may identify factors (such as  
5 dirty desks) that skew our value judgments. All of these discoveries would be  
6 relevant to constructing normative theories. Normative ethics can systematize the  
7 values we already possess, and reduce inconsistencies by determining which of two  
8 competing values is more important to us. I think we should actively pursue all  
9 methods of discovering what our values are.

10 Put differently, the argument for deriving oughts takes on a special role in the  
11 first-person case. When we discover what we value, we discover the norms that bind  
12 us—the things that, by our own standards, we ought to do. In other words, when we  
13 empirically investigate the things that arouse our emotions of blame, we are con-  
14 tributing to a normative project. The argument shows that we can discover norma-  
15 tive oughts. When the agent in question is me, the argument shows that I can  
16 discover what I ought to do on the basis of empirically discovered facts. The adage  
17 that descriptive methods cannot contribute to normative projects is a mistake. The  
18 psychological sciences reveal not only what we do value but also, thereby, what we  
19 ought to do. The idea that there are moral oughts that exist independently of what  
20 we do in fact care about may be a myth. If so, we may benefit from reconceiving  
21 normative ethics as a kind of auto-anthropology.

## 22 REFERENCES

- 23  
24  
25 Baumeister, R. F., Stillwell, A. M., and Heatherton, T. F. 1994. "Guilt: An Interpersonal Approach." *Psychological Bulletin* 115: 243–67.
- 26  
27 Blair, R. J. R. 1995. "A Cognitive Developmental Approach to Morality: Investigating the Psycho-  
28 path." *Cognition* 57: 1–29.
- 29  
30 Ferguson, T. J., Stegge, H., and Damhuis, I. 1991. "Children's Understanding of Guilt and Shame." *Child Development* 62: 827–39.
- 31  
32 Fodor, J. A. 1990. *A theory of content*. Cambridge, MA: MIT Press.
- 33  
34 Greene, J. D., Sommerville, R. B., Nystrom, L. E., Darley, J. M., and Cohen, J. D. 2001. "An fMRI  
35 Investigation of Emotional Engagement in Moral Judgment." *Science* 293: 2105–8.
- 36  
37 Haidt, J. 2003. "The Moral Emotions." In *Handbook of affective sciences*, eds. R. J. Davidson, K. R.  
38 Scherer, and H. H. Goldsmith, 852–70. Oxford: Oxford University Press.
- 39  
40 Haidt, J., Rozin, P., McCauley, C., and Imada, S. 1997. "Body, Psyche, and Culture: The Relationship  
41 of Disgust to Morality." *Psychology and Developing Societies* 9: 107–31.
- 42  
43 Harman, G. 1975. "Moral Relativism Defended." *Philosophical Review* 84: 3–22.
- 44  
45 Heald, S. 1990. "Joking and Avoidance, Hostility and Incest: An Essay on Gisu Moral Categories." *Man* 25: 377–92.
- 46  
47 Kelly, D., Stich, S., Haley, K., Eng, S., and Fessler, D. 2007. "Harm, Affect and the Moral/  
48 Conventional Distinction." *Mind and Language* 22: 117–31.
- 49  
50 King, J. A., Blair, R. J. R., Mitchell, D. G. V., Dolan, R. J., and Burgess, N. 2006. "Doing the Right  
51 Thing: A Common Neural Circuit for Appropriate Violent or Compassionate Behavior." *Neuroimage* 30: 1069–76.
- 52  
53 McDowell, J. 1985. "Values and Secondary Qualities." In *Morality and objectivity*, ed. T. Honder-  
54 ich, ●●–●●. London: Routledge & Kegan Paul.
- 55  
56 Mele, A. R. 1996. "Internalist Moral Cognitivism and Listlessness." *Ethics* 106: 727–53.
- 57  
58 Moll, J., de Oliveira-Souza, R., Eslinger, P. J., Bramati, I. E., Mourao-Miranda, J., Andreiuolo, P. A.,  
59 and Pessoa, L. 2002. "The Neural Correlates of Moral Sensitivity: A Functional Magnetic  
60



- 1           Resonance Imaging Investigation of Basic and Moral Emotions." *Journal of Neuroscience* 22:  
2           2730–36.
- 3           Nichols, S. 2004. "After Objectivity: An Empirical Study of Moral Judgment." *Philosophical*  
4           *Psychology* 17: 5–28.
- 5           Nisbett, R. E., and Cohen, D. 1996. *Culture of Honor: The Psychology of Violence in the South*.  
6           Boulder, CO: Westview Press.
- 7           Nucci, L. P. 2001. *Education in the Moral Domain*. Cambridge: Cambridge University Press.
- 8           Prinz, J. J. 2002. *Furnishing the Mind: Concepts and Their Perceptual Basis*. Cambridge, MA: MIT  
9           Press.
- 10          ———. 2007. *The Emotional Construction of Morals*. Oxford: Oxford University Press.
- 11          Rozin, P., Lowery, L., Imada, S., and Haidt, J. 1999. "The CAD Traid Hypothesis: A Mapping  
12          Between Three Moral Emotions (Contempt, Anger, Disgust) and Three Moral Codes (Com-  
13          munity, Autonomy, Divinity)." *Journal of Personality and Social Psychology* 76: 574–86.
- 14          Schnall, S., Haidt, J., and Clore, G. Submitted. "Irrelevant Disgust Makes Moral Judgment More  
15          Severe, for Those Who Listen to Their Bodies." 9
- 16          Sinnott-Armstrong, W. 2006. "Moral Intuitionism Meets Empirical Psychology." In *Metaethics*  
17          *after Moore*, eds. T. Horgan and M. Timmons, 339–65. New York: Oxford University Press.
- 18          Valdesolo, P., and DeSteno, D. 2006. "Manipulations of Emotional Context Shape Moral Judg-  
19          ment." *Psychological Science* 17: 476–77.
- 20          Wheatley, T., and Haidt, J. 2005. "Hypnotically Induced Disgust Makes Moral Judgments More  
21          Severe." *Psychological Science* 16: 780–84.
- 22          Wiggins D. 1987. "A Sensible Subjectivism." In *Needs Values, Truth: Essays in the Philosophy of*  
23          *Value*, ed. ••, 185–214. Oxford: Blackwell. 10