PHILOSOPHICAL PSYCHOLOGY, VOL. 17, NO. 1, 2004

Routledge
Taylor & Francis Group

After objectivity: an empirical study of moral judgment

SHAUN NICHOLS

ABSTRACT *This paper develops an empirical argument that the rejection of moral objectivity leaves important features of moral judgment intact. In each of five reported experiments, a number of participants endorsed a nonobjectivist claim about a canonical moral violation. In four of these experiments, participants were also given a standard measure of moral judgment, the moral/conventional task. In all four studies, participants who respond as nonobjectivists about canonical moral violations still treat such violations in typical ways on the moral/conventional task. In particular, participants who give moral nonobjectivist responses still draw a clear distinction between canonical moral and conventional violations. Thus there is some reason to think that many of the central characteristics of moral judgment are preserved in the absence of a commitment to moral objectivity.*

> Were I not afraid of appearing too philosophical, I should remind my reader of that famous doctrine, supposed to be fully proved in modern times, "That tastes and colors, and all other sensible qualities, lie not in the bodies, but merely in the senses." The case is the same with beauty and deformity, virtue and vice … And as it is certain, that the discovery above-mentioned in natural philosophy, makes no alteration on action and conduct; why should a like discovery in moral philosophy make any alteration? (Hume, 1742/1987, p. 166)

1. Introduction

If we become convinced that morality lacks objective moorings, what will the consequences be for our commonsense moral judgment? Would a rejection of moral objectivism engender rampant nihilism? Or would the rejection of moral objectivism leave our normative lives relatively unfazed? Even at a popular level, these questions have real currency. People worry that the abandonment of moral objectivity threatens to unravel the moral tissue of society. Meanwhile, professional philosophers debate the implications of nonobjectivism for the status of commonsense moral claims. "Error theorists" maintain that commonsense is mistakenly committed to the view that morality is objective and that, as a consequence, lay moral concepts are

Shaun Nichols, Department of Philosophy, College of Charleston, Charleston, SC 29424 USA, email: nichols@cofc.edu

ISSN 0951-5089/print/ISSN 1465-394X/online/04/010000–00 © 2004 Taylor & Francis Ltd
DOI: 10.1080/0951508042000202354

empty and all commonsense moral judgments are false. In this paper, I will make some initial steps towards exploring these issues in a resolutely empirical manner.

The paper will have a rather unusual structure, with experiments folded into the argumentative flow, so I will supply an unusually detailed introductory summary. To begin, the rejection of objectivity only looms important against the background assumption that commonsense is committed to the view that morality is objective. Analytic philosophers who discuss the issue typically maintain that commonsense is indeed committed to moral objectivity. However, there has been some resistance to this claim, on the grounds that many undergraduates seem to profess moral nonobjectivism. Such undergraduates are often dismissed as being confused, and this is allegedly reflected by the fact that undergraduates will also profess a broad metaphysical nonobjectivism. In Section 3, I'll report preliminary evidence indicating that a significant number of undergraduates endorse a nonobjectivist claim about a canonical moral violation without endorsing a broad metaphysical nonobjectivism. However, it's still plausible, I argue, that there is an important sense in which commonsense is committed to moral objectivity. Evidence from developmental psychology indicates that children are objectivists about moral properties (but not about conventional or response-dependent properties). As a result, I suggest that moral objectivity is a default setting on commonsense metaethics, but the setting is, as indicated by the evidence on undergraduates, defeasible. In Section 5, I will turn to the debate over error theory. I'll argue that in order to determine whether error theorists are right that lay moral concepts are empty, we would need to settle long-standing disputes in philosophy of mind. However, even without settling those disputes, we can begin to address in an empirical way how the rejection of moral objectivity affects moral cognition. In Sections 6–9, I report four experiments exploring moral judgment in undergraduates who respond as nonobjectivists about a canonical moral violation. In all four studies, the moral judgment of those who respond as nonobjectivists seems very much like the moral judgment of those who respond as moral objectivists. Thus there is some reason to think that moral judgment remains largely intact even in the absence of a commitment to objectivity. Even in the population of individuals who respond as moral nonobjectivists, many of the central characteristics of moral judgment are preserved. Of course, a host of natural objections arise for the empirical case made here. In Section 10, I briefly discuss some of these objections.

2. Are people moral objectivists?

Like most great philosophical questions, the issue of moral objectivity has its roots in commonsense. Attacks on moral objectivism are thought to threaten a commonsense conviction that morality is objective. Analytic philosophers have explicitly maintained that people are moral objectivists in some sense. Unfortunately, the term "moral objectivity" has accumulated multiple definitions, and we will need to do some tidying before long. But first, let's consider some characteristic statements among analytic philosophers to the effect that lay people are moral objectivists. In

The Moral Problem, Michael Smith (1994, p. 6) claims that one of the central features of commonsense morality is the "objectivity of moral judgment,"

> we seem to think moral questions have correct answers; that the correct answers are made correct by objective moral facts; that moral facts are wholly determined by circumstances and that, by engaging in moral conversation and argument, we can discover what these objective moral facts determined by the circumstances are.

Similarly, Darwall (1998, p. 25) writes, "Ethical thought and feeling have 'objective purport.' From the inside, they apparently aspire to truth or correctness and presuppose that there is something of which they can be true or false." He defines "objective purport" as follows: "Seeming to be of something objective and independent of the perceiver (e.g. some objective fact or an objective property of some substance)" (Darwell, 1998, p. 239).

Many of the philosophers who maintain that commonsense is committed to moral objectivity are themselves sympathetic to moral objectivity (e.g. Brink, 1989; Darwall, 1998; Smith, 1994). However, even philosophers who are not so sympathetic maintain that lay people are moral objectivists. Indeed, perhaps the best-known opponent of moral objectivism in recent years, J.L. Mackie, goes to some lengths to establish that ordinary people are committed to the "claim to objectivity":

> The ordinary user of moral language means to say something about whatever it is that he characterizes morally, for example a possible action, as it is in itself, or would be if it were realized, and not about, or even simply expressive of, his, or anyone else's relation to it. But the something he wants to say is not purely descriptive, certainly not inert, but something that involves a call for action or for the refraining from action, and one that is absolute, not contingent upon any desire or preference or policy or choice, his own or anyone else's. (Mackie, 1977, p. 33)

Thus, among analytic philosophers of different minds about the status of moral objectivity, we find an important island of agreement about the commonsense view—both objectivists and their opponents can agree that commonsense is committed to moral objectivity [1].

At this point, we can no longer forestall the task of stipulating a sense of "moral objectivism" that will sustain us throughout the remainder of the paper. The first part of Mackie's above analysis will work nicely for a start. To claim that an action is *objectively* immoral is to claim that the action is wrong "as it is in itself" and not in relation to subjects. Cashing out "as it is in itself" is no pleasant undertaking, but we might focus on the negative claim, which Mackie seems to regard as an entailment. If morality is objective, then being morally wrong is *not* determined by a relation subjects bear to the action. Thus, to claim that an action is objectively immoral is to claim that the action is wrong "as it is in itself" and not in relation to other subjects. There are various quibbles that might be made over this characterization, but the underlying idea is familiar. According to the objectivist, if a particular action is morally wrong, then it is wrong *simpliciter*. So morally wrong actions are not

merely wrong relative to certain populations. This point is easiest to appreciate by focusing on a particular example. Let's say that a teenage boy, Tom, intentionally kicks a small dog. It cannot turn out, according to the objectivist, that Tom's kicking the dog was morally wrong for some populations but not for other populations. If the action is morally wrong, it's wrong full stop. Thus, objectivism about moral wrongness is committed to the view that (i) a morally wrong action is wrong *simpliciter*, not merely relative to certain populations and (ii) some actions are morally wrong. Now, if we look back over the philosophical characterizations of moral objectivity, it is plausible that Smith, Darwall and Mackie all regard commonsense as committed to something like this notion of moral objectivism.

Although analytic philosophers routinely speak on behalf of commonsense's commitment to moral objectivity, empirically-minded philosophers have expressed skepticism about whether analytic philosophers pay adequate attention to how morality is actually viewed by lay people, and in particular, undergraduates (e.g. Harman, 1985; Stich & Weinberg, 2001). Many undergraduates seem explicitly to disavow moral objectivism at least for some standard moral violations. Nonobjectivism among undergraduates is commonly dismissed by drawing attention to the fact that some college students will proclaim a broad metaphysical nonobjectivism in which they maintain that there is no fact of the matter even about whether the earth is flat. It is likely, the dismissal continues, that this endorsement of metaphysical nonobjectivism ensues from a naïve confusion, and that this confusion lies behind the students' apparent departure from moral objectivism. Hence, we can safely disregard the undergraduate nonobjectivist.

Undergraduates who endorse metaphysical nonobjectivism should, perhaps, often be regarded as confused. But this parry does not adequately address the issue. For it is trivial to construct experimental materials that will allow one to exclude students who espouse nonobjectivism about whether the earth is flat. Under those circumstances, one can easily see whether a substantial number of students respond as nonobjectivists about canonical moral violations without responding as metaphysical nonobjectivists. In an ongoing series of recent experiments, I've undertaken this simple task, drawing samples from several different classes.

Before presenting the first experiment, a bit of background on the materials is in order. Over the last two decades, psychologists have amassed an impressive body of data on the capacity to distinguish moral violations from conventional violations. People, from preschool through adulthood, distinguish moral violations (e.g. hitting others) from conventional violations (e.g. drinking soup out of a bowl) on a number of dimensions. Moral violations are treated as more serious and less permissible than conventional violations. Moral violations are also regarded as less contingent on authority than conventional violations. For example, children tend to say that if the teacher said it was okay to leave class without permission, that action would be okay, but even if the teacher said it was okay to hit, that action would not be okay. In addition, moral violations also elicit different kinds of justifications from conventional violations. Whereas the explanations for why moral violations are wrong typically appeal to welfare considerations (e.g. "because it's not fair" or "because it hurt the person"), the explanations for why conventional violations are wrong appeal

instead to social-conventional considerations (e.g. "you're not supposed to do that at social functions") (for reviews, see Nucci, 2001; Smetana, 1993; Tisak, 1995).

In the next section, I report a study in which canonical moral and conventional cases drawn from the literature are used in assessing judgments about moral objectivity. In addition, a case of a disgusting violation was included since recent work indicates that disgusting violations (e.g. spitting in one's water glass before drinking from it), like moral violations, are treated as less permissible, more serious and less authority contingent than conventional violations (Nichols, 2002).

3. Experiment 1

This experiment was designed to explore participants' views about the objectivity of morals, conventions, and ordinary facts. One common way that views about moral objectivity are elicited in introductory philosophy is by considering cases of cultural disagreement (e.g. Rachels, 2002). This experiment exploited that device to assess participants' views about objectivity in different domains.

3.1. Method

3.1.1. Participants. 46 undergraduates at the College of Charleston participated in this study. 24 of the participants were in an introductory philosophy class; 22 were in an introductory mathematics class. 21 participants were female and 25 were male.

3.1.2. Procedure. Participants were given questionnaires in a classroom. The questionnaires described 4 scenarios of cultural disagreement: a moral scenario, a conventional scenario, a disgust scenario, and a flat earth scenario. For each case, participants were asked to indicate whether they agreed with one of the disputants or whether they thought that there was no fact of the matter in the dispute. For instance, in the moral disagreement scenario, participants were presented with the following:

> John and Fred are members of different cultures, and they are in an argument. John says, "It's okay to hit people just because you feel like it," and Fred says, "No, it is not okay to hit people just because you feel like it." John then says, "Look you are wrong. Everyone I know agrees that it's okay to do that." Fred responds, "Oh no, you are the one who is mistaken. Everyone I know agrees that it's not okay to do that."
>
> Which of the following do you think best characterizes their views? (Check one and give a brief justification for your answer.)
>
> —It is okay to hit people just because you feel like it, so John is right and Fred is wrong.
> —It is not okay to hit people just because you feel like it, so Fred is right and John is wrong.
> —There is no fact of the matter about unqualified claims like "It's okay to

> hit people just because you feel like it." Different cultures believe different things, and it is not absolutely true or false that it's okay to hit people just because you feel like it.

In the flat-earth case, participants were presented with the following (the other items are in the appendix):

> Teresa and Heather are members of different cultures, and they are in an argument. Teresa says, "The earth is flat," and Heather says, "No, the earth is not flat." Teresa then says, "Look you are wrong. Everyone I know agrees that it is flat." Heather responds, "Oh no, you are the one who is mistaken. Everyone I know agrees that it is not flat."
>
> Which of the following do you think best characterizes their views? (Check one and give a brief justification for your answer.)
>
> —The earth is flat, so Teresa is right and Heather is wrong.
> —The earth is not flat, so Heather is right and Teresa is wrong.
> —There is no fact of the matter about unqualified claims like "The earth is flat." Different cultures believe different things, and it is not absolutely true or false that the earth is flat.

3.2. Results

3.2.1. Scoring. A response is coded as nonobjectivist if it indicates agreement with the statement that begins: "There is no fact of the matter …" If the response indicates agreement with either of the other options, it is coded as objectivist.

3.2.2. Analysis. 6 of the 46 participants gave nonobjectivist responses to the flat earth case. These participants were excluded from subsequent analyses [2].

Not surprisingly, participants were much more likely to give nonobjectivist responses to the conventional case than to the moral case (McNemar's test, $N = 40$, $p < 0.001$) (all tests two-tailed). Participants were also much more likely to give nonobjectivist responses to the disgust case than to the moral case (McNemar's test, $N = 40$, $p < 0.001$). Participants were somewhat more likely to give nonobjectivist responses to the conventional case than to the disgusting case (McNemar's test, $N = 40$, $p < 0.05$). Participants tended to respond as nonobjectivists to both the conventional case (χ^2 goodness-of-fit ($N = 40$, $df = 1$) $= 36.1$, $p < 0.001$) and the disgust case (χ^2 goodness-of-fit ($N = 40$, $df = 1$) $= 14.4$, $p < 0.001$), but their responses to the moral case did not differ significantly from chance (χ^2 goodness-of-fit ($N = 40$, $df = 1$) $= 0.9$, $p = 0.342$, n.s.). What is of particular importance for our purposes is that even with the flat-earth nonobjectivists excluded, a sizable number of participants gave nonobjectivist responses to the moral case (17 of the 40 participants) [3].

4. The claim to objectivity refigured

The import of the foregoing evidence for our purposes is simply that a nontrivial population of undergraduates endorse a nonobjectivist claim about a moral violation without endorsing a full blown metaphysical nonobjectivism. For many philosophers, this will not be the least bit surprising. It merely serves as some corroboration for the observations made by Harman, Stich and Weinberg. In keeping with their remarks, the point is just that the prevailing view in analytic philosophy that commonsense is committed to objectivity is at least too crude. However, I suspect that there is an important nugget of truth in the standard claim that commonsense is committed to moral objectivity. And there is a bit of evidence from developmental psychology that suggests a way to shore up the standard view.

In Section 2, I noted that psychologists have found that children distinguish moral from conventional violations on a number of dimensions. One dimension that I neglected to mention plays on the "generalizability" of moral judgment. Researchers have found that children distinguish moral from conventional violations along this dimension as well. For instance, children tend to maintain that moral violations are wrong in other countries, but they are less likely to maintain that conventional violations are wrong in other countries. That is, young children tend to regard moral violations as generalizably wrong (e.g. Smetana, 1993). Recent evidence indicates that young children also recognize that certain properties, like *yummy*, are response dependent, but children tend not to treat moral properties as response dependent in this way (Nichols & Folds-Bennett, forthcoming). In both the generalizability and response-dependent studies, one natural interpretation is that children are displaying their commitment to moral objectivism—they regard immoral actions as wrong *simpliciter*, not merely wrong relative to certain individuals. The data do not, it's worth stressing, constitute decisive evidence that children are moral objectivists. Many possible interpretations of the data that do not impute objectivism to children have yet to be explored. But all of the evidence we have thus far on judgments of generalizability and response-dependence fits well with the claim that children are moral objectivists, and I will accordingly adopt the working hypothesis that moral objectivism is a default setting on folk metaethics. In any case, even if moral objectivism is a default position, the preceding experiment (and the ones to follow) suggests that this default position is *defeasible*. For there is a significant population of individuals who indicate that canonical moral violations are not objectively wrong. Nonetheless, the developmental evidence suggests that these individuals have moved away from the default view.

5. After moral objectivism: error theory and inessentialism

Moral objectivity, then, is plausibly a default setting on commonsense metaethics. As a result, the rejection of moral objectivity exacts a revision of commonsense metaethics. There is, of course, a thriving tradition in philosophy that argues against moral objectivism. I will not rehearse any of the arguments against objectivity (e.g. Harman, 1985; Hume, 1739/1964; Mackie, 1977; Nichols, forthcoming), but in the

remainder of this paper, I want to explore the impact that rejecting objectivity would have on lay moral judgment.

The wholesale rejection of moral objectivism leaves open radically different accounts of the status of commonsense moral judgment. Perhaps the most notorious proposal in contemporary analytic philosophy is Mackie's error theory. According to error theory, ordinary moral judgments are so infused with the commitment to objectivity that the truth of nonobjectivism renders all such judgments false. Mackie maintains that the presupposition that moral values are objective "is part of what our ordinary moral statements mean: the traditional moral concepts of the ordinary man … are concepts of objective value" (Mackie, 1977, p. 35). He goes on to say:

> I do not think it is going too far to say that this assumption has been incorporated in the basic, conventional meanings of moral terms … The claim to objectivity, however ingrained in our language and thought, is not self-validating. It can and should be questioned. But the denial of objective values will have to be put forward not as the result of an analytic approach, but as an "error theory", a theory that although most people in making moral judgments implicitly claim, among other things, to be pointing to something objectively prescriptive, these claims are all false. (Mackie 1977, p. 35)

Hence, on Mackie's view, it would seem that once one comes to reject moral objectivism, one should regard common moral judgments, even one's own earlier moral judgments, as uniformly false. For all of those judgments are infected by the objectivist presupposition. The "traditional moral concepts" employed in lay moral judgments are, like the concept *witch*, empty.

One might reject objectivism and try to resist error theory by maintaining that the belief in moral objectivity is simply not essential to lay moral concepts. On this "inessentialist" view, lay moral concepts might not be empty even if morality is not objective. Gilbert Harman defends moral relativism in a way that suggests an inessentialist account. According to Harman, when people make moral claims, the claims are best regarded as statements that have tacit restrictive quantification. As a result, even though objectivity is rejected, common moral judgments can often be true, since they are best read as implicitly relativized to a group. Harman intimates that it would be "mean-spirited" to suggest that common moral judgments are all false (Harman, 1996, p. 4). Although lay people might not intend their moral **Q1** proclamations to be elliptical for relativized statements, it is best to interpret them thus. So, Harman writes, "a judgment of the form, *P ought morally to D*, has to be understood as elliptical for a judgment of the form, *in relation to moral framework M, P ought morally to D*" (Harman, 1996, p. 4 fn 2). Now, then, when the Hopi say "it **Q1** was okay for Fred to hurt that bird" and we say "it was not okay for Fred to hurt that bird," we might both be saying something true. The Hopi's claim is indexed to one moral framework and our claim is indexed to a different moral framework. Again, Harman takes pains to explain that he is not saying that this is what people intend (e.g. Harman, 1996, p. 17). But after one has relinquished objectivism, this **Q1** may be the best way to interpret them.

Thus, according to error theorists, common moral judgments are all false. According to inessentialists, common moral judgments need not all be false; rather, commonsense moral judgments might be true judgments that are relativized to some moral framework. How do we decide who is right in this debate? Are lay moral concepts so infected with the metaethical claim of objectivity that the rejection of objectivity is tantamount to denying that lay moral concepts apply to anything? The question can take on a more concrete shape in light of the empirical results discussed thus far. In Section 4, I maintained that moral objectivism is a default setting on commonsense metaethics; in earlier sections, we saw that many college students apparently reject objectivism. If, as seems possible, some of these college students fully reject moral objectivism, it generates the longitudinal result that, at some point in their development, a number of individuals convert from objectivism to nonobjectivism. Such individuals still use terms like "morally wrong," and let's charitably assume that their post-objectivist uses of "morally wrong" are sometimes true. For such individuals who reject moral objectivity but still claim that certain actions are "morally wrong" do we want to say (1) or (2)?

1. Their term "morally wrong" expresses a different concept after the rejection of objectivity.
2. Their term "morally wrong" expresses the same concept that it did before the rejection of objectivity.

The error theorist will maintain (1). According to the error theorist, the traditional moral concepts are empty, so if post-objectivist uses of "morally wrong" accurately apply to certain actions (under certain conditions), the term "morally wrong" must express a concept different from the old concept. The inessentialist, on the other hand, can embrace (2). For according to the inessentialist, even if objectivity is false, the moral concepts of moral objectivists are not empty. As a result, according to the inessentialist, objectivists and nonobjectivists might well deploy the same moral concepts when they assert such things as "it's morally wrong to hit people."

To resolve this dispute between error theorists and inessentialists, we want to know the extent to which abandoning objectivity affects one's moral cognition. But there is a prior and more fundamental issue. To determine whether an individual deploys the same moral concepts before and after the rejection of objectivism, we need to know the answer to a longstanding problem in philosophy of mind: how are concepts individuated? In particular, to know whether "morally wrong" expresses the same concept before and after the rejection of objectivity, one needs to know the basic criteria for what makes one concept the same concept across changes in the cognitive economy. To see why, we need to pause to review recent debates on concepts and concept individuation (for a review, see Laurence & Margolis, 1999) [4]. Most theorists agree that two concept tokens are of the same type only if they have the same content; however, there is radical division about the conditions under which two concepts have the same content. Some theorists maintain that the content of a concept is partly determined by the concept's "conceptual role," that is, its pattern of causal interactions with other mental states (e.g. Block, 1986; Murphy & Medin, 1985). On these accounts, two concepts have the same content only if they

have the same (or roughly the same) conceptual roles. As a result, if two concept tokens have even minor differences in their conceptual roles, then they are, on some theories, not tokens of the same type of concept. Hence, on such an account of concept individuation, when the objectivist and nonobjectivist each says "hitting is morally wrong," they would almost certainly be expressing different concepts, since there will at least be minor differences in the conceptual roles of their respective "moral" concepts. At the other end of the spectrum, Jerry Fodor's widely discussed account of concepts suggests exactly the opposite (e.g. Fodor, 1998). According to Fodor, the content of a concept is not determined by conceptual role, but rather by a certain nomological relation between the concept and a property. If this theory is right, and if content identity determines concept identity, then two concept tokens with radically different conceptual roles can still be instances of the same concept. So two people might be deploying the same concept even though they hold drastically different beliefs that involve the concept. As a result, nonobjectivists and objectivists might be exploiting tokens of the same moral concepts even if those tokens have radically different conceptual roles.

On some accounts of concept individuation, then, even if nonobjectivists have radically different moral cognition than objectivists, it might still be the case that they are all using the same moral concepts. On other accounts of concept individuation, even if nonobjectivists have only minimally different moral cognition from objectivists, it might still be the case that uses of "morally wrong" express concepts of different types in objectivists and nonobjectivists. Roughly speaking, conceptual role theories tend to be biased towards error theory, and Fodorian theories tend to be biased against error theory. These are just a couple of locations in the sophisticated and highly contentious theoretical landscape on concept individuation. However, they serve to indicate that, without settling issues about concept individuation, we will be unable to determine whether error theory is right.

This is hardly the place to try to arrive at a considered view of concept individuation. Yet, one needs a theory of concept individuation to assess whether lay moral concepts get preserved as the same concepts after a nonobjectivist epiphany. As a result, it seems premature to pronounce on whether nonobjectivism entails error theory or inessentialism. While we await the resolution of how to individuate concepts (it may be awhile), we can, however, consider the subsequent issue about the extent to which rejecting objectivity impacts moral cognition. In fact, we can ask a nice experimental question at this juncture. We have seen that a significant population of college students respond as nonobjectivists about canonical moral violations. And we know that moral judgments carry a typical suite of responses, probed by the moral/conventional task (Section 2). This suite of responses plausibly reflects a capacity for a kind of "core moral judgment" (Nichols, forthcoming). Hence we might investigate whether students who respond as moral nonobjectivists continue to distinguish moral violations from conventional violations on the classic dimensions.

Four experiments are presented exploring this issue. Note that there is a Humean prediction flowing from the epigraph to this paper about how these experiments will turn out. The Humean prediction is that giving up objectivity will

have little effect on the characteristics of moral judgment, so we should expect nonobjectivists to respond much like objectivists. More recent affect-based accounts of moral judgment (e.g. Blair, 1995; Nichols, 2002) would similarly predict that responses to the moral/conventional task will likely be preserved even in (clinically normal) adults who reject moral objectivity.

6. Experiment 2

In this experiment, participants were given a standard moral/conventional task, followed by a set of objectivity questions drawn from experiment 1. The objectivity questions included the moral case (whether it's okay to hit others just because you feel like it), the conventional case, and the factual case (whether the earth is flat) (Section 3).

6.1. Method

6.1.1. Participants. 13 undergraduates in an introductory philosophy class at the College of Charleston participated in this study. 9 of the participants were female, 4 were male.

6.1.2. Procedure. Participants were given questionnaires in a classroom. The questionnaires first presented a standard moral/conventional task. There were four scenarios: two moral cases (hitting; pulling hair) and two conventional cases (talking in class; chewing gum). After each scenario was described, four questions were asked to explore judgments about permissibility, seriousness, justification, and authority contingency. The scenarios and questions were drawn from the literature (e.g. Blair, 1995; Nichols, 2002). For instance, one of the moral violation items went as follows:

> Ben and Frank are boys in grade school. While Ben is sitting at his desk working, Frank walks up and hits Ben and Ben begins to cry.
>
> 1. Was it O.K. for Frank to hit Ben?
> 2. On a scale of one to ten (with ten highest), how bad was it for Frank to hit Ben?
> 3. Why was it bad for Frank to hit Ben?
> 4. Now what if the teacher had said before the lesson, before Frank hit Ben, that "At this school, anybody can hit someone if they want to." Would it O.K. for Frank to hit Ben if the teacher says he can?

The order of the moral and conventional items was counterbalanced. Following the moral/conventional task, participants were given the moral, factual, and conventional objectivity questions from experiment 1.

6.2. Results

6.2.1. Scoring. On the moral/conventional task, the scoring procedure follows that

TABLE 1. Description of justification categories

Other's welfare (harm): Any reference to the welfare of the victim (e.g. "it will hurt him")
Violence: Any reference to the wrongness of violence (e.g. "violence is bad")
Normative reference (rule): Any reference, even implicit (e.g. "it's not acceptable to do that"), to rules
Disorder statements (disorder): Any reference to the disruption caused by the transgression (e.g. "It will distract the class")
Rudeness (manners): Any reference to the rudeness of the transgression (e.g. "it's bad manners")
Other: Any other response

of the previous studies in the literature (Smetana & Braeges, 1990; Blair, 1995). Questions 1 and 4 were scored binomially, with each No answer being given a score of 1, so the cumulative score for each domain could range from 0 to 2. Question 2 was coded by the value (between 1 and 10) given to the seriousness of the transgression. Question 3 was scored by an independent coder according to the justification categories in Table 1. Responses coded as harm or violence were treated as welfare-based, all other responses were treated as non-welfare based. Scoring for the objectivity questions was the same as in experiment 1.

6.2.2. Analysis. As in experiment 1, several students (3) said that there was no fact of the matter about whether the earth is flat. Those students were excluded from further analysis. There remained 10 participants, 7 of whom maintained that there is a fact of the matter about whether the earth is flat, but there is no fact of the matter about the moral case. This group is labeled "moral nonobjectivists."

Mean responses were calculated for each type of question for each domain (see Table 2). Four 2 (group) $\times$ 2 (domain) split-plot ANOVAs were performed on the responses for all 4 question types. There was a main effect for domain on permissibility ($F(1,8) = 11.2$, $p < 0.05$), seriousness ($F(1,8) = 72.019$, $p < 0.001$), justification ($F(1,8) = 9.618$, $p < 0.05$), and authority contingence ($F(1,8) = 20.084$, $p < 0.01$). There were no significant interactions, and there were no significant group effects.

More importantly, single-factor repeated measures ANOVAs revealed that nonobjectivists drew the moral/conventional distinction on all questions. Four

TABLE 2. Means and standard deviations for experiment 2

	Permissibility		Seriousness		Welfare justification		Authority contingence	
	M	C	M	C	M	C	M	C
Objectivists								
M	2.00	1.00	18.667	4.33	1.00	0.00	2.00	0.67
SD	0.00	1.00	1.15	2.08	1.00	0.00	0.00	1.15
Nonobjectivists								
M	2.00	1.00	12.71	4.14	1.14	0.00	1.71	0.29
SD	0.00	0.82	4.75	1.86	0.90	0.00	0.76	0.49

single-factor repeated measures ANOVAs were performed with domain as the repeated factor. There was a main effect for permissibility ($F(1,6) = 10.5$, $p < 0.05$), seriousness ($F(1,6) = 30.337$, $p < 0.01$), justification ($F(1,6) = 11.294$, $p < 0.05$), and authority contingence ($F(1,6) = 23.077$, $p < 0.01$).

The fact that the moral/conventional distinction is found with only 7 participants indicates that even among nonobjectivists, there is a powerful tendency to distinguish moral from conventional violations. However, given that there are so few participants, it's reasonable to worry about the replicability of this finding. Hence, a variant of the task was presented to a new set of participants in the next experiment.

7. Experiment 3

7.1. Method

7.1.1. Participants. 54 undergraduates in an introductory philosophy class at the College of Charleston participated in this study. 28 of the participants were female, 26 were male. 1 additional participant failed to complete half of the form and was excluded.

7.1.2. Procedure. The procedure was the same as in experiment 2, with a few small changes. First, the order of the moral/conventional task and the objectivity questionnaire was counterbalanced. Approximately half the participants received the moral/conventional task first and the other half received the objectivity questions first. Second, in experiment 2, one of the conventional items (chewing gum in class) was regarded as fully permissible by many participants, so it was replaced with an alternate conventional item (leaving class without permission). Third, since it's possible that whether a person endorses moral nonobjectivism is partly a function of education, in this experiment information about year in school was also requested. Finally, at the end of the objectivity questionnaire, another question about moral objectivity was added. Intuitions about objectivity are often probed by considering thought experiments involving aliens. Participants were given such an alien case focusing on the permissibility of puppy torture:

> Most people think that it's wrong for anyone to torture puppies for the fun of it. Imagine that there are intelligent aliens who know all about the earth and all about people and animals on the planet. The aliens think that it's okay for anyone to torture puppies for the fun of it. In this case, which of the following do you think best characterizes the views of humans and the aliens? (Check one and give a brief justification for your answer.)
>
> —It is okay to torture puppies, so the aliens are right and the humans are wrong.
> —It is not okay to torture puppies, so the humans are right and the aliens are wrong.
> —There is no fact of the matter about unqualified claims like "It's okay to

TABLE 3. Means and standard deviations for experiment 3

	Permissibility		Seriousness		Welfare justification		Authority contingence	
	M	C	M	C	M	C	M	C
Objectivists								
M	2.00	1.85	14.44	9.36	0.93	0.00	1.93	0.39
SD	0.00	0.37	4.03	5.32	0.83	0.00	0.38	0.69
Nonobjectivists								
M	2.00	1.89	13.30	8.70	1.30	0.11	1.44	0.40
SD	0.00	0.33	4.83	3.71	0.82	0.33	0.88	0.52

torture puppies." Humans and aliens believe different things, and it is not absolutely true or false that it's okay to torture puppies.

7.2. Results

7.2.1. Scoring. Same as in experiment 2.

7.2.2. Analysis. As in experiments 1 and 2, several students (12) said that there is no fact of the matter about whether the earth is flat, and 2 more participants failed to answer the question. All of these participants were excluded from further analysis, leaving 39 remaining participants.

10 of the remaining 39 participants gave nonobjectivist responses to the moral case of hitting others. Interestingly, there was a significant correlation between year in college and nonobjectivist response to the hitting-others case ($r = -.366$, $N = 39$, $p < 0.05$). Participants with more years in college were more likely to respond as nonobjectivists on this question.

Mean responses were calculated for each type of question for each domain (see Table 3). Four 2 (group) $\times$ 2 (domain) split-plot ANOVAs were performed on the responses for all 4 question types. There was a marginally significant main effect for domain on the permissibility question ($F(1,32) = 3.663$, $p = 0.065$). There were clear main effects for seriousness ($F(1,35) = 35.632$, $p < 0.001$), justification ($F(1,34) = 48.478$, $p < 0.001$), and authority contingence ($F(1,34) = 80.295$, $p < 0.001$). There were no significant interactions or group effects.

When we turn to focus on the moral nonobjectivists, the results of experiment 2 were largely replicated. Four single-factor repeated measures ANOVAs were performed with domain as the repeated factor. On permissibility, there was no significant difference ($F(1,8) = 1$, $p = 0.347$, n.s.). This is not surprising, since a less permissible item (leaving class without permission) was explicitly sought and added. However, there were clear main effects for seriousness ($F(1,9) = 30.519$, $p < 0.001$), justification ($F(1,8) = 21.333$, $p < 0.01$), and authority contingence ($F(1,8) = 18$, $p < 0.01$).

As noted in the Procedure section, an additional moral objectivity question, using alien attitudes about puppy torture, was added to this experiment. A some-

what larger number of participants responded as nonobjectivists to this question (N = 15). Using this category to divide groups, an additional four 2 (group) × 2 (domain) split-plot ANOVAs were performed on the responses for all 4 question types to check for group effects and interactions. Again there were no group effects or interactions. Focusing on the group of participants responding as nonobjectivists about puppy torture, four single-factor repeated measures ANOVAs were performed with domain as the repeated factor. In this group as well, the moral/conventional distinction is clearly made. Again, there was not a distinction on the permissibility question ($F(1,12) = 2.182$, $p = 0.165$, n.s.). However, there were main effects for seriousness ($F(1,12) = 36.958$, $p < 0.001$), justification ($F(1,11) = 12.439$, $p < 0.01$), and authority contingence ($F(1,14) = 58.414$, $p < 0.001$).

Thus, this experiment replicates the findings of experiment 2. Participants who give nonobjectivist responses to moral cases still draw the moral/ conventional distinction. Of course, the number of nonobjectivist respondents in this experiment, as in experiment 2, is small. However, the fact that the finding is replicated provides fairly compelling evidence that the phenomenon is real. Furthermore, the fact that the effect is evident with so few participants indicates that the tendency to distinguish moral from conventional transgressions is quite pronounced even in people who respond as moral nonobjectivists. However, the questions used to assess whether the participant is a moral objectivist do not quite map on to the moral violations in the moral/conventional task. This was controlled for in the next experiment.

8. Experiment 4

8.1. Method

8.1.1. Participants. 53 undergraduates in a developmental psychology class at the College of Charleston participated in this study. 50 of the participants were female, 3 were male.

8.1.2. Procedure. The procedure was largely the same as in experiment 2, using the alternate conventional item from experiment 3. The key difference was that the same kinds of moral violations (hitting and shoving) were used in both the moral/conventional task and in the moral objectivity question. Thus, the moral objectivity question was as follows:

> John and Fred are members of different cultures, and they are in an argument. John says, "It's okay to hit or shove people just because you feel like it," and Fred says, "No, it is not okay to hit or shove people just because you feel like it." John then says, "Look you are wrong. Everyone I know agrees that it's okay to do those things." Fred responds, "Oh no, you are the one who is mistaken. Everyone I know agrees that it's not okay to do those things."

TABLE 4. Means and standard deviations for experiment 4

	Permissibility		Seriousness		Welfare justification		Authority contingence	
	M	C	M	C	M	C	M	C
Objectivists								
M	2.00	1.96	17.28	9.90	0.53	0.00	1.97	0.52
SD	0.00	1.92	2.63	3.60	0.76	0.00	0.19	0.63
Nonobjectivists								
M	2.00	1.88	13.90	7.67	0.67	0.00	2.00	0.38
SD	0.00	0.35	4.60	2.40	0.87	0.00	0.00	0.52

Which of the following do you think best characterizes their views? (Check one and give a brief justification for your answer.)

—It is okay to hit or shove people just because you feel like it, so John is right and Fred is wrong.
—It is not okay to hit or shove people just because you feel like it, so Fred is right and John is wrong.
—There is no fact of the matter about unqualified claims like "It's okay to hit or shove people just because you feel like it." Different cultures believe different things, and it is not absolutely true or false that it's okay to hit or shove people just because you feel like it.

8.2. Results

8.2.1. Scoring. Same as in experiment 2.

8.2.2. Analysis. As in the other experiments, several students said that there is no fact of the matter about whether the earth is flat. 12 participants answered this way in the present experiment. All of these participants were excluded from further analysis, leaving 41 remaining participants. Nine of the remaining 41 participants responded as moral nonobjectivists. In this experiment, there was no correlation between year in school and moral nonobjectivism ($r = 0.098$, $N = 39$, $p = 0.554$, n.s.).

Mean responses were calculated for each type of question for each domain (see Table 4). Four 2 (group) $\times$ 2 (domain) split-plot ANOVAs were performed on the responses for all 4 question types. There was no main effect for domain on permissibility ($F(1,33) = 2.909$, $p = 0.097$, n.s.). However, there were main effects for seriousness ($F(1,39) = 117.104$, $p < 0.001$), justification ($F(1,39) = 16.401$, $p < 0.001$), and authority contingence ($F(1,34) = 154.981$, $p < 0.001$). There were no significant interactions, and there were no significant group effects, except on the seriousness questions ($F(1,39) = 6.987$, $p < 0.05$).

As in the previous experiments, single-factor repeated measures ANOVAs revealed that nonobjectivists drew the moral/conventional distinction on seriousness, justification, and authority contingence. Four single-factor repeated measures

ANOVAs were performed with domain as the repeated factor. There was no main effect for permissibility ($F(1,7) = 1$, $p = 0.357$, n.s.). However, there were main effects for seriousness ($F(1,8) = 25.210$, $p < 0.01$), justification ($F(1,8) = 5.333$, $p = 0.05$), and authority contingence ($F(1,7) = 78.867$, $p < 0.001$). Thus the current experiment shows that people who claim that hitting and shoving are not objectively wrong still treat instances of hitting and shoving in much the same way as people who claim that hitting and shoving are objectively wrong.

9. Experiment 5

9.1. Method

9.1.1. Participants. 22 undergraduates in an introductory philosophy class at the College of Charleston participated in this study. 15 of the participants were female, 7 were male. One additional participant was excluded for failing to answer the objectivity questions. The experiment was conducted at the end of the semester, by which time all participants had some exposure to introductory metaethics.

9.1.2. Procedure. The procedure was largely the same as in experiment 4, except that the objectivity questions were focused on the exact same transgressions described in the moral/conventional task. Thus, the moral objectivity question was as follows:

> John and Ted are members of different cultures, and they are in an argument about some of the scenarios described above. John says, "It was okay for Frank to hit Ben and for Lisa to shove Nancy," and Ted says, "No, those things were not okay." John then says, "Look you are wrong. Everyone I know agrees that those kinds of things are okay." Ted responds, "Oh no, you are the one who is mistaken. Everyone I know agrees that those kinds of actions are not okay."
>
> Which of the following do you think best characterizes their views? (Check one and give a brief justification for your answer.)
>
> —It is an objective fact, independent of what different people think, that it was not wrong for Frank to hit Ben and for Lisa to shove Nancy. So John is right and Ted is wrong.
> —It is an objective fact, independent of what different people think, that it was wrong for Frank to hit Ben or for Lisa to shove Nancy. So Ted is right and John is wrong.
> —There is no objective fact, independent of what different people think, about whether it was wrong for Frank to hit Bill or Lisa to shove Nancy. These actions were "wrong for Ted" and maybe "wrong for me," but they aren't *objectively wrong* independent of what people think about them.

The flat earth question was altered to parallel this as well.

TABLE 5. Means and standard deviations for experiment 5

	Permissibility		Seriousness		Welfare justification		Authority contingence	
	M	C	M	C	M	C	M	C
Objectivists								
M	2.00	1.83	16.00	12.00	1.50	0.00	2.00	0.33
SD	0.00	0.41	2.61	3.41	0.84	0.00	0.00	0.52
Nonobjectivists								
M	2.00	1.83	15.17	7.67	1.33	0.00	2.00	0.17
SD	0.00	0.58	4.02	3.34	0.65	0.00	0.00	0.39

9.2. Results

9.2.1. Scoring. Same as in experiment 2.

9.2.2. Analysis. As in the other experiments, several students (4) said that there is no fact of the matter about whether the earth is flat. All of these participants were excluded from further analysis, leaving 18 remaining participants. 12 of those 18 responded as moral nonobjectivists.

Mean responses were calculated for each type of question for each domain (see Table 5). Four 2 (group) $\times$ 2 (domain) split-plot ANOVAs were performed on the responses for all 4 question types. There was no main effect for domain on permissibility ($F(1,16) = 1.580$, $p = 0.227$, n.s.). However, there were main effects for seriousness ($F(1,16) = 58.372$, $p < 0.001$), justification ($F(1,16) = 62.912$, $p < 0.001$), and authority contingence ($F(1,16) = 261.333$, $p < 0.001$). There were no interactions except on the serious dimension ($F(1,16) = 5.407$, $p < 0.05$); but interestingly, this can be explained by the fact that objectivists gave higher scores than nonobjectivists for the *conventional* violations ($t(16) = 2.579$, $p < 0.05$). There were no significant group effects.

As in the previous experiments, single-factor repeated measures ANOVAs revealed that nonobjectivists drew the moral/conventional distinction on seriousness, justification, and authority contingence. Four single-factor repeated measures ANOVAs were performed with domain as the repeated factor. There was no main effect for permissibility ($F(1,11) = 1$, $p = 0.339$, n.s.). However, there were main effects for seriousness ($F(1,11) = 66.892$, $p < 0.001$), justification ($F(1,11) = 50.286$, $p < 0.001$), and authority contingence ($F(1,11) = 266.200$, $p < 0.001$).

10. Discussion

The results of experiments 2–5 bear out the Humean prediction that the rejection of objectivity will leave moral cognition largely intact. On the standard version of the moral/conventional task used here, students responding as moral nonobjectivists about standard moral violations drew the moral/conventional distinction just like

objectivists, and indeed, there is little to distinguish objectivists from nonobjectivists in any of these experiments.

The results should, of course, be taken as preliminary. In particular, the measures of whether a student is a moral nonobjectivist are far from perfect. I should hope that future research will develop more sensitive ways to determine participants' commitment to moral nonobjectivism. I would, however, like to discuss two of the more important objections that can be raised against the present experiments. First, I'll consider several versions of the objection that the "moral nonobjectivists" in these experiments were not really nonobjectivists [5]. Second, I'll address the subsequent objection that if they really are nonobjectivists, then their performance on the moral/conventional task shows them to be deeply confused.

Many analytic philosophers are generally skeptical of the claim that a significant number of undergraduates are nonobjectivists. This general skepticism might be directed against the experiments here. For instance, several colleagues in philosophy have suggested that we should not take students at their nonobjectivist word because it's easy to talk students out of their nonobjectivist rhetoric. It is no doubt true that professors often talk their students out of their nonobjectivist rhetoric, but this is a seriously tainted set of anecdotes. An effective professor can talk students out of many of their genuine commitments. Indeed, it's likely that an effective professor could talk many objectivist students (at least among those who are not theists) out of their commitment to moral objectivism. As a result, to investigate these matters, we need to take precautions to minimize the sort of biases that are endemic to the professor/student relationship. Presumably anonymous experimental procedures are a step in the right direction.

A different approach to rejecting the claim that some undergraduates are nonobjectivists is to maintain that students' real views come out if they are asked about extreme moral violations, like racist murder. It is likely that far fewer students would maintain nonobjectivism about such cases. Thus, the objection goes, this shows that "deep down" they are really thoroughgoing moral objectivists. While perhaps most of the nonobjectivists in the above studies would be objectivists about racist murder, this does not immediately show that they are really thoroughgoing moral objectivists. Rather, an interpretation that would cling more tightly to the students' predicted responses is that they think that moral violations vary with respect to objectivity—some moral violations are objectively wrong and some are not. Of course, this suggestion triggers the familiar philosophical reply that if the person thinks some of these violations are not objectively wrong, then he really does not think those violations are *morally* wrong. However, this move sucks us back into the quagmire over concept individuation. To assess this view, one would need to determine whether objectivity is essential to the concept of *moral wrong*. And to do that, we would need a theory of concept individuation. Unfortunately, as noted earlier, we manifestly do not have a settled story about concept individuation. We have put aside the issue of concept individuation to see whether the rejection of moral objectivity radically affects moral cognition. To explore this, it suffices to consider participants who are nonobjectivists about canonical moral violations like hitting. For the developmental evidence suggests that basically all children are

objectivists about canonical moral violations like hitting. So college students who are nonobjectivists about such violations have presumably abandoned their earlier objectivist views about these transgressions.

All of the above is by way of defending the *possibility* that the "nonobjectivists" really are nonobjectivists about certain moral violations. But there is a bit that might be said more positively. In the experiments reported here, participants were invited to explain briefly why they answered as they did to the objectivity questions. Many nonobjectivists supplied explanations reflecting the view that what is morally right and wrong depends on the culture, but scientific facts do not. For instance one nonobjectivist participant, who responded that there is no fact of the matter about the moral case, wrote, "Morality is subjective to culture." In explaining his objectivist response to the flat earth case, he wrote, "This is a cold, unwavering scientific fact." More generally, it's plausible that many of these nonobjectivists think that whether some kind of act is morally wrong depends on the culture, but when it comes to the shape of the earth, as one participant put it, "In this case culture does not matter." These kinds of claims, while not exactly sophisticated, do conform to nonobjectivist views about morality. It's worth emphasizing here that the pattern of responses of nonobjectivists in these experiments was precisely that they treat moral disagreements quite differently from scientific disagreements, and the idea that moral disagreements have a radically different status than scientific disagreements is part of the considered views of many philosophical nonobjectivists (e.g. Harman, 1977; Mackie, 1977). The recognition of moral disagreement also forms the basis for one of the most influential arguments against moral realism, the "argument from disagreement" (e.g. Doris & Stich, forthcoming; Loeb, 1998; Mackie, 1977). We could hardly expect students who proclaim nonobjectivism to have compelling arguments for the view (anymore than we can expect the objectivists to have compelling arguments for objectivism), but neither does it seem appropriate to dismiss out of hand their stated commitments.

Yet another complaint against treating the "nonobjectivists" as nonobjectivists is that students who give nonobjectivist responses might really be objectivists who regard the relevant action-descriptions as too coarse grained. The students might have the following interpretation: actions that meet this description can have crucially different properties in different contexts and these different properties can make the actions objectively wrong in some instances and objectively permissible in others. This interpretation of the responses is not excluded by the measures used in experiments 2–4. But neither does this seem like a natural interpretation of all nonobjectivist responses. For example, it seems unlikely that subjects think that puppy torture is objectively permissible in some cases (perhaps for *very* naughty puppies?). More importantly, experiment 5 was designed to address this issue directly. In that experiment, several participants maintained that the particular actions described in the moral scenarios (Fred hitting Bill and Lisa shoving Nancy) were not objectively wrong, independent of what people think. Yet there was still a clear moral/conventional distinction in this group.

A final complaint about calling participants "nonobjectivists" is that it is too coarse-grained in a different way. The participants are probably not *global* nonobjec-

tivists about morality. Rather, subjects might be described as nonobjectivists about some kinds of moral transgressions but not about others. Obviously, this is in keeping with my remarks two paragraphs back. But one might raise this as a worry about the experiments. If it is allowed that the objectivity probes do not demonstrate that participants are global nonobjectivists, then it's too crude to say that the experiments show that nonobjectivists draw the moral/conventional distinction. Rather, in experiments 2 and 3, participants who were nonobjectivists about action-type *X* (hitting others because you feel like it) still treated instances of action-types *Y* (pulling hair) and *Z* (hitting) as wrong in a way that differed from conventional wrongs. So it's possible that the participants are objectivists about, for instance, the wrongness of pulling hair. This is a fair point, but one obvious parry is to note that hitting others because you feel like it is much like the hitting violation used in the moral/conventional task. More importantly, experiments 4 and 5 were designed specifically to control for this worry. For in those experiments, participants who responded as nonobjectivists about shoving and hitting still treated instances of shoving and hitting as distinctively wrong. Apparently people can be nonobjectivists about certain types of transgressions while still treating such transgressions otherwise very much like moral violations, and very much in the way that objectivists treat them. This is the key lesson of the experiments.

The second central objection can now be framed in terms of the key result just noted—if the students really are not objectivists about hitting, then their responses to the moral/conventional task are confused. That is, the patterns of the nonobjectivists' responses reveal that they hold an inconsistent normative view. More generally the worry is that a nonobjectivist cannot be justified in drawing the moral/conventional distinction [6]. People who think morality is not objective have no business distinguishing morality from convention.

To address this objection, it's important to pull apart two different claims about the moral/conventional distinction. One might view the moral/conventional distinction as a sharp analytic distinction on which morality is entirely nonconventional. Alternatively, one might view the moral/conventional distinction as a measure of moral cognition in terms of a cluster of differential responses to moral and conventional violations. On the first, analytic, view of the moral/conventional distinction, the objection is that it is inconsistent to view morality as entirely nonconventional while judging it to be nonobjective. Although there is room to dispute this claim there is no need. For even if it is true that nonobjectivists cannot be thoroughgoing nonconventionalists, that does not affect the argument developed here. The moral/conventional experiments reported here were not intended to show that the participants draw an analytic distinction between morality and convention. Rather the moral/conventional distinction was used as a measure of moral cognition.

Now then, is it inconsistent for moral nonobjectivists to make differential responses to moral and conventional violations? Are nonobjectivists unjustified in giving the normal responses to the moral/conventional task? It's hard to see why. It seems presumptuous to say that giving up objectivity means that I *should not* judge harmful violations to be more serious, that I should not judge the wrongfulness of hitting another as independent of the teacher's authority, or that I should not think

that the actions are wrong because they are harmful. As philosophical sentimentalists have long maintained, you do not have to believe that an action is objectively wrong to have a deep and abiding opposition to such actions.

11. Conclusion

The goal of this paper has been to bring an empirical approach to bear on issues surrounding moral nonobjectivism. Obviously the results here are preliminary. Nonetheless, all five of the experiments reported here provide prima facie evidence that a sizable group of undergraduates are nonobjectivists about some canonical moral transgressions. Participants who respond as nonobjectivists about hitting, for example, still say that hitting is impermissible, just like moral objectivists do. The longstanding dispute over concept individuation prevents us from determining whether nonobjectivists use the same moral concepts as objectivists in such cases. However, we do have the beginnings of an answer to perhaps a more important issue. One popular fear about anti-objectivism is that the rejection of moral objectivism will usher in a flagrant nihilism. Nonobjectivists, the worry goes, will only follow moral rules opportunistically, like Hume's sensible knave. The studies reported here point in exactly the opposite direction from this worry. Participants who treat canonical moral violations as nonobjective also treat such violations in typical ways on the standard measure of moral judgment. They treated moral violations as more serious than conventional violations. They treated moral violations as less contingent on authority than conventional violations. Further, they offered standard welfare-based justifications for why the moral violations were wrong, but they offered standard social-conventional explanations for why the conventional violations were wrong.

In the passage from "The Sceptic" that begins this paper, Hume intimates that relinquishing objectivity will have little effect on our moral lives. One need not agree with the details of Hume's moral theory to appreciate his insight that the rejection of objectivity will not radically alter our moral lives. The present paper might be viewed as supplying some empirical confirmation for Hume's speculation. The evidence indicates that, on the prevailing measure of moral judgment, abandoning objectivity will scarcely dent the rest of commonsense moral judgment.

Acknowledgements

I would like to thank William Bechtel, Michael Gill, Edouard Machery, three anonymous reviewers, and especially Ned Hettinger for very helpful comments on this paper. Thanks also to the many participants who suffered through the materials.

Notes

[1] Noncognitivists, of course, maintain that moral utterances (e.g. "it's wrong to hit people") do not express judgments or beliefs of the speaker but only noncognitive states (like disapproval). As a result, noncognitivists typically reject the claim that commonsense is committed to objectivity. Blackburn and other noncognitivists have developed tremendously clever devices for accommo-

dating the apparent objectivism of commonsense morality in a noncognitivist theory (e.g. Blackburn, 1984, 1985). These noncognitivist pyrotechnics will not be engaged here. For the goal in this paper is to begin by taking commonsense moral thought at face value. On this approach we want to avoid, at least initially, invoking the kinds of subtle reinterpretations of moral discourse that are offered by noncognitivists. So, although it's possible that a noncognitivist account can be given of all the results presented in this paper, I will not take up the issue here.

[2] Although these participants gave nonobjectivist responses about whether the earth is flat, it's not at all clear that they are really metaphysical nonobjectivists. Perhaps a more plausible interpretation is that they misunderstand the question. Of course, that interpretation provides all the more reason to exclude these participants. For if these individuals respond as moral nonobjectivists, this might be attributed to their failure to understand the question.

[3] It might turn out that for more outrageous moral violations, like racist murder, the participants would not sustain their nonobjectivism. For present purposes, it suffices that the participants seem to think that a canonical example of an immoral action is not objectively wrong. For that is enough to call into question the idea that people are thoroughgoing moral objectivists. This issue is revisited briefly in Section 10.

[4] The following treatment is shaped by the discussion of a related point in Nichols and Stich (2003).

[5] I'm indebted to Bill Bechtel and to an anonymous referee for pressing me on this issue.

[6] This objection was raised by an anonymous referee.

References

Blackburn, S. (1984). *Spreading the word: groundings in the philosophy of language*. Oxford: Oxford University Press.

Blackburn, S. (1985). Errors and the phenomenology of value. In T. Honderich (Ed.), *Morality and objectivity* (pp. 1–22). London: Routledge & Kegan Paul.

Blair, R. (1995). A cognitive developmental approach to morality. *Cognition*, *57*, 1–29.

Block, N. (1986). Advertisement for a semantics for psychology. In P. French, T. Uehling Jr. & H. Wettstein (Eds) *Midwest studies in philosophy, vol. 10: studies in the philosophy of mind* (pp. 73–121). Minneapolis, MN: University of Minnesota Press.

Brink, D. (1989). *Moral realism and the foundation of ethics*. Cambridge: Cambridge University Press.

Darwall, S. (1998). *Philosophical ethics*. Boulder, CO: Westview Press.

Doris, J. & Stich, S. (forthcoming). As a matter of fact: empirical perspectives on ethics. In F. Jackson & M. Smith (Eds) *The Oxford handbook of contemporary analytic philosophy*. Oxford: Oxford University Press.

Fodor, J. (1998). *Concepts*. Oxford: Oxford University Press.

Harman, G. (1977). *The nature of morality*. Oxford: Oxford University Press.

Harman, G. (1985). Is there a single true morality? In D. Copp & D. Zimmerman (Eds) *Morality, reason and truth: new essays on the foundations of ethics* (pp. 27–48). Totowa, NJ: Rowman & Allanheld.

Harman, G. & Thomson, J. (1996). *Moral relativism and moral objectivity*. Oxford: Blackwell. Q3

Hume, D. (1739/1964). *A treatise of human nature*. Oxford: Clarendon Press.

Hume, D. (1742/1987). The sceptic. In E. Miller (Ed.) *Hume's essays, moral, political, and literary* (pp. 159–180). Indianapolis, IN: Liberty Classics.

Laurence, S. & Margolis, E. (1999). Concepts and cognitive science. In E. Margolis & S. Laurence (Eds) *Concepts: core readings* (pp. 3–81). Cambridge, Mass: MIT Press.

Loeb, D. (1998). Moral realism and the argument from disagreement. *Philosophical Studies*, *90*, 281–303.

Mackie, J. (1977). *Ethics: inventing right and wrong*. London: Penguin.

Murphy, G. & Medin, D. (1985). The role of theories in conceptual coherence. *Psychological Review*, *92*, 289–316.

Nichols, S. (2002). Norms with feeling: towards a psychological account of moral judgment. *Cognition*, *84*, 221–236.

NICHOLS, S. (forthcoming). *Sentimental rules: on the natural foundations of moral judgment*. New York: Oxford University Press.
NICHOLS, S. & FOLDS-BENNETT, T. (forthcoming). Are children moral objectivists? Children's judgments about moral and response-dependent properties. *Cognition*.
NICHOLS, S. & STICH, S. (2003). *Mindreading*. Oxford: Oxford University Press.
NUCCI, L. (2001). *Education in the moral domain*. Cambridge: Cambridge University Press.
RACHELS, J. (2002). *Elements of moral philosophy*. New York: McGraw Hill.
SMETANA, J. (1993). Understanding of social rules. In M. BENNETT (Ed.) *The development of social cognition: the child as psychologist* (pp. 111–141). New York: Guilford Press.
SMETANA, J. & BRAEGES, J. (1990). The development of toddlers' moral and conventional judgments. *Merrill-Palmer Quarterly*, *36*, 329–346.
SMITH, M. (1994). *The moral problem*. Oxford: Blackwell.
STICH, S. & WEINBERG, J. (2001). Jackson's empirical assumptions. *Philosophy and Phenomenological Research*, *62*, 637–643.
TISAK, M. (1995). Domains of social reasoning and beyond. In R. VASTA (Ed.) *Annals of child development, vol. 11* (pp. 95–130). London: Jessica Kingsley.

Appendix

Disgusting violation item for experiment 1:

> Susan and Mary are members of different cultures, and they are in an argument. Susan says, "It's okay to drink your own vomit so long as you microwave it first to kill the germs," and Mary says "No, it is not okay to drink your own vomit regardless of whether you microwave it first." Susan then says, "Look you are wrong. Everyone I know agrees that it's okay to do that." Mary responds, "Oh no, you are the one who is mistaken. Everyone I know agrees that it's not okay to do that."
>
> Which of the following do you think best characterizes their views? (Check one and give a brief justification for your answer.)
>
> —It is okay to drink your own vomit as long as you microwave it first, so Susan is right and Mary is wrong.
> —It is not okay to drink your own vomit even if you microwave it first, so Mary is right and Susan is wrong.
> —There is no fact of the matter about unqualified claims like "It's okay to drink your own vomit so long as you microwave it first." Different cultures believe different things, and it is not absolutely true or false that it's okay to drink your own vomit.

Conventional violation item:

> Mark and William are members of different cultures, and they are in an argument. Mark says, "It's okay to drink soup out of a bowl," and William says "No, it is not okay to drink soup out of a bowl." Mark then says, "Look you are wrong. Everyone I know agrees that it's okay to do that." William responds, "Oh no, you are the one who is mistaken. Everyone I know agrees that it's not okay to do that."
>
> Which of the following do you think best characterizes their views? (Check one and give a brief justification for your answer.)
>
> —It is okay to drink soup out of a bowl, so Mark is right and William is wrong.
> —It is not okay to drink soup out of a bowl, so William is right and Mark is wrong.
> —There is no fact of the matter about unqualified claims like "It's okay to drink soup out of a bowl." Different cultures believe different things, and it is not absolutely true or false that it's okay to drink soup out of a bowl.

Author Query Sheet

Manuscript Information	
Journal Acronym	CPHP
Volume and issue	17/1
Author name	Nichols
Manuscript No. (if applicable)	

AUTHOR: The following queries have arisen during the editing of your manuscript. Please answer the queries by marking necessary corrections **at the appropriate positions on the PROOFS. Do not answer the queries on the query sheet itself. Please also return a copy of the query sheet with your corrected proofs.**

QUERY NO.	QUERY DETAILS
AQ1	Reference in text given as Harman 1996. Full reference details in Reference section is given to Harman & Thomson 1996. Should text reference be amended to Harman & Thomson 1996?
AQ2	Should 'p' in the expression p<.05, be in italics?
AQ3	See AQ1, no text ref to Harman & Thomson 1996. Delete from References if necessary.