

Folk Moral Relativism

[Forthcoming in *Mind & Language*]

Hagop Sarkissian¹, John Park², David Tien³, Jennifer Cole Wright⁴ and Joshua Knobe^{5,6}

¹Department of Philosophy, Baruch College; ²Department of Philosophy, Duke University; ³Department of Philosophy, National University of Singapore; ⁴Department of Psychology, College of Charleston; ⁵Program in Cognitive Science, Yale University; ⁶Department of Philosophy, Yale University

ABSTRACT: It has often been suggested that people's ordinary folk understanding of morality involves a rejection of moral relativism and a belief in objective moral truths. The results of six studies call this claim into question. Participants did offer apparently objectivist intuitions when confronted with questions about individuals from their own culture, but they offered increasingly relativist intuitions as they were confronted with questions about individuals from increasingly different cultures or ways of life. In light of these data, the authors hypothesize that people do not have a fixed commitment to moral objectivism but instead tend to adopt different views depending on the degree to which they consider radically different perspectives on moral questions.

Suppose that two individuals are discussing mathematics. One of them claims that the number 2,377 is prime, while the other claims that it is not prime. In a case like this, it is usually assumed that one of the two individuals must be wrong. There is a fact of the matter about whether 2,377 is prime, and anyone who holds the other opinion has to be mistaken.

But now suppose we switch to a different topic. Two individuals are talking about the seasons. One of them claims that January is a winter month, while the other claims that it is a summer month. Faced with this latter case, we might well reach a different conclusion. There is no single objective fact about whether January is a winter month or a summer month. Rather, it can only be a winter month or a summer month relative to a specific hemisphere. So if one of them is talking about the northern hemisphere and the other is talking about the southern hemisphere, they can make seemingly opposite claims but still both be correct.

What about moral claims? Suppose that two individuals are talking about the moral status of a particular action. One claims the action is morally bad, while the other claims it is not morally bad. Must one of these individuals be wrong, or could it turn out that they are both right?

Within the philosophical literature, this question remains controversial. Some philosophers say that there is a single objective truth about whether a particular action is morally bad, so that if two individuals hold opposite opinions, one of them must be mistaken (Railton, 1986; Shafer-Landau, 2003; Smith, 1994). Other philosophers say that moral claims can only be assessed relative to a particular

moral framework or set of values, so that different moral claims could be right when asserted by different individuals (Dreier, 1990; Harman, 1975; Pinillos, 2010; Prinz, 2007; Wong, 1984; 2006). The debate between these two views has persisted at least since the ancient Greeks and shows no sign of letting up.

Our aim here is to explore what ordinary people think about this age-old philosophical question. Do people believe in objective moral truth, or do they accept some form of moral relativism?¹

Prior Work

Regardless of the position being defended, the usual assumption within the philosophical literature is that people subscribe to some form of moral objectivism. For example, Michael Smith writes that ordinary folk

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. (Smith, 1994, 6).

This claim that ordinary folk are moral objectivists enjoys a surprising degree of consensus in moral philosophy, and can be found in the works of a diverse range of moral philosophers with disparate theoretical commitments (e.g. Blackburn, 1984; Brink, 1989; Mackie, 1977; Shafer-Landau, 2003; Smith, 1994). Of course, philosophers hold very different views about how the study of folk intuitions can contribute to moral philosophy (Appiah, 2008; Kagan, 2009; Kauppinen, 2007; Knobe & Nichols, 2008; Ludwig, 2007), and even about whether folk intuitions have any relevance at all (Singer, 1974; 2005). Yet in spite of these important differences, the claim that the folk believe in some form of moral objectivism is widespread in moral philosophy.

¹ Note on terminology: The word 'relativism' is used in different ways in different disciplines. In moral philosophy, it is used broadly to describe any view according to which moral claims can only be assessed relative to a particular culture or system of values (e.g. Harman, 1975; Wong, 1996; 2006), whereas in formal semantics, it is used to pick out one very specific way of understanding this relativity (e.g. MacFarlane, 2007). We will be using the term in the broader, non-technical sense in which it appears within moral philosophy and will not be exploring the more specific questions that arise within formal semantics.

Now, this claim of folk moral objectivism is, on the face of it, an empirical claim—one that is amenable to systematic investigation. Researchers have for some time been exploring this very question, and the traditional philosophical view has enjoyed considerable empirical support. Results of many studies have thus far suggested that people reject relativism about morality, and believe instead in some type of absolute moral truth.

For example, in a pioneering study, Nichols (2004) presented participants with a story about two individuals—John and Fred – who appeared to hold different moral views. 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.” Participants were then asked to choose which of the following options best represented their own views:

1. It is okay to hit or shove people just because you feel like it, so John is right and Fred is wrong.
2. It is not okay to hit or shove people just because you feel like it, so Fred is right and John is wrong.
3. 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 people just because you feel like it.

Approximately 78% of participants chose either the first or second option. In other words, the majority of participants seemed to reject the idea that, e.g., hitting or shoving might be both good and bad—good relative to one person yet bad relative to another. On the contrary, they appeared to suggest that there was some definite fact of the matter as to whether hitting or shoving was morally good or bad, and that anyone who held the opposite opinion must be mistaken.

In a subsequent study, Goodwin and Darley (2008) provided participants with a large number of statements purporting to make claims about how things are in factual matters (“the earth is not at the center of the known universe”), moral matters (“consciously discriminating against someone on the basis of race is morally wrong”), conventional matters (“talking loudly and constantly to the person next to you during a lecture is a permissible action”), and matters of taste (“classical music is better than rock music”). After being asked whether they agreed with these statements, participants were told that none of the statements had, in fact, elicited full agreement; whatever judgment they happened to hold,

there was some fellow participant holding the opposite view. Once again, participants were asked how they would interpret such disagreements. This time, the options were:

1. The other person is surely mistaken.
2. It is possible that neither you nor the other person is mistaken.
3. It could be that you are mistaken, and the other person is correct.
4. Other.

In the moral cases, 70% of participants answered that the other individual had to be incorrect in her moral judgments (option 1). For Goodwin and Darley, this suggests that people are highly objectivist about many canonical moral transgressions—indeed, only somewhat less objectivist than they are about factual questions (such as whether the earth is at the center of the known universe).

Importantly, this same result has been found across different age groups. Wainryb and colleagues (2004) presented children ages 5, 7, and 9 with cases of disagreement in a number of domains, including the moral domain. For example, subjects were told that Sarah believes it's okay to hit and kick other children, whereas Sophie believes it is not okay to hit and kick other children. They were then asked whether both Sarah and Sophie could be correct and, if not, which of the two was correct and why. The results strongly supported the view that people are objectivists about morality: 100% of the 5 and 7 year olds and 94% of the 9 year olds thought that, in cases of such disagreement, only one of the individuals could be correct. In fact, children were as objectivist about moral disagreements as they were about purely factual disagreements (e.g., disagreement about whether pencils fall down or shoot up when you drop them).

Across all these studies, the same method has revealed the same result, time and again. It seems that ordinary, pre-philosophical folk reject the notion that moral disagreements can admit of many different answers, each of them correct relative to the person who is making the judgment. Rather, the folk appear to believe that when individuals hold opposite opinions about a moral question, only one of those individuals can be correct.

A New Hypothesis

In our view, however, the appearances here are deceiving. It is true that existing studies have consistently elicited apparently objectivist responses. However, we will argue that this pattern only

arises because of a particular feature of the existing experimental procedures. When the studies are conducted in a slightly different way, a more complex pattern of responses begins to emerge.

To get a better sense for the issue here, consider people's ordinary way of talking about the seasons. Many people understand that at a time of year it can only be said to be a particular season relative to a given hemisphere. So they understand that if individuals from different hemispheres make seemingly opposite claims about the seasons, there is no reason to conclude that one of these individuals has to be wrong. It is always possible that both of them are saying something perfectly true.

Now suppose we decided to study people's attitudes toward the seasons using a fairly straightforward experimental design. Participants would be told that one individual says 'January is a winter month' while another says 'January is not a winter month.' How might participants respond in a case like that? Would they say that one of these individuals had to be wrong, or would they say that both could actually be right?

Well, it depends. Do the participants assume that the two individuals are in different hemispheres, or do they assume that they are both in the same hemisphere? As long as experimental materials are designed in such a way that participants tend to think, e.g., that both individuals are living in the United States, they will presumably conclude that one of the individuals has to be wrong. Yet such a response would not show that people are objectivists about the seasons—that they think there is some objective truth about what season it is at any given time, full stop. Rather, the lesson to draw here is that people's deeper understanding of the relativity about the seasons will only come out clearly if they are encouraged to take a broader view and consider a number of possible perspectives—for example, if they are asked about a case involving both Americans and Australians.

Or consider a more extreme example. Suppose we asked people about the claim that there are twenty-four hours in a day. Faced with a case like this, people might initially think, 'That claim is objectively true, and anyone who holds the opposite opinion must surely be mistaken.' But now suppose we tried to broaden their perspective. Suppose we encouraged them to think about other planets, emphasizing that different planets take different amounts of time to rotate around their axes, with some taking less than twenty-four hours and some taking more. People might then begin to have a different intuition. They might begin to think, 'I guess there just happen to be twenty-four hours in a day on this particular planet, but if there are individuals on other planets, they might be perfectly right to say that the number of hours in a day was higher or lower.'

We want to suggest that a similar effect arises in the domain of morality. People do sometimes display what appear to be objectivist intuitions about morality, but it would be a mistake just to

conclude straightaway that people are moral objectivists. The more accurate thing to say is that people's intuitions depend on the precise way in which they are thinking of the question. As long as they are thinking only about individuals who are fairly similar to themselves—say, individuals from their own cultural groups—their intuitions might look more or less objectivist. But we will argue that people's intuitions do not always have this objectivist character. On the contrary, people's intuitions undergo a systematic shift as they begin considering different sorts of individuals. So as they come to think more and more seriously about individuals who are deeply dissimilar—individuals with radically different cultures, values, or ways of life—their intuitions move steadily toward a kind of relativism. They gradually come to feel that even if two individuals have opposite opinions on some moral question, it could still turn out that neither one would have to be wrong.

Study 1

As an initial test of this hypothesis, we conducted a study in which participants randomly received either a straightforward case like those used in earlier studies, or a case that specifically encouraged them to consider different cultures and ways of life.

All participants read vignettes that described an agent performing a behavior and other individuals judging that behavior. The difference between conditions lay only in the description of the individuals judging the behavior. Some participants were told that the individuals were Americans much like themselves; others were told that one of the individuals was from a very different sort of society with a radically different form of life. (Instead of simply stating that the individual was from a different society, the stimulus materials were designed to present this other society in as vivid and compelling a way as possible.) The hypothesis was that as people were encouraged to think about ever more distant forms of life, they would come to have ever more relativist intuitions.

Methods

Participants. Two hundred twenty-three students (112 females) taking introductory philosophy courses at Baruch College voluntarily completed a questionnaire after class.

Materials and Procedure. Each subject was randomly assigned to one of three conditions: *same-culture*, *other-culture* or *extraterrestrial*.

Participants in the same-culture condition were asked to imagine a person named Sam. They were told that Sam was a fairly ordinary student at their own college who enjoyed watching college

football and hanging out with friends. They were then asked to consider Sam's opinions about two moral transgressions. For one of these transgressions, they were given the sentence:

Horace finds his youngest child extremely unattractive and therefore kills him.

They were told to imagine that one of their classmates thinks that this act is morally wrong but that Sam thinks that the act is morally permissible. The instructions then were:

Given that these individuals have different judgments about this case, we would like to know whether you think at least one of them must be wrong, or whether you think both of them could actually be correct. In other words, to what extent would you agree or disagree with the following statement concerning such a case

Since your classmate and Sam have different judgments about this case, at least one of them must be wrong.

Responses to this question were recorded on a scale from 1 to 7. To test the generality of the effect, we also included a second transgression:

Dylan buys an expensive new knife and tests its sharpness by randomly stabbing a passerby on the street.

Instructions for this second transgression were exactly the same as for the first. (We did not predict any significant differences between the two transgressions.)

Participants in the other-culture condition were asked to imagine an isolated tribe of people called the Mamilons. They were told that this tribe lives in the Amazon rainforests and has preserved a traditional warrior culture, with quite different values from those of the people in the surrounding society. Participants in this condition then received exactly the same two questions as those received by participants in the same-culture condition, except that they were asked to imagine that the individual regarding the transgressions as morally permissible was a Mamilon.

Finally, participants in the extraterrestrial condition were asked to imagine a race of extraterrestrial beings called Pentars. They were told that the Pentars have a very different sort of psychology from human beings, that they are not at all interested in friendship or love and that their main goal is simply to increase the total number of equilateral pentagons in the universe. These participants then received the two questions, this time with a Pentar as the individual who regards the transgressions as permissible.

Results

The data were analyzed using a mixed-model ANOVA, with condition as a between-participants factor and transgression as a within-subject factor. There was no main effect of transgression and no interaction effect. However, there was a significant main effect of condition, $F(2, 218) = 20.7, p < .001, \eta^2 = .16$.

Ratings for the two transgressions were highly correlated ($r = .82, p < .001$) and could therefore be averaged to form a scale. The mean for each condition is displayed in Figure 1.

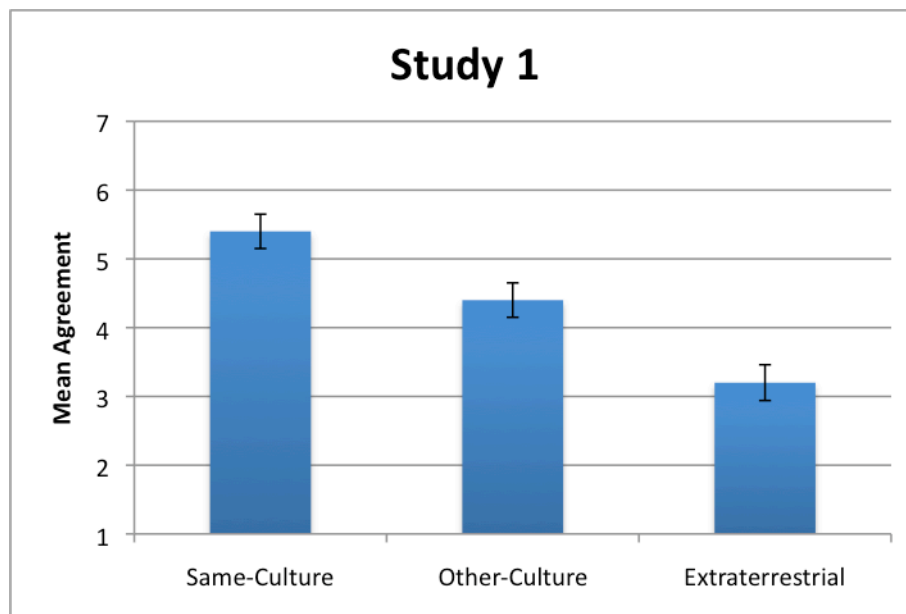


Figure 1. Mean agreement with the claim that ‘At least one must be wrong’ by condition. Error bars show SE mean.

Participants in the *same-culture* condition tended to agree that at least one person had to be wrong ($M = 5.4, SD = 2.15$), those in the *other-culture* condition were approximately at the midpoint ($M = 4.4, SD = 2.05$), and those in the *extraterrestrial* condition tended to say that both could actually be right ($M = 3.2, SD = 2.28$). Post-hoc Tukey’s tests showed significant differences both between responses in the other-culture condition and the same-culture condition ($p < .05$) and between responses in the extraterrestrial condition and the other-culture condition ($p < .01$).

Discussion

The results of this first experiment allow us to locate the results from earlier studies in a broader framework. Those earlier studies demonstrated that when two individuals hold opposite moral views, people think that at least one of those individuals has to be wrong. The present study replicates that basic finding, but also shows that it only arises under certain quite specific conditions—namely, when the individuals are from the same culture. As the individuals under discussion become ever more dissimilar, people become less and less inclined to agree with the claim that one of them has to be wrong, so that when the two individuals become dissimilar enough, people were far more willing to say that both of them can be right.

Overall, then, the responses we find in the same-culture condition do not appear to reflect any kind of general, across-the-board commitment to moral objectivism. On the contrary, it seems that people's intuitions are highly sensitive to the nature of the question posed. The more people are encouraged to consider radically different cultures and ways of life, the more they are drawn to more relativist responses.

But, of course, it would be a big mistake to draw any far-reaching conclusions from just this one experiment. The only way to get a proper understanding of what these results mean is to conduct follow-up studies designed to examine specific hypotheses about the nature of the effect observed here.

Study 2

To begin with, the participants in Study 1 were all American undergraduates. It might therefore be supposed that these initial results are not revealing anything important about the nature of moral cognition per se but are simply showing us something about the idiosyncrasies of contemporary American culture. (For example, American students might think it would be politically incorrect to pass judgments on individuals from other cultures and declare their moral judgments to be wrong.) To address this worry, we conducted a follow-up experiment with a population of subjects from a different culture, one that focuses less on individuals and more on group level, communal norms.

Methods

Participants. One hundred and fifty-one students (71 females) taking introductory philosophy courses at National University of Singapore voluntarily completed a questionnaire at the start of class.

Materials and procedure. The materials and procedure were the same as used in Study 1.

Results. A mixed-model ANOVA was run, with culture of the subject as a between-subjects factor and transgression as a within-subject factor. There was a significant main effect of the culture of the subject, $F(2, 146) = 3.1, p = .05, \eta^2 = .04$. There was no significant difference between transgressions and no significant interaction effect. The means per condition are reported in Figure 2.




Figure 2. Mean agreement with the claim that ‘At least one must be wrong’ by condition. Error bars show SE mean.

Discussion

The pattern of responses in this Singaporean sample mirrored the pattern of results found with Americans. Here again, the more participants were encouraged to think about individuals with very different cultures or ways of life, the more they were inclined to endorse relativist claims.

Of course, the fact that this same pattern emerged in two different cultures should not be taken as proof that the pattern is some sort of cross-cultural universal. It is possible, and indeed quite likely, that different cultures have quite different intuitions about the objectivity of moral claims. Still, the present results do provide strong reason to suspect that the pattern obtained here is not simply a reflection of some idiosyncratic feature of contemporary American culture. It appears that we actually are getting at something of importance about moral cognition.

Study 3

In these first two studies, participants gave different responses depending on whether they received the same-culture, other-culture or extraterrestrial cases. We now wanted to know whether participants would continue to show this effect even if they received all three cases side-by-side. Previous research has shown that in cases where people themselves regard an aspect of their moral judgments as a kind of 'bias' or 'error,' differences between conditions disappear when participants are able to see all of the conditions at once (Hsee, Loewenstein, Blount & Bazerman, 1999). We wanted to know whether a similar pattern would arise here or whether participants would affirm the difference between conditions even when they were presented with all three at the same time.

Methods

Participants. 61 students (52 females), taking Introduction to Psychological Science courses at the College of Charleston signed up for the study for research participation credit.

Materials & Procedures. Participants in this study received all three conditions (*same-culture*, *other-culture*, and *extraterrestrial*) in counterbalanced order (1: Sam/Mamilon/Pentar, 2: Mamilon/Pentar/Sam, 3: Pentar/Sam/Mamilon). Otherwise, the stories and questions were identical to those used in Study 1. Students signed up for the study and then were sent an online survey link to the questionnaires, which they were instructed to complete on their own in a quiet setting.

Results

The mean responses for each cultural condition are reported in Figure 3.




Figure 3. Mean agreement with the claim that 'At least one must be mistaken' by condition. Error bars show SE mean.

The data were analyzed using a mixed-model ANOVA, with cultural condition and transgression as within-participant factors and counterbalancing order as a between-participants factor, revealing a significant main effect for culture, $F(2,116) = 22.3, p < .001, \eta^2 = .28$. As in Study 1, participants were significantly more likely to give an objectivist response (at least one of the judgments had to be wrong) in the *same-culture* condition ($M = 5.2, SD = 2.02$) than in the *other-culture* condition ($M = 4.3, SD = 2.11$) and even less likely in the *extraterrestrial* condition ($M = 3.7, SD = 1.97$).

Discussion

Even when each participant received all three cases side-by-side, they continued to offer different responses depending on which sort of individual they were considering. Just as in the two earlier studies, participants were more inclined to reject the claim that at least one individual must be wrong as they moved toward individuals who were more deeply different in their culture or way of life.

These results show us something important about the nature of the effect under study here. It is not just that people respond differently in the different conditions; they seem actually to think that it is *right* to offer these different responses. So, even after they have given a clearly relativist response in one condition, they are happy to say in another condition that if the two individuals hold opposite opinions, one of them has to be wrong.

Study 4

The experimental results thus far seem to indicate that participants' intuitions change as they begin to consider individuals with radically different cultures or ways of life judging a single moral transgression. However, some might wonder whether important details are left out of the descriptions of the transgressions themselves. After all, there are *two* characters in our vignettes. First, there is the *agent*, the individual who actually commits the moral transgression (e.g., Horace, who kills his own child). Then, second, there is the *judge*, the individual who makes a judgment about whether the agent's action was morally bad or morally permissible (e.g., the Mamilon, an Amazonian tribesman). So far, we've only asked people to consider the latter of these two. But, details with respect to the former could arguably be relevant to people's moral judgments as well. For example, where is Horace located? Might people's judgments change if they are told that Horace is a member of their own culture—perhaps even a peer? Might people hold everyone to the same moral standards when judging an act taking place in their own culture or backyard? Conversely, would people apply very different standards when judging a similar act if took place far away, on another continent and in a different cultural context? Our claim has been that participants' intuitions about the rightness or wrongness of a judgment might actually depend on the identity of the *judge*, so that different judges could rightly arrive at different judgments of the very same agent. But what about the identity of the agent? Could this sway people's intuitions?

To get at this question more directly, we conducted a study in which we independently varied the cultural identity of both the agent and of the judge. Hence, the agent who performed the transgression could be either an American or someone from another culture, and the judge could be either an American, someone from another culture or an extraterrestrial. This design enabled us to clarify what impact the agent's identity might have on people's intuitions (apart from the judge's identity).

Methods

Participants. One hundred and eighteen students (91 females) taking Introduction to Psychological Science courses at the College of Charleston signed up for the study for research participation credit.

Materials & Procedure. Participants were given the same within-participants questionnaire as in Study 2 (counterbalancing the three judges: *same-culture*, *other-culture*, and *extraterrestrial*), only now they were divided into two groups: half received the *local agent transgression* condition, in which the

child-killer was represented as being American and the knife-stabber was represented as being a College of Charleston student, and half received the *foreign agent transgression* condition, in which the child-killer was represented as being Algerian and the knife-stabber was represented as being University of Algiers student.

Results

The data were analyzed using a mixed-design ANOVA, with the identity of the judge (same-culture vs. other-culture vs. extraterrestrial) and transgression as within-participant factors and the identity of the agent (local vs. foreign) as between-participant factors. There was a main effect for the identity of the judge, $F(2,224) = 34.7, p < .001, \eta^2 = .24$, but no main effect of the identity of the agent. There was also a significant interaction effect, $F(2,224) = 3.1, p = .048, \eta^2 = .03$.

The means for each condition are displayed in Figure 4. Inspection of these means indicated that the interaction effect arose because participants gave especially objectivist responses when presented with the case involving a local agent and a same-culture judge.

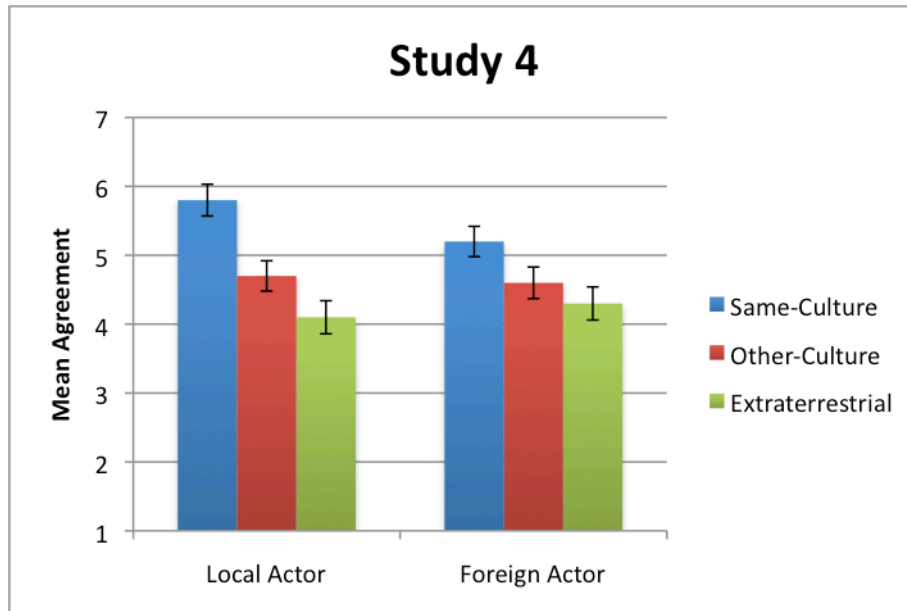


Figure 4. Mean agreement with the claim that 'At least one must be mistaken' by condition. Error bars show SE mean.

Discussion

Although we found no main effect from our manipulation of the identity of the agent, we did find a main effect from our manipulation of the identity of the judge. In other words, the pattern of intuitions observed here does not seem to be arising simply because people have different reactions to different kinds of agents. Rather, it seems that people can have different reactions to judgments about the very same act, performed by the very same agent, so long as we vary the identity of the judge.

There was a significant interaction effect, yet this was due, in large part, to the high levels of objectivist responses to the local actor transgression when judged by the two classmates (in the same-culture condition). This result is in line with the hypothesis of this paper. After all, it is made explicit that the agent and the judges are located in the same location and within the same moral framework. When the judges and the actor are all within the same moral framework, it seems reasonable to assume that participants would be least likely to agree with the claim that the differing responses of the judges could both be correct. It would be more likely to think that in such a situation people should be held to the same standards, that moral norms apply to them all in the same way. When frameworks align in such a fashion, it seems natural for people to assume that there can be an objective fact of the matter as to whether the actor's transgression is permissible—relative to that framework.

This pattern of intuitions suggests that people are evaluating each moral judgment relative to a set of standards that apply specifically to the judge in question, rather than the (potentially different) set of standards that apply to the agents who acted. Such a finding may initially seem a bit puzzling or bizarre, but work in natural language semantics has shown similar sorts of effects in other domains that have nothing to do with morality (Egan, Hawthorne & Weatherson, 2005; Lasersohn, 2009; MacFarlane, 2007; Stephenson, 2007), and it seems reasonable to suspect that the effects that arise in these other domains might apply in the moral domain as well.

A great deal of controversy remains about how exactly such effects are to be understood, and we cannot hope to resolve those difficult issues here. Nonetheless, the data do seem to be suggesting that the sorts of effects that arise concerning, e.g., judgments of taste can also be found when we turn to moral judgments. People appear to reject the idea that a single absolute standard can be applied to all moral judgments of a given agent and to operate instead with a system that applies different standards to different judges.

Study 5

In studies 1-4 subjects appear to be responding in ways suggesting that they are applying different sets of standards to the persons with differing judgments about the permissibility of a given

action. Our hypothesis has been that this is because ordinary folk might be tacitly committed to moral relativism. However, it is possible that subjects in these studies are not expressing a relativistic view about *morality* in particular. Instead, they might be expressing a more general, mad-dog sort of relativism—a relativism that applies not only to moral questions but also to purely descriptive questions, such as those that come up in science, history or mathematics.

In order to ascertain whether the folk really are relativists in this more radical sense, we conducted a follow-up experiment that made it possible to compare intuitions about moral claims with intuitions about non-moral claims.

Methods

Participants. Eighty-eight students (26 females) taking introductory philosophy courses at Duke University voluntarily completed a questionnaire at the end of class.

Materials and Procedure. Survey materials were distributed for participants to fill out. All participants were presented with the Other-Culture vignette used in Study 1. Participants were then randomly assigned to one of two conditions.

Participants in the *moral* condition were asked to consider the Mamilon's opinions about two moral transgressions. For the first transgression, they were given the sentence: 'Jason robs his employer, the Red Cross, in order to pay for a second holiday for himself.' They were then asked to imagine that one of their classmates believed what Jason did was morally wrong, but that the Mamilon thought what Jason did was not morally wrong. After reading about these different judgments, they were asked the same question as in Study 1—namely, whether they agreed with the statement 'Since your classmate and the Mamilon have different judgments about this case, at least one of them must be wrong.' Participants rated this sentence on a scale from 1 ('disagree') to 7 ('agree'). For the second transgression, participants were given the sentence: 'Emily promises to take Molly's sick child to the hospital for an important surgical procedure, but instead decides she'd rather go shopping.' Instructions for this other transgression were exactly the same as for the first. The order of transgressions was counterbalanced.

Participants in the *non-moral* condition were asked to consider the Mamilon's opinion about two non-moral cases. For the first case, participants were told the following: 'A group of individuals are discussing where pasta comes from. Alejandro thinks pasta is made by combining flour, water, and eggs, whereas Marya thinks pasta grows on trees and is harvested by special farmers called "Pastafarians" once every 5 years.' They were then asked to imagine that a classmate agrees with

Alejandro, and that a Mamilon agrees with Marya. After reading about these different judgments, they were asked the same question as in the *moral* condition. The second case was as follows: ‘A group of individuals are discussing the military strategies of Napoleon Bonaparte. Anita thinks that Napoleon rode into battle on a horse, whereas Fabio thinks Napoleon flew into battle in a helicopter.’ They were then asked to imagine a classmate agrees with Anita, whereas a Mamilon agrees with Fabio. Instructions for this other transgression were exactly the same as for the first. The order of transgressions was counterbalanced.

Results

Ratings for the two transgressions were highly correlated, both for the non-moral cases ($r = .67$, $p < .001$) and for the moral cases ($r = .75$, $p < .001$), and there were no significant differences between them. They could therefore be averaged to form a scale. The mean for each condition is displayed in Figure 5.

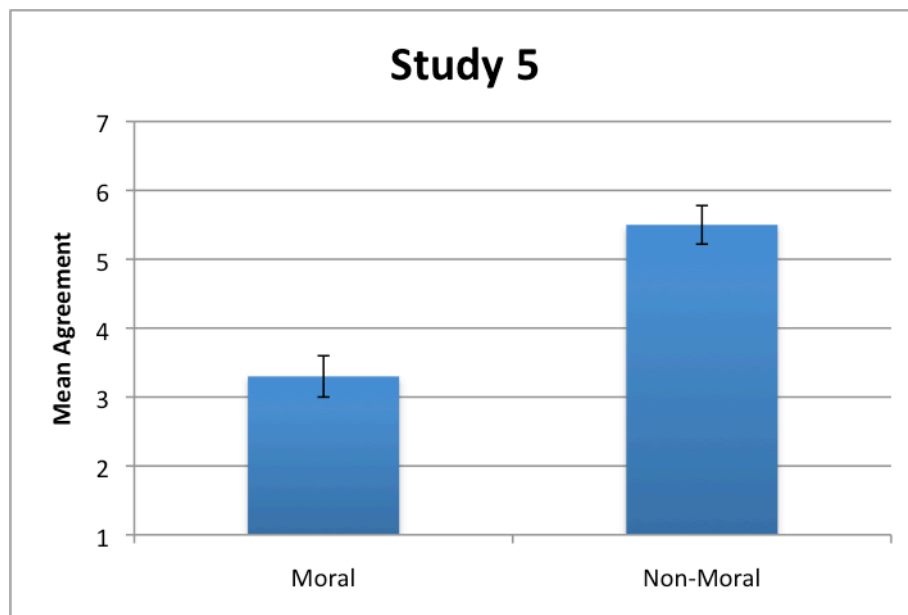


Figure 5. Mean agreement with the claim that ‘At least one must be wrong’ by condition. Error bars show SE mean.

Participants were strongly objectivist when the differing judgments were about non-moral cases ($M = 5.5$, $SD = 1.84$) but were significantly more inclined to embrace relativism when the differing judgments were about moral cases ($M = 3.3$, $SD = 1.99$), $t(88) = 5.3$, $p < .001$.

Discussion

Participants showed the usual tendency to endorse relativism about moral questions, but they did not show this same tendency when confronted with non-moral questions. Instead, their responses in the non-moral domain took a more objectivist turn. They tended to agree with the claim that if two individuals held opposite opinions about historical facts, then one of those individuals has to be wrong.

In short, it does not appear that people subscribe to a blanket relativism that applies equally to all issues. People seem instead to be distinguishing between different domains, endorsing relativist claims about morality but not about straightforward matters of fact.

Study 6

Finally, one might worry about how people are interpreting the question they receive in the stimulus materials. Participants in these studies are presented with stories about two judges with seemingly opposite opinions and then asked whether they agree with the claim that at least one of the judges must be 'wrong' or 'mistaken.' But how exactly are participants understanding the words 'wrong' or 'mistaken' in this context? Thus far, we have been assuming that participants understand these words to mean that the opinions of the judges are *not true*. So our assumption has been that when people disagree with the claim, they are saying something like: 'Even though the two judges have opposite opinions, it could be that both of their opinions are true.' This really would be an endorsement of moral relativism.

But it seems that the claim could actually be understood in other, very different ways. For example, one could interpret the words 'wrong' or 'mistaken' to mean something like *not a reasonable inference, given the available evidence*. On this alternative construal, participants are not saying anything deeply controversial about the truth of the judges' opinions. They are simply saying something straightforward of the form: 'Even though the two judges have opposite opinions, it could be that both have good reason to arrive at those inferences, given the available evidence.' Such a statement would not amount to an endorsement of genuine moral relativism. (In the jargon of philosophy, it would be classified as a purely 'epistemic' claim.)

To address this last worry, we conducted a study in which participants were asked both about whether the various judges' beliefs were 'incorrect' and about whether they had 'no good reason' to

believe what they did. This method makes it possible to disentangle intuitions about the truth of the judges' opinions from intuitions about whether those opinions were reasonable inferences from the available evidence.

Methods

Participants. Participants were ninety undergraduate students (45 females) taking introductory philosophy courses at Duke University.

Materials and Procedures. Survey materials were distributed for participants to fill out. As in Study 5, all participants were presented with the Other-Culture vignette used in Study 1. Participants were then randomly assigned to one of four conditions in a 2x2 design.

Half of the subjects received one of two 'truth' conditions. In these conditions, they were given either two moral cases (the robbing and promise-renegeing transgressions used in Study 5) or two non-moral cases (the Napoleon and pasta cases used in Study 5). They were then asked to imagine that a classmate and a Mamilon have differing judgments about whether, for example, the moral transgression was permissible, or whether Napoleon rode a horse or a helicopter, and were asked to what extent they agreed with the following statement: 'Given that your classmate and the Mamilon have these particular beliefs, at least one of their beliefs must be incorrect.' Participants rated this sentence on a scale from 1 ('disagree') to 7 ('agree').

The other half of the subjects received one of two 'justification' conditions. In these conditions, they were given either the same moral or non-moral cases from the 'truth' condition. When told that the classmate and the Mamilon had different judgments about the case, they were given the following question: "These individuals have different beliefs about this case. We would like to know whether you think only one of them has good reason to believe what he or she does, or whether they both have good reasons." They were then asked to what extent they agreed with the following statement: "Given the particular beliefs that your classmate and the Mamilon have, at least one of them must not have good reason to believe as he or she does." Participants rated this sentence on a scale from 1 ('disagree') to 7 ('agree').

Results

The data were subjected to an ANOVA, with question type (truth vs. justification) and vignette type (moral vs. non-moral) as between-subject factors, transgression as a within subject factor, and

gender and number of previous philosophy courses as covariates. There was no main effect of either question type or vignette type. However, there was a significant interaction effect $F(1,84) = 10.4, p = .002, \eta^2 = .11$.

Inspection of the means revealed opposite patterns of responses for the two question types (see Figure 6). On the question about truth, people tended to reject the claim that at least one of the judges had to be incorrect in the moral cases ($M = 2.93, SD = 2.15$) but not in the non-moral cases ($M = 4.67, SD = 2.17$), $t(43)=2.7, p=.01$. By contrast, on the question about justification, people tended to reject the claim that at least one of the judges must not have had good reason in the non-moral cases ($M = 3.0, SE = 1.98$) but not in the moral cases ($M = 4.4, SD = 2.08$), $t(43)=2.3, p=.03$.

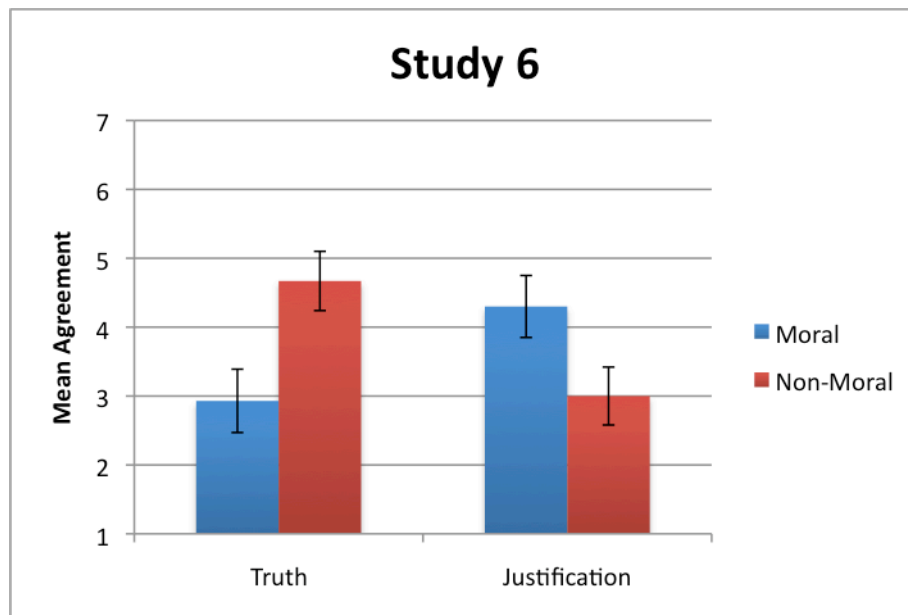


Figure 6. Mean agreement with the claim that 'at least one must be incorrect' (the 'truth question') and 'at least one must not have good reason to believe' (the 'justification question'). Error bars show SE mean.

Discussion

In this final study, participants showed a clear distinction between the intuitions about truth and intuitions about justification. On the question about truth, participants showed the same basic pattern of judgments they displayed in Study 5: they were less willing to say that at least one judge had to be 'incorrect' in the moral condition than they were in the non-moral condition. However, on the question about justification, they showed exactly the opposite pattern of judgments: they were

actually more willing to say that at least one judge had 'no good reason' in the moral condition than in the non-moral condition. In other words, people seem to be willing to grant that people from different cultures may legitimately hold different moral beliefs—that these beliefs need not be considered illicit or false—yet they are not willing to accept that such beliefs are equally justified—that there are equally good reasons supporting them.

Of course, the precise pattern of justification intuitions depends in part on the nature of the vignettes themselves, and we could doubtless have obtained a very different pattern if we had used different vignettes. The important point, however, is simply that participants are drawing a clear distinction between truth and justification. Hence, when we see that participants disagree with the claim that at least one of the beliefs must be incorrect, we have some reason to conclude that they are in fact expressing the relativist view that two judges can make opposite judgment without either of them thinking anything false.

General Discussion

The present studies offer a complex picture of people's intuitions about whether morality is objective or relative. People do have apparently objectivist intuitions in certain cases, but our results suggest that one cannot accurately capture their views in a simple claim like: 'People are committed to moral objectivism.' On the contrary, people's intuitions take a strikingly relativist turn when they are encouraged to consider individuals from radically different cultures or ways of life.

Overall, then, the pattern of people's intuitions about morality appears to resemble the pattern of their intuitions about the seasons. In the course of a typical conversation, people might give little thought to the idea that it can only be winter or summer relative to a given hemisphere. Indeed, people who live in the Northern Hemisphere might ordinarily suppose that a certain time simply *is* winter and that anyone who thinks otherwise must be mistaken. But this pattern of ordinary intuitions would not necessarily make it right for us to draw a conclusion of the form: 'People are committed to seasonal objectivism.' The thing to focus on is people's ability, when suddenly confronted with an individual from another hemisphere, to recognize that a single time might be winter around here but summer somewhere else.

Ordinary intuitions about morality appear to work in much the same way. Day to day, people might give little thought to issues of relativity. It might appear to them, when they witness a typical transgression, that the agent has clearly done something morally bad and that anyone who thinks otherwise must surely be wrong. Yet this pattern of intuitions is not itself sufficient to make them count

as moral objectivists. The key question is whether they hold on to that view even in the face of radically different perspectives. The results of the present studies suggest that they do not.

1. *Relation to previous work*

It may initially appear that our chief conclusions are opposed to those drawn by earlier researchers. After all, earlier researchers found that most participants were moral objectivists, whereas we are suggesting that many of those participants might actually have had a tendency toward moral relativism.

Our own view, however, is that the disagreement here is only a superficial one and that there is a deeper sense in which our conclusions are actually in harmony with those of earlier researchers. The main goal of research in this area is not to figure out precisely what percentage of people are moral objectivists or moral relativists but rather to reach a better understanding of the psychological processes that can pull people in these different directions. Our claim has been that people are drawn toward moral relativism by one specific type of psychological process: namely, active engagement with radically different perspectives and ways of life. So what we really want to know is whether prior research provides evidence for this same psychological hypothesis.

Though earlier studies found that a majority of participants gave apparently objectivist responses, each of these studies also found a minority who gave relativist responses. Correlational data from these studies can therefore be used to get a sense for the psychological differences between these two groups of participants. The results offer us a remarkably coherent picture. Relativists were higher in the personality trait of openness to experience (Cokely & Feltz, 2010). They scored higher on a measure of 'disjunctive thinking,' which is the ability to unpack alternative possibilities when problem solving (Goodwin & Darley, 2010). They were more likely to fall in a particular age range – namely, in their twenties (Beebe & Sackris, 2010). They were more able to explain alternative views (Goodwin & Darley 2010) and to be tolerant of people with opposite opinions (Wright, Cullum & Schwab, 2008; Wright, McWhite & Grandjean, 2010). All in all, these separate studies seem to be converging on a surprisingly unified portrait of relativist participants. Specifically, it appears that the relativists were precisely the people who were *most open to alternative perspectives*.

So perhaps it will be possible to develop a single unified account that explains the full pattern of data. The basic idea would be that people's intuitions are influenced by the degree to which they actively consider alternative perspectives. One factor that leads people to consider alternative

perspectives is the wording of the question itself (as in the present studies); another is the personality traits of the participants (as in prior studies). But regardless of the factors that trigger it, the impact of this psychological process is always the same. The more people engage with radically different perspectives, the more they are drawn to moral relativism.

2. Relation to philosophical metaethics

As noted at the outset of this paper, many philosophers think both that the folk are objectivists about morality and that the folk view bears on the philosophical truth about morality. In the words of Michael Smith, “the task of the philosopher in meta-ethics is to make sense of ordinary moral practice”—that is, ordinary folk objectivism (Smith, 1994, 5). Many philosophers have explicitly tried to fulfill this task.

For present purposes, we can divide the ways in which philosophers have sought to discharge the task into three general categories: 1) affirm that the folk are generally right; 2) explain why the folk are entirely wrong; 3) come with some more complicated interpretation of what the folk are committed to. Each of these represents an attempt to accommodate folk objectivism. We'll explore them in turn.

One way to account for the purported folk view is to say that this view is actually correct: there truly are such things as real, objective moral facts. Philosophers who adopt this ‘realist’ approach have sometimes claimed that they have a leg up on opposing theories because they can more easily account for the phenomenon of folk objectivism (Brink, 1989; Shafer-Landau, 2003). In the words of Shafer-Landau, only an account of this basic type “straightforwardly preserves ordinary talk of moral truth” (Shafer-Landau, 2003, 23).

A second approach would be to deny the existence of real, objective moral facts. Philosophers who adopt this approach maintain that the sorts of moral properties presupposed by ordinary folk simply do not exist, so the folk err in being objectivists about morality (Joyce, 2001; Mackie, 1977).

Mackie (1977, 33), for example, argued that such absolute moral properties – “not contingent upon any desire or preference or policy or choice” – are simply nonexistent. Nonetheless, in formulating his theory, Mackie ends up affirming the very same claim about ordinary folk morality. He says that ordinary people tend to “objectify values,” that “ordinary moral judgments include a claim to objectivity” (1977, 35).

Finally, some philosophers have taken a different tack, proposing that we could adopt a more complex interpretation of people's ordinary moral discourse. According to this interpretation, people's

ordinary moral claims don't mean precisely what they might appear to mean on first glance. Thus, when people ordinarily make claims like: 'That action is morally bad, and anyone who says otherwise must surely be mistaken,' they are not literally saying that there is some kind of objective moral truth in the way that 'moral realists' have supposed (Blackburn, 1984; Gibbard, 1992). Yet even while offering such a complex analysis, these philosophers maintain the usual view about the shape of ordinary folk discourse, suggesting that philosophical theories must endeavor to explain "why our discourse has the shape it does... to explain, and justify, the realistic-seeming nature of our talk of evaluations—the way we think we can be wrong about them, that there is a truth to be found, and so on" (Blackburn, 1984, 180).

Each of these positions tries to accommodate folk objectivism in one way or another: by claiming to best capture it, by dismissing it as riddled with error, or by providing a nuanced interpretation of it. But the results of the present studies raise some questions about whether this task is needed at all. If the folk are not, in fact, moral objectivists, then perhaps there is simply no need to continue engaging in philosophical work aimed at making sense of folk objectivism.

Nevertheless, a task remains. Philosophers are undoubtedly correct in their commitment to make sense of ordinary moral practice; the one mistake was to suppose that people's ordinary moral practice is a straightforwardly objectivist one. So perhaps the real philosophical task here is to make sense of a different sort of practice: one in which people's views differ depending on the extent to which they explore alternative perspectives.

3. Conclusion

Recent experimental research has investigated people's intuitions about a whole host of different philosophical questions – free will, moral obligation, personal identity, the nature of knowledge. Even though these questions are in many ways quite different, the outcomes of the various experimental research programs have been surprisingly convergent. Again and again, we find that it is not possible to capture the full pattern of people's intuitions just by constructing a coherent philosophical theory and then claiming that this theory captures the 'folk view.' Instead, it has emerged in each case that people's intuitions show certain kinds of tensions or contradictions, with different psychological processes pulling people in different directions. (For a few representative cases, see Greene, 2008; Nahmias, forthcoming; Nichols & Bruno, forthcoming; Nichols & Knobe, 2007; Phelan, 2010; Phelan & Sarkissian, 2009; Sinnott-Armstrong, 2008). But perhaps that is exactly what we should

have expected. Philosophers do not write about questions where the answer seems obvious; they write about issues that provoke conflict or confusion. What the recent experimental work has done is, in part, to trace such conflict and confusion back to certain tensions within people's ordinary intuitions.

The results of the present studies are very much in line with this general trend. It is, we think, a mistake just to say that ordinary people subscribe to some form of moral objectivism. People do have objectivist intuitions in certain cases, but these intuitions are the product of psychological processes that can, in other cases, lead to strikingly relativist intuitions. Future research might proceed not by asking whether 'people are objectivists' or 'people are relativists' but rather by trying to get a better grip on the different psychological processes at work here and the conflicts and tensions that these processes can create.

REFERENCES

- Appiah, K. A. 2008: *Experiments in Ethics*. Cambridge: Harvard University Press.
- Beebe, J. R., and Sackris, D. 2010: Moral objectivism across the lifespan. Unpublished Manuscript. University of Buffalo.
- Blackburn, S. 1984: *Spreading the Word: Groundings in the Philosophy of Language*. New York: Oxford University Press.
- Brink, D. O. 1989: *Moral Realism and the Foundations of Ethics*. New York: Cambridge University Press.
- Cokely, E. T., and Feltz, A. 2010: The truth about right and wrong: predicting moral objectivism with a heritable personality trait. Unpublished.
- Dreier, J. 1990: Internalism and speaker relativism. *Ethics*, 101, 6-26.
- Egan, A., Hawthorne, J., and Weatherson, B. 2005: Epistemic modals in context. In *Contextualism in Philosophy: Knowledge, Meaning, and Truth*. G. Preyer and P. George, eds., 131–170. New York: Oxford University Press.
- Gibbard, A. 1992: *Wise Choices, Apt Feelings: A Theory of Normative Judgment*. Cambridge: Harvard University Press.
- Goodwin, G. P., and Darley, J. M. 2008: The psychology of meta-ethics: Exploring objectivism. *Cognition*, 106, 1339.
- Goodwin, G. P., and Darley, J. M. 2010: The perceived objectivity of ethical beliefs: psychological findings and implications for public policy. *Review of Philosophy and Psychology*, 1, 1-28.
- Greene, J. D. 2008: The secret joke of Kant's soul. In *Moral Psychology Volume 3: The Neuroscience of Morality: Emotion, Disease, and Development*. W. Sinnott-Armstrong, ed. Cambridge: MIT Press.
- Harman, G. 1975: Moral relativism defended. *The Philosophical Review*, 84, 3-22.
- Hsee, C. K., Loewenstein, G. F., Blount, S., and Bazerman, M. H. 1999: Preference reversals between joint and separate evaluations of options: A review and theoretical analysis. *Psychological Bulletin*, 125, 576-590.
- Joyce, R. 2001: *The Myth of Morality*. Cambridge: Cambridge University Press.
- Kagan, S. 2009: Thinking about cases. *Social Philosophy and Policy*, 18, 44-63.
- Kauppinen, A. 2007: The rise and fall of experimental philosophy. *Philosophical explorations*, 10, 95-118.
- Knobe, J., and Nichols, S. 2008: An experimental philosophy manifesto. In *Experimental Philosophy*. J. Knobe and S. Nichols, eds., 3-14. New York: Oxford University Press.

- Lasersohn, P. 2009: Relative truth, speaker commitment, and control of implicit arguments. *Synthese*, 166, 359-374.
- Ludwig, K. 2007: The epistemology of thought experiments: first person versus third person approaches. *Midwest Studies in Philosophy*, 31, 128-159.
- MacFarlane, J. 2007: Relativism and disagreement. *Philosophical Studies*, 132, 17-31.
- Mackie, J. L. 1977: *Ethics: Inventing Right and Wrong*. New York: Penguin.
- Nahmias, E. forthcoming: Intuitions about free will, determinism, and bypassing. In *The Oxford Handbook on Free Will*. R. Kane, ed. New York: Oxford University Press.
- Nichols, S. 2004: After objectivity: An empirical study of moral judgment. *Philosophical Psychology*, 17, 3-26.
- Nichols, S., and Bruno, M. forthcoming: Intuitions about personal identity: An empirical study. *Philosophical Psychology*.
- Nichols, S., and Knobe, J. 2007: Moral responsibility and determinism: the cognitive science of folk intuitions. *Nous*, 41, 663-685.
- Phelan, M. T. 2010: Evidence that stakes don't matter for evidence. Unpublished Manuscript. Yale University.
- Phelan, M. T., and Sarkissian, H. 2009: Is the trade-off hypothesis worth trading for? *Mind & Language*, 24, 164-180.
- Pinillos, N. Á. 2010: Knowledge and moral relativism. Unpublished Manuscript. Arizona State University.
- Prinz, J. 2007: *The Emotional Construction of Morals*. New York: Oxford University Press, USA.
- Railton, P. 1986: Moral realism. *The Philosophical Review*, 95, 163-207.
- Shafer-Landau, R. 2003: *Moral Realism: A Defence*. New York: Oxford University Press.
- Singer, P. 1974: Sidgwick and reflective equilibrium. *The Monist*, 58, 490-517.
- Singer, P. 2005: Ethics and intuitions. *The Journal of Ethics*, 9, 331-352.
- Sinnott-Armstrong, W. 2008: Abstract + Concrete = Paradox. In *Experimental Philosophy*. J. Knobe and S. Nichols, eds., 209-230. New York: Oxford University Press.
- Smith, M. 1994: *The Moral Problem*. Oxford: Blackwell.
- Stephenson, T. 2007: Judge dependence, epistemic modals, and predicates of personal taste. *Linguistics and Philosophy*, 30, 487-525.
- Wong, D. B. 1984: *Moral Relativity*. Berkeley: University of California Press.
- Wong, D. B. 1996: Pluralistic relativism. *Midwest Studies in Philosophy*, 20, 378-99.
- Wong, D. B. 2006: *Natural Moralities: A Defence of Pluralistic Relativism*. New York: Oxford University Press.
- Wright, J. C., Cullum, J., and Schwab, N. 2008: The cognitive and affective dimensions of moral conviction: implications for attitudinal and behavioral measures of interpersonal tolerance. *Personality and Social Psychology Bulletin*, 34, 1461-1476.
- Wright, J. C., McWhite, C., and Grandjean, P. T. 2010: The cognitive mechanisms of intolerance: Do people's meta-ethical commitments matter? Unpublished Manuscript. College of Charleston.