

## How Different Kinds of Disagreement Impact Folk Metaethical Judgments

James R. Beebe\*

Although the empirical study of folk metaethical judgments is still in its infancy, a variety of interesting and significant results have been obtained.<sup>1</sup> Goodwin and Darley (2008), for example, report that individuals tend to regard ethical statements as more objective than conventional or taste claims and almost as objective as scientific claims, although there is considerable variation in metaethical intuitions across individuals and across different ethical issues. Goodwin and Darley (2012) also report (i) that participants treat statements condemning ethical wrongdoing as more objective than statements enjoining good or morally exemplary actions, (ii) that perceived consensus regarding an ethical statement positively influences ratings of metaethical objectivity, and (iii) that moral objectivism is associated with greater discomfort with and more pejorative attributions toward those with whom individuals disagreed. Beebe and Sackris (under review) found that folk metaethical commitments vary across different life stages, with decreased objectivism during the college years.

Sarkissian et al. (2011) found that folk intuitions about metaethical objectivity vary as a function of cultural distance, with increased cultural distance between disagreeing parties leading to decreased attributions of metaethical objectivity. Wright et al. (forthcoming) found that not only is there significant diversity among individuals with regard to the objectivity they attribute to ethical claims, there is also significant diversity of opinion with respect to whether individuals take certain issues such as abortion or anonymously donating money to charity to be ethical issues at all, despite the fact that philosophers overwhelmingly regard these issues as ethical.<sup>2</sup> Wright et al. (forthcoming) provide the following useful summary of the current set of findings on folk metaethical intuitions:

People do not appear to conceive of morality as a unified (meta-ethically speaking) domain, but rather as a domain whose normative mandates come in different shapes and sizes. They view the wrongness of some moral actions as clear and unquestionable, unaltered (and unalterable) by the feelings/beliefs/values of the individual or culture. They view the wrongness of other actions (though still genuinely moral in nature) as more sensitive to, and molded by, the feelings/beliefs/values of the actor and/or the people whose lives would be (or have been) affected by the action. This possibility is one we've not seen seriously considered in the meta-ethical literature—and perhaps it is time that it was.

The present article reports a series of experiments designed to extend the empirical investigation of folk metaethical intuitions by examining how different kinds of ethical disagreement can impact attributions of objectivity to ethical claims.

Study 1 reports a replication of Beebe and Sackris' work on metaethical intuitions, in order to establish a baseline of comparison for Studies 2 through 4. In Study 2, societal disagreement about ethical issues was made salient to participants before they answered metaethical questions about the objectivity of ethical claims, and this was found to decrease attributions of objectivity to those claims. In Studies 3 and 4, the parties with whom participants were asked to consider having an ethical disagreement were made more concrete than in Studies 1 and 2, using either verbal descriptions or facial pictures. This manipulation was found to increase attributions of metaethical objectivity. In a final study, metaethical judgments were shown to vary with the moral valence of the actions performed by the disagreeing party—in other words, a Knobe effect for metaethical judgments was found. These studies aim to increase our understanding of the complexity of the folk metaethical landscape.

## Study 1

### Method

#### **Participants**

Study 1 was an attempt to replicate Beebe and Sackris' (under review) initial study with a population of participants that was limited to the same university

student population from which participants for Studies 2 and 3 would be drawn. Participants were 192 undergraduate students (average age = 20, 53% female, 40% Anglo-American) from the University at Buffalo (a large, public university in the northeastern United States) in exchange for extra credit in an introductory course.

## **Materials**

Beebe and Sackris asked two and a half thousand participants between the ages 12 and 88 to indicate the degree to which they agreed or disagreed with the claims that appear in Table 9.1 and the extent to which they thought that “people in our society” disagreed about whether they are true. The same set of claims was used in Studies 1 through 3.

## **Procedure**

The items from Table 9.1 were divided into three questionnaire versions, and participants indicated their agreement or disagreement with them on a six-point scale, where “1” was anchored with “Strongly Disagree” and “6” with “Strongly Agree.” Participants rated the extent to which they thought people in our society disagreed about the various claims on a six-point scale anchored with “There is no disagreement at all” and “There is an extremely large amount of disagreement.”

In order to capture one kind of objectivity that participants might attribute to the various claims in Table 9.1, participants were asked, “If someone disagrees with you about whether [one of these claims is true], is it possible for both of you to be correct or must one of you be mistaken?” The answer “At least one of you must be mistaken” was interpreted as an attribution of objectivity, and an answer of “It is possible for both of you to be correct” was taken to be a denial of objectivity.

## **Results**

As can be seen from Figure 9.1, the items in Table 9.1 are ordered within each subcategory in terms of increasing proportions of participants who attributed objectivity to them.

**Table 9.1** Factual, ethical, and taste claims used in Beebe and Sackris (under review) and in Studies 1 through 4

---

*Factual*


1. Frequent exercise usually helps people to lose weight.
2. Global warming is due primarily to human activity (for example, the burning of fossil fuels).
3. Humans evolved from more primitive primate species.
4. There is an even number of stars in the universe.
5. Julius Caesar did not drink wine on his 21st birthday.
6. New York City is further north than Los Angeles.
7. The earth is only 6,000 years old.
8. Mars is the smallest planet in the solar system.

*Ethical*

9. Assisting in the death of a friend who has a disease for which there is no known cure and who is in terrible pain and wants to die is morally permissible.
10. Before the third month of pregnancy, abortion for any reason is morally permissible.
11. Anonymously donating a significant portion of one's income to charity is morally good.
12. Scientific research on human embryonic stem cells is morally wrong.
13. Lying on behalf of a friend who is accused of murder is morally permissible.
14. Cutting the American flag into pieces and using it to clean one's bathroom is morally wrong.
15. Cheating on an exam that you have to pass in order to graduate is morally permissible.
16. Hitting someone just because you feel like it is wrong.
17. Robbing a bank in order to pay for an expensive vacation is morally bad.
18. Treating someone poorly on the basis of their race is morally wrong.

*Taste*

19. Classical music is better than rock music.
  20. Brad Pitt is better looking than Drew Carey.
  21. McDonald's hamburgers taste better than hamburgers made at home.
  22. Gourmet meals from fancy Italian restaurants taste better than microwavable frozen dinners.
  23. Barack Obama is a better public speaker than George W. Bush.
  24. Beethoven was a better musician than Britney Spears is.
-



**Figure 9.1** Proportions of participants who attributed objectivity to the 24 items in Study 1. Error bars in all figures represent 95 percent confidence intervals.

Goodwin and Darley (2008) and Beebe and Sackris both found that more participants attributed objectivity to factual claims than to ethical or taste claims. In Study 1, a greater proportion of participants attributed objectivity to factual claims (0.64, averaged across all claims in the factual subcategory) than to ethical (0.34) or taste (0.10) claims. Chi-square tests of independence reveal that the difference between the factual and ethical proportions was significant,  $\chi^2(1, N = 926) = 80.523, p < 0.001$ , Cramér's  $V = 0.30$ , and the difference between the ethical and taste proportions was significant as well,  $\chi^2(1, N = 826) = 61.483, p < 0.001$ , Cramér's  $V = 0.27$ .<sup>3</sup> Study 1 also replicates earlier findings that objectivity attributions are positively associated with strength of belief about an issue ( $\chi^2(2, N = 1,224) = 67.276, p < 0.001$ , Cramér's  $V = 0.23$ ) but negatively associated with the extent of perceived disagreement about the issue ( $\chi^2(5, N = 1,218) = 89.517, p < 0.001$ , Cramér's  $V = 0.27$ ). In other words, participants tended to attribute more objectivity to claims that they had stronger opinions about than to claims they had weaker

opinions about, but they tended to attribute less objectivity to claims they recognized were widely disputed in society. Somewhat surprisingly, higher ratings of perceived disagreement about an issue were positively associated with participants' strength of opinion about the issue,  $\chi^2(10, N = 1,212) = 100.897$ ,  $p < 0.001$ , Cramér's  $V = 0.20$ .

## Discussion

Like Goodwin and Darley (2008) and Beebe and Sackris, Study 1 found that participants attribute more objectivity to some ethical claims than to some factual claims and that there is significant variation concerning the degree of objectivity attributed to different claims within each subcategory.<sup>4</sup> Thus, Study 1 reinforces the conclusion already established by Goodwin and Darley (2008) and Beebe and Sackris that the question of whether ordinary individuals are moral objectivists is not going to have a simple "Yes" or "No" answer.

## Study 2

### Method

#### **Participants**

A total of 195 undergraduate students (average age = 19, 47% female, 69% Anglo-American) from the University at Buffalo participated in Study 2 in exchange for extra credit in an introductory course.

#### **Materials and procedure**

The primary purpose of Study 1 was to construct a baseline of data with which the results of Studies 2 through 4 could be compared. These latter studies each introduce some kind of modification to the research materials used in Study 1 in order to see how folk metaethical judgments will be affected. The manipulation in Study 2 was simply a change in the order of the tasks participants were asked to complete.

As noted above, Study 1 followed Beebe and Sackris in having participants perform the following tasks in the following order:

Task 1: Participants indicated the extent to which they agreed or disagreed with a given claim.

Task 2: Participants answered the question “If someone disagrees with you about whether [the claim in question is true], is it possible for both of you to be correct or must at least one of you be mistaken?”

Task 3: Participants rated the extent to which they thought people in our society disagreed about whether the claim in question is true.

Thus, the last thing participants were asked to do was to consider about the extent of societal disagreement with respect to the claims. Given the negative association between perceived disagreement and objectivity attributions, it was hypothesized that if participants were directed to think about societal disagreement before completing Task 2, their attributions of metaethical objectivity would decrease. Disagreement was not hypothesized to have a similar effect on factual and taste claims.

## Results

As expected, the overall proportion of objectivity attributions in the ethical subcategory was lower in Study 2 (0.29) than in Study 1 (0.34). This difference was significant,  $\chi^2(1, N = 1045) = 4.015, p < 0.05$ , Cramér's  $V = 0.06$ . There were no significant differences in the factual and taste subcategories. Thus, it appears that making disagreement about ethical issues salient to participants can have a modest effect on the metaethical judgments they make. The fact that this result was obtained in the ethical domain but not in the factual domain is consistent with the widespread view among philosophers that ethical disagreement—because of its seemingly intractability—poses a significant challenge to the objectivity of ethical claims in a way that disagreement about factual matters fails to do for the objectivity of factual claims.<sup>5</sup>

## Discussion

The findings of Study 2 are consistent not only with the correlational data obtained by Goodwin and Darley (2008) and Beebe and Sackris but also with

the experimental data obtained by Goodwin and Darley (2012). The latter manipulated participants' perceived consensus about ethical issues by giving them bogus information about the percentage of students from the same institution who agreed with them. Participants who were told that a majority of their peers agreed with them about some ethical statement were more likely to think there was a correct answer as to whether or not the statement was true than participants who were told that significantly fewer of their peers agreed with them. These studies show that perceived disagreement or consensus can be a causal and not a merely correlational factor in folk metaethical decision-making.

### Study 3

Various studies of folk intuitions about moral responsibility have shown that individuals hold agents more responsible for their actions when the situations of those agents are described concretely than when they are described abstractly. Nichols and Knobe (2007), for example, obtained significantly higher ratings of moral responsibility for "Bill," who was attracted to his secretary and killed his wife and three children in order to be with her, than for "a person" whose actions were left unspecified. Small and Loewenstein (2003, 2005) showed that the subtlest change in the concreteness of the representation of an individual can lead to surprising differences in judgments or decisions regarding them. When their participants were given the opportunity to punish randomly selected defectors in an economic game, participants selected significantly harsher punishments for anonymous defectors whose numbers had just been chosen than for anonymous defectors whose numbers were about to be chosen. Because increased concreteness appears to heighten or intensify the engagement of cognitive and affective processes associated with attributions of blame and responsibility and to lead participants to treat the actions of concrete individuals as more serious than abstractly represented ones,<sup>6</sup> it was hypothesized that increasing the concreteness of those with whom participants were asked to imagine they disagreed would lead participants to take the disagreements more seriously and to increase attributions of metaethical objectivity.



## Method

### Participants

A total of 108 undergraduate students (average age = 19, 59% female, 66% Anglo-American) from the University at Buffalo participated in Study 3 in exchange for extra credit in an introductory course.

### Materials and procedure

In Beebe and Sackris' materials, which serve as the basis for Studies 1 and 2, each participant was asked "If someone disagrees with you about whether [one of these claims is true], is it possible for both of you to be correct or must one of you be mistaken?" In Study 3, this unspecified "someone" was replaced with "Joelle P., a junior nursing major at UB," "Mike G., a freshman computer science major at UB," or some other student from the participant's university, whose first name, last initial, class, and major were specified.

In between completing Tasks 1 and 3 (which were described above) for 8 of the 24 claims found in Table 9.1, each participant completed a modified version of Task 2 such as the following:

Madeline B., a senior biology major at UB, believes it is permissible to lie on behalf of a friend who is accused of murder. If you disagree with Madeline B., is it possible for both of you to be correct or must one of you be mistaken?

\_\_\_ It is possible for both of you to be correct.

\_\_\_ At least one of you must be mistaken.

**[If you agree with Madeline B., please skip to the next question.]**

## Results

In accord with my expectations, having more concrete parties with which to disagree resulted in a significantly greater overall proportion of objectivity attributions to ethical claims in Study 3 (0.43) than in Study 1 (0.34),  $\chi^2(1, N = 826) = 5.399, p < 0.05$ , Cramér's  $V = 0.08$ . The proportions were

numerically higher for eight of the ten ethical claims. Having more concrete parties in Study 3 did not, however, result in any significant difference in the objectivity attributed to factual or taste claims.

## Discussion

The results from Study 3 are consistent with those obtained by Sarkissian et al. (2011), who found that strong objectivity ratings were obtained when participants were asked to consider disagreeing with a concretely presented individual from their same culture (vs. a concretely presented individual from a different culture). The fact that the concreteness of the disagreeing parties used in Study 3 led to increased metaethical objectivity attributions may also explain why the objectivity ratings obtained in Study 1 fell below those obtained by Goodwin and Darley (2008), even though both used samples of university students. The Task 2 objectivity question in Study 1 asked participants to consider a situation of hypothetical disagreement (“If someone disagrees with you . . .”). Goodwin and Darley (2008, 1344), however, instructed participants, “We have done prior psychological testing with these statements, and we have a body of data concerning them. None of the statements have produced 100% agreement or disagreement.” Each of Goodwin and Darley’s objectivity questions then reiterated that some individuals who had been previously tested disagreed with participants about the relevant issue. Goodwin and Darley thus constructed situations of disagreement that were more concrete than those in Studies 1 and 2 by having (allegedly) actual rather than merely hypothetical individuals who disagreed with participants.

## Study 4

Study 3 made the parties with whom experimental participants were asked to consider disagreeing concrete by providing them with given names, surname initials, academic classes, and majors. In Study 4, the disagreeing parties were made concrete by having pictures of their faces shown. Faces (and parts of faces) have been shown to have a variety of effects on morally relevant

behavior. For example, Bateson et al. (2006) found that academics paid 276 percent more for the tea they took from a departmental tea station when an image of eyes was displayed by the station than when an image of flowers was displayed. Rezlescu, Duchaine, Olivola, and Chater (2012) found that unfakeable facial features associated with trustworthiness attracted 42 percent greater investment in an economic game that required trust.<sup>7</sup>

## Method


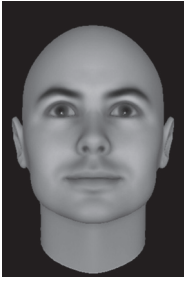

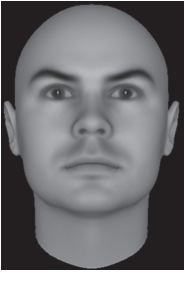

### Participants

A total of 360 participants (average age = 32, 38% female, 82% Anglo-American) were recruited through Amazon's Mechanical Turk ([www.mturk.com](http://www.mturk.com)) and were directed to complete online questionnaires hosted at [vovici.com](http://vovici.com).<sup>8</sup>

### Materials and procedure

Combining behavioral studies and computer modeling, Oosterhof and Todorov (2008) found that individuals make surprisingly consistent judgments about socially relevant traits of individuals on the basis of differences in their facial characteristics. They claim that the two most important dimensions of face evaluation are trustworthiness/untrustworthiness and dominance/submissiveness. Judgments concerning the first dimension are reliably associated with judgments about whether an individual should be approached or avoided and with attributions of happiness or anger. Judgments concerning dominance or submissiveness were found to be reliably associated with judgments of the maturity, masculinity, and physical strength of an individual. Both untrustworthy and dominant faces were associated with potential threat.<sup>9</sup> By exaggerating features specific to one of these evaluative dimensions, Oosterhof and Todorov (2008) created the set of faces represented in Table 9.2.<sup>10</sup> Each of the non-neutral faces was plus or minus three standard deviations from the mean along the relevant dimension. The faces in Table 9.2 were used in Study 4, along with a control condition in which no face was displayed.

**Table 9.2** Faces used in Study 4

		
	Dominant	
		
Trustworthy	Neutral	Untrustworthy
		
	Submissive	

Claims (12), (13), and (14) from Table 9.1—concerning embryonic stem cell research, lying for a friend accused of murder, and treating a national flag disrespectfully—were selected for use in Study 4. The degrees of objectivity attributed to them in Studies 1 through 3 fell in the middle range, suggesting that judgments about them could be more easily manipulated than judgments near the floor or ceiling. The first screen contained one of the pictures from Table 9.2, along with the following (Task 1) question:

Mark (pictured above<sup>11</sup>) believes that [statement (12), (13), or (14) is true].  
Please indicate whether you agree or disagree with Mark's belief.

Agree

Disagree

If participants selected "Agree" in response to one Task 1 question, they would be directed to answer the Task 1 question for one of the other target claims. However, if participants selected "Disagree," they were directed to answer the following (Task 2) metaethical question about their disagreement before moving on to the next Task 1 question:

You disagree with Mark about whether [the target claim is true]. Is it possible for both of you to be correct about this issue or must at least one of you be mistaken?

It is possible for both of you to be correct.

At least one of you must be mistaken.

Each screen that presented the metaethical question included the same picture (if any) that participants saw at the top of their Task 1 question. Each participant was presented with claims (12), (13), and (14) in counterbalanced order. The same picture (if any) of Mark appeared above each of these questions. Thus, no participant saw more than one version of Mark's face.

It was hypothesized that the five facial conditions would engage online processes of social cognition to a greater degree than the control condition and that this would result in higher attributions of metaethical objectivity. On the basis of Oosterhof and Todorov's (2008) finding that untrustworthy and dominant faces were associated with potential threat, it was also hypothesized that untrustworthy and dominant faces would elicit lower objectivity attributions than their dimensional pairs, since participants might be more tentative or anxious about disagreeing with potentially threatening interlocutors.

## Results

The proportion of objective attributions was significantly higher in the Neutral (0.65), Dominant (0.61), Submissive (0.60), Trustworthy (0.67),

and Untrustworthy (0.66) face conditions than it was in the No Face (0.46) condition. The proportions of objectivity attributions in the five face conditions did not differ significantly from each other.

## Discussion

Thus, it appears that having a face—any face, perhaps—makes the situation of moral judgment more concrete and engages moral cognitive processes in a way that increases attributions of objectivity. Because there were no significant differences between the Trustworthy and Untrustworthy conditions and the Dominant and Submissive face conditions, the second hypothesis concerning the effect of specific kinds of faces on folk metaethical intuitions failed to receive confirmation from Study 4.<sup>12</sup>

## Study 5

### Method

#### **Participants**

Using a between-subjects design, 160 participants (average age = 34, 38% female, 80% Anglo-American) were recruited through Amazon's Mechanical Turk and were directed to complete online questionnaires hosted at [vovici.com](http://vovici.com).<sup>13</sup>

#### **Materials and procedure**

A final study was constructed to see if the moral valence of the actions that disagreeing parties were described as performing would have an effect on folk metaethical judgments. Building upon work on the well-known Knobe effect in experimental philosophy,<sup>14</sup> in which individuals' folk psychological attributions have been shown to depend in surprising ways upon the goodness or badness of agents' actions, the following four descriptions were constructed:

1. The CEO of a company that helps and preserves the environment believes that it is morally wrong to harm the environment.
2. The CEO of a company that helps and preserves the environment believes that it is not morally wrong to harm the environment.
3. The CEO of a company that harms and pollutes the environment believes that it is morally wrong to harm the environment.
4. The CEO of a company that harms and pollutes the environment believes that it is not morally wrong to harm the environment.

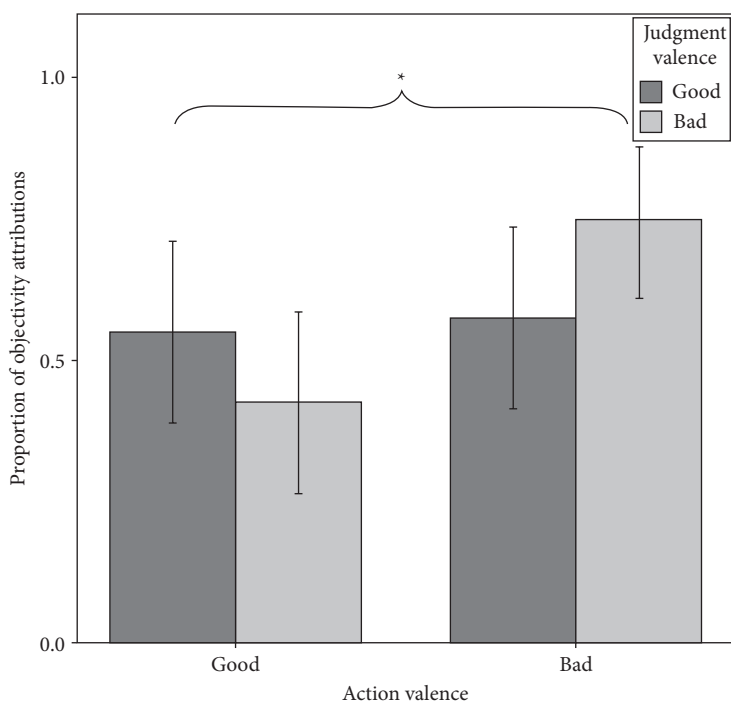
In (1) and (2), the CEO is depicted doing something morally good, namely, helping and preserving the environment, whereas the CEO's actions in (3) and (4) are morally bad. In (1) and (3), the CEO is described as having a morally good belief about the environment, namely, that it should not be harmed; in (2) and (4), the CEO has the corresponding morally bad belief. The crossing of good and bad actions with good and bad beliefs results in the actions and beliefs of the CEO being congruent in (1) and (4) and incongruent in (2) and (3).

Participants were first asked to indicate in a forced-choice format whether they agreed or disagreed with the CEO's belief. They were then asked, "If someone disagreed with the CEO about whether it is morally wrong to harm the environment, would it be possible for both of them to be correct or must at least one of them be mistaken?" Participants were then directed to choose between "It is possible for both of them to be correct" and "At least one of them must be mistaken."

## Results

The results of Study 5 are summarized in Figure 9.2.

Participants were more inclined to attribute objectivity to the ethical beliefs in question when the protagonist performed morally bad actions than when he performed morally good ones. This difference was significant,  $\chi^2(1, N = 160) = 5.013, p < 0.05$ , Cramér's  $V = 0.18$ . Neither belief valence nor the congruence between action and belief significantly affected folk metaethical judgments. However, it is noteworthy that the highest proportion of objectivity attributions was obtained in the "double bad" (i.e., Bad Action/



**Figure 9.2** Mean objectivity attributions in the Good Action/Good Belief (0.55), Good Action/Bad Belief (0.43), Bad Action/Good Belief (0.57), and Bad Action/Bad Belief (0.75) conditions of Study 5.

Bad Belief) condition, since it is badness (rather than goodness or neutrality) that has been shown to be the driving force behind the various forms of the Knobe effect.

## Discussion

As with other findings from the Knobe effect literature, the moral valence of a protagonist's action significantly affected participants' responses to probe questions. However, unlike other results in this literature, the responses in question were not folk psychological ascriptions. They were second-order attributions of objectivity to ethical beliefs held by the protagonist. These results provide further evidence that individuals' assessments of metaethical disagreements are significantly affected by a variety of factors in the situation of disagreement.



## General discussion

The foregoing studies show (i) that making disagreement salient to participants before asking them to make metaethical judgments can decrease objectivist responses, (ii) that increasing the concreteness of the situation of disagreement participants are directed to consider can increase objectivist responses, and (iii) that the moral valence of the actions performed by agents whose ethical beliefs participants are asked to consider affected attributions of objectivity to those beliefs. Because philosophical discussion—whether in the classroom or at professional conferences—often takes place in a somewhat rarefied atmosphere of abstractions, philosophers should be aware that intuitive agreement or disagreement with their metaethical claims can be affected by the very abstractness of those situations and that the amount of agreement or disagreement they encounter might be different in other situations. In spite of the fact that an increasing number of philosophers are familiar with the Knobe effect and its seemingly unlimited range of applicability, many philosophers continue to give little thought either to the moral valence of the actions depicted in their favored thought experiments and/or to the consequences this might have.

An important question raised by the studies reported above concerns the coherence of folk metaethical commitments. Most philosophers assume that the correct semantics for ordinary ethical judgments must show them to be uniformly objective or subjective.<sup>15</sup> Yet, Studies 2 through 5—in addition to work by Goodwin and Darley (2008), Beebe and Sackris (under review), and Sarkissian et al. (2011)—reveal that there are several kinds of variation in folk metaethical judgments. The lack of uniformity in the objectivity attributed to ethical claims might make us wonder how well ordinary individuals grasp the ideas of objectivism and subjectivism (and perhaps the related ideas of relativism and universalism). It might also lead us to question their reasoning abilities. Goodwin and Darley (2008, 1358, 1359), for example, suggest that “individuals were not particularly consistent in their meta-ethical positions about various ethical beliefs” and that “requirements of judgmental consistency across ethical scenarios are not considered.” However, this attribution of inconsistency seems both uncharitable and unwarranted.

Why should we believe that the ordinary use of ethical terms requires a semantics that assumes uniform objectivity or subjectivity? Because

armchair philosophers who have gathered no empirical evidence about the actual practice of using ethical terms say so? It seems that the practice should dictate the semantics, and not the other way around. If we find variability in the practice, we should look for semantic theories that can accommodate such variation. Furthermore, a variety of semantic theories can do so. For example, in Beebe (2010) I offer a relevant alternatives account of ethical judgments that borrows heavily from the semantic machinery of the epistemic contextualists (e.g., Lewis 1996; DeRose 2011). I argue that treating ethical terms as context-sensitive yields a better interpretation of ordinary normative and metaethical judgments than interpretations that treat them as context-invariant. Without delving into the details of the view, the upshot for present purposes is that attributions of inconsistency or incoherence to folk metaethical practice are premature when there are more charitable interpretive options available.

Another important issue raised by the above studies concerns my hypothesis that it is concreteness that is driving the effects observed in Studies 3 and 4. An alternative possibility is that when undergraduates at the University at Buffalo are told that Madeline B., a senior biology major at UB, believes that some action is morally permissible, it may be Madeline's cultural proximity or group affiliation that leads participants to make more objectivist judgments. Signaling that someone from the same university believes that *p* may suggest to participants that they should believe it as well, if they are to remain members in good standing in the relevant group. And it is of course possible that some other kind of social influence might be operative as well. Further research is required to determine whether it is concreteness or other social factors that push individuals in the direction of greater objectivism.<sup>16</sup>

The studies reported above show that not only are there differences in folk metaethical judgments that track the content of ethical claims (Goodwin and Darley 2008; Beebe and Sackris, under review; Study 1), how contested they are (Goodwin and Darley 2012; Study 2), and the cultural distance between disagreeing parties (Sarkissian et al. 2011); there are also differences that track the goodness or badness of disagreeing parties (Study 5) and possibly their concreteness as well (Studies 3 and 4). It is hoped that the present research sheds useful light on the multi-dimensional variation that characterizes the folk metaethical landscape.

## Notes

- \* James R. Beebe, Department of Philosophy, University at Buffalo. The author would like to thank Mike Giblin and Anthony Vu for assistance in data collection and Hagop Sarkissian and Jen Cole Wright for helpful comments on a previous version of this manuscript. Correspondence should be addressed to James R. Beebe, 135 Park Hall, Department of Philosophy, University at Buffalo, Buffalo, NY 14260. Email: jbeebe2@buffalo.edu.
- 1 However, cf. Beebe and Sackris (under review, sec. 1) for critical discussion of many of the measures of metaethical commitments that are employed in the published literature.
  - 2 Forty-one percent of participants in one study and 38 percent of participants in another classified first trimester abortion as a personal rather than a moral issue; 89 percent and 73 percent of participants did the same for anonymously donating money to charity.
  - 3 All statistical tests reported in this chapter are chi-square tests of independence. On the conventional interpretation of Cramér's  $V$ , an effect size of 0.1 to 0.29 counts as small, one 0.3 to 0.49 counts as medium, and one 0.5 or larger counts as large.
  - 4 The only gender differences in the data were that females held slightly less strong opinions than males on factual matters ( $\chi^2(2, N = 303) = 6.124, p < 0.05$ , Cramér's  $V = 0.14$ ) and reported greater societal disagreement than males concerning matters of taste ( $\chi^2(2, N = 225) = 11.296, p < 0.05$ , Cramér's  $V = 0.22$ ).
  - 5 Cf. Sidgwick (1907/1981, 342), Mackie (1977, 36–8), Wong (1984), and Tersman (2006). Because salient disagreement impacted participants' second-order (metaethical) judgments in Study 2, a follow-up study was performed to see if salient disagreement might have a similar impact upon participants' first-order judgments—that is, upon the degree of agreement they expressed in response to various ethical claims in Task 1. Participants were directed to complete Task 3 immediately before Task 1, and it was hypothesized that salient disagreement would result in less confident Task 1 judgments. However, this manipulation failed to have a significant impact on participants' Task 1 judgments.
  - 6 Nahmias et al. (2007) found that this was especially true if wrongdoing is involved.
  - 7 Thanks to Mark Alfano for bringing this work to my attention.
  - 8 Participants were required to reside in the United States and to have at least a 95 percent approval rating on more than 500 mturk tasks. Each participant was paid \$.30.

- 9 Untrustworthy faces were associated with potentially harmful intentions, while dominant faces were associated with the capacity to cause harm.
- 10 Oosterhof and Todorov constructed the faces using FaceGen Modeller 3.2 (Singular Inversions 2007).
- 11 This phrase of course did not appear in the No Face condition.
- 12 It may be that how people respond to these kinds of faces depends upon whether they themselves are dominant, submissive, etc. The method of the present study did not allow this factor to be explored. Thanks to Hagop Sarkissian for raising this point.
- 13 Participants were required to reside in the United States and to have at least a 95 percent approval rating on more than 500 mturk tasks. Each participant was paid \$.30.
- 14 See Alfano et al. (2012) for an overview of this literature.
- 15 See Sinnott-Armstrong (2009) for further discussion of this point.
- 16 Thanks to Hagop Sarkissian and Jen Cole Wright for pressing these points.

## References

- Alfano, M. (2010). Social cues in the public good game. Presentation at *KEEL 2010 Conference: How and Why Economists and Philosophers Do Experiments*, Kyoto Sangyo University, Kyoto, Japan.
- Alfano, M., Beebe, J. R., and Robinson, B. (2012). The centrality of belief and reflection in Knobe effect cases: A unified account of the data. *The Monist*, 95, 264–89.
- Bateson, M., Nettle, D., and Roberts, G. (2006). Cues of being watched enhance cooperation in a real-world setting. *Biology Letters*, 12, 412–14.
- Beebe, J. R. (2010). Moral relativism in context. *Noûs*, 44, 691–724.
- Beebe, J. R., and Sackris, D. (Under review). Moral objectivism across the lifespan.
- DeRose, K. (2011). *The Case for Contextualism: Knowledge, Skepticism, and Context*, vol. 1. Oxford: Oxford University Press.
- Goodwin, G. P., and Darley, J. M. (2008). The psychology of meta-ethics: Exploring objectivism. *Cognition*, 106, 1339–66.
- (2012). Why are some moral beliefs perceived to be more objective than others? *Journal of Experimental Social Psychology*, 48, 250–6.
- Lewis, D. K. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74, 549–67.
- Nahmias, E., Coates, D., and Kvaran, T. (2007). Free will, moral responsibility, and mechanism: Experiments on folk intuitions. *Midwest Studies in Philosophy*, 31, 214–42.

- Nichols, S., and Knobe J. (2007). Moral responsibility and determinism: The cognitive science of folk intuitions. *Noûs*, 41, 663–85.
- Oosterhof, N., and Todorov, A. (2008). The functional basis of face evaluations. *Proceedings of the National Academy of Sciences of the USA*, 105, 11087–92.
- Rezlescu, C., Duchaine, B., Olivola, C. Y., and Chater, N. (2012). Unfakeable facial configurations affect strategic choices in trust games with or without information about past behavior. *PLoS ONE*, 7, e34293.
- Sarkissian, H., Parks, J., Tien, D., Wright, J. C., and Knobe. J. (2011). Folk moral relativism. *Mind & Language*, 26, 482–505.
- Sidgwick, H. (1907/1981). *The Methods of Ethics*. Indianapolis: Hackett.
- Sinnott-Armstrong, W. (2009). Mixed-up meta-ethics. *Philosophical Issues*, 19, 235–56.
- Small, D. A., and Loewenstein, G. (2003). Helping a victim or helping *the* victim: Altruism and identifiability. *Journal of Risk and Uncertainty*, 26, 5–16.
- (2005). The devil you know: The effects of identifiability on punitiveness. *Journal of Behavioral Decision Making*, 18, 311–8.
- Tersman, F. (2006). *Moral Disagreement*. Cambridge: Cambridge University Press.
- Wong, D. B. (1984). *Moral Relativity*. Berkeley: University of California Press.
- Wright, J. C., Grandjean, P. T., and McWhite, C. B. (forthcoming). The meta-ethical grounding of our moral beliefs: Evidence for metaethical pluralism. *Philosophical Psychology*.

AQ: Please provide publication year and complete details for the reference “Wright et al. (forthcoming).”