

Exploring Metaethical Commitments: Moral Objectivity and Moral Progress

Kevin Uttich, George Tsai, and Tania Lombrozo*

People have beliefs not only about specific moral issues, such as the permissibility of slavery, but also about the *nature* of moral beliefs. These beliefs, or “metaethical commitments,” have been the subject of recent work in psychology and experimental philosophy. One issue of study has been whether people view moral beliefs in more objectivist or relativist terms (i.e., as more like factual beliefs or more like personal preferences).

In this chapter, we briefly review previous research on folk moral objectivism. We then present the results of an experiment that compares two different approaches to measuring moral objectivism (those of Goodwin and Darley 2008, and Sarkissian et al. 2011) and consider the relationship between objectivism and two additional metaethical beliefs: belief in moral progress and belief in a just world. By examining the relationships between different metaethical commitments, we can better understand the extent to which such commitments are (or are not) systematic and coherent, shedding light on the psychological complexity of an important area of moral belief and experience.

To preview our results, we find that different metaethical beliefs are reliably but weakly associated, with different measures of moral objectivism generating distinct patterns of association with belief in moral progress and belief in a just world. We highlight some of the challenges in reliably measuring metaethical commitments and suggest that the distinctions that have been useful in differentiating philosophical positions may be a poor guide to folk moral judgment.

Recent work on metaethical beliefs

Moral objectivity

Moral objectivity is a complex idea with multiple variants and diverse proponents (for useful discussions see Goodwin and Darley 2010; Knobe et al. 2012; Sinnott-Armstrong 2009). For our purposes, to accept moral objectivism is to believe that some moral claims are true in a way that does *not* depend on people's decisions, feelings, beliefs, or practices. Thus, to reject the objectivity of moral claims one can either deny that moral claims have a truth value or allow that moral claims can be true, but in a way that does depend on decisions, feelings, beliefs, or practices (e.g., Harman 1975; Sinnott-Armstrong 2009). Non-cognitivism is typically an instance of the former position, and cultural or moral relativism of the latter.

Recently, there have been a few attempts to examine empirically what people believe about moral objectivity (Goodwin and Darley 2008, 2010; Forsyth 1980; Nichols 2004; Sarkissian et al. 2011; see Knobe et al. 2012, for review). Goodwin and Darley (2008) asked participants to rate their agreement with statements that were factual, ethical, social-conventional, or about personal taste, and then asked them whether these statements were true, false, or "an opinion or attitude." For example, one of the ethical statements was "Robbing a bank in order to pay for an expensive holiday is a morally bad action," while one of the social-conventional statements was that "Wearing pajamas and bath robe to a seminar meeting is wrong behavior." Responding that these were either true or false was considered a more objectivist response than selecting "an opinion or attitude." Participants were later asked whether the fact that someone disagreed with them about a given statement meant that the other person was wrong, that neither person was wrong, that they themselves were wrong, or something else entirely. On this measure, responding that one of the two people must be wrong was taken as a more objectivist response.

Using a composite of these two measures, Goodwin and Darley found evidence that people treat statements of ethical beliefs as more objective than either social conventions or taste. They also found a great deal of variation in objectivism across both ethical statements and individuals. Strongly held ethical beliefs were seen as more objective than beliefs that people did not hold

as strongly, and those who said they grounded their ethical beliefs in religion, moral self-identity, or the pragmatic consequences of failing to observe norms were more likely to be objectivist about ethical statements. Subsequent work has suggested that variation in objectivist beliefs is not an artifact of variation concerning which issues participants themselves take to be moral, nor of misunderstanding moral objectivism (Wright et al. 2012).

More recently, Sarkissian et al. (2011) have argued that relativist beliefs are more prevalent than suggested by Goodwin and colleagues, but that these beliefs are only observed when participants are comparing judgments made by agents who differ from each other in important ways. In their studies, participants were presented with two agents who disagreed about a moral claim and were asked whether one of them must be wrong. For example, participants were asked to imagine a race of extraterrestrial beings called “Pentars” who “have a very different sort of psychology from human beings.” Participants were then presented with a hypothetical case in which a classmate and a Pentar had differing views on a moral case, and were asked to rate their agreement with the statement that “at least one of them must be wrong.” Participants provided more objectivist answers (“one of them must be wrong”) when comparing judgments made by agents from the same culture, but more relativist answers (denying that “at least one of them must be wrong”) when comparing judgments made by agents from different planets (i.e., a human and a Pentar). Sarkissian et al. argue that engaging with radically different perspectives leads people to moral relativism.

What are the implications of this research? On the one hand, the findings from Goodwin and Darley (2008) and Sarkissian et al. (2011) suggest that metaethical beliefs are not particularly developed or unquestionably coherent. They certainly challenge the idea that those without philosophical expertise can be neatly classified as “moral objectivists” versus “moral relativists.” Instead, judgments vary considerably depending on the moral claim in question and the way in which objectivism is assessed—in particular, whether a case of disagreement involves similar or dissimilar agents.

On the other hand, a growing body of research suggests that moral objectivism is systematically related to aspects of cognition and behavior that go beyond metaethical beliefs. For example, Goodwin and Darley (2012) found that moral claims were judged more objective when there was greater

perceived consensus. They also found that participants judged those who held opposing beliefs as less moral and harder to imagine interacting with when disagreement concerned a claim that was considered objective (see also Wright et al. in press). Finally, Young and Durwin (2013) found that participants primed to think in more objective terms were more likely to give to charity.

These findings, among others, suggest that despite intra- and interpersonal variation in judgments, moral objectivism relates to factual beliefs (e.g., about consensus), attitudes (e.g., tolerance of others), and decisions (e.g., about whether to give to charity). We aim here to better understand the ways in which metaethical beliefs are and are not systematic and coherent by considering the relationship between three different metaethical beliefs: belief in moral objectivism, belief in moral progress, and belief in a just world.

Moral progress

A belief in moral progress is a commitment to the idea that history tends toward moral improvement over time. This notion, which postulates a certain directionality in human history, can be contrasted with the notion of mere moral change. Although moral progress has been defended by philosophers in the history of philosophy, notably Marx and Hegel, the notion also finds expression in people's ordinary thinking. For example, Martin Luther King famously proclaimed, "the arc of the moral universe is long but it bends towards justice" (King 1986).

It is worth noting that a belief in a historical tendency toward moral progress can be consistently held while maintaining that moral progress can be imperceptible, occurring over long stretches of time. Sometimes moral improvement can be dramatic and rapid, while at other times not. Thus, belief in a tendency toward moral progress does not require committing to a particular rate of moral progress. Additionally, to hold that there is a basic tendency toward moral progress in human history is also compatible with allowing that these tendencies do not inevitably or necessarily prevail. Believing in some tendency need not require belief in inevitability. For example, one could believe that 6-year-old children tend to grow physically larger (e.g., that a child at 14 years of age will be larger than that very same child at age 6) without claiming that they inevitably or necessarily get physically larger (serious illness

or death could prevent their continuing to grow in size). Likewise, in the case of moral progress, one could still allow that there could be exogenous forces such as environmental and biological catastrophes or foreign invasions that prevent the historical development toward moral progress.

One reason to focus on moral progress is that the notion is commonly invoked, reflecting ideas in the broader culture. There is therefore reason to suspect that people have commitments concerning its truth, and it is natural to ask, with philosopher Joshua Cohen (1997), “Do [ideas of moral progress] withstand reflective examination, or are they simply collages of empirical rumination and reified hope, held together by rhetorical flourish?” (p. 93). In particular, we might ask whether moral progress typically involves a commitment to moral objectivism, as objective norms might be thought to causally contribute to progress or simply provide a metric against which progress can be assessed.

It is also important to note that the notion of moral progress does not merely contain metaethical content but also a kind of descriptive content: to believe in moral progress involves believing something about the nature of human history and the character of the social world. This suggests that our metaethical beliefs, including beliefs about moral objectivity, do not stand alone, compartmentalized from other classes of beliefs. Not only might they be rationally and causally related to each other, in some cases these beliefs are inseparable, expressing a union between the ethical and the descriptive. Thus, a second reason for our interest in considering moral progress in tandem with moral objectivity is that it may reveal important connections between different types of metaethical beliefs as well as connections between metaethical beliefs and other beliefs (such as descriptive beliefs about consensus, or explanatory beliefs about social phenomena).

Belief in a just world

While previous research has not (to our knowledge) investigated beliefs about moral progress directly, there is a large body of research on a related but distinct metaethical belief, “belief in a just world” (e.g., Lerner 1980; Furnham 2003). Belief in a just world refers to the idea that good things happen to good people while bad things happen to bad people. The belief that people experience consequences that correspond to the moral nature of their actions or character

is potentially consistent with belief in moral progress, although the relationship is complex. For example, it is not obvious that the world is morally improved when a criminal experiences a string of bad luck, unless retribution or the deterrence of future criminal activity is itself the moral payoff. Nonetheless, we focus on belief in a just world as a third metaethical belief for two reasons. First, doing so allows us to examine empirically whether belief in a just world is in fact related to belief in moral progress, and thus relate our novel measures to existing research. Second, investigating a third metaethical commitment can help us differentiate two possibilities: that relationships between metaethical commitments are relatively selective, such that (for example) moral objectivity and moral progress might be related but have no association with belief in a just world, or alternatively, that the relationship reflects a single and more general tendency, such that individuals with strong metaethical commitments of one kind will typically have strong metaethical commitments of all kinds.

Method

We present a subset of results from a larger experiment investigating people's beliefs about moral objectivity using modified versions of both the Goodwin and Darley's (2008) and Sarkissian et al.'s (2011) measures, as well as people's beliefs about moral progress and belief in a just world. We also solicited explanations for social changes to investigate the relationship between metaethical beliefs and ethical explanations. In the present chapter, we focus on the relationships between different metaethical beliefs. In ongoing work, we consider the relationship between these beliefs and explanations (Uttich et al. in prep).

Participants

Three hundred and eighty-four participants (223 female; mean age = 33) were recruited from Amazon Mechanical Turk, an online crowd-sourcing platform. Participants received a small payment for their participation. All participants identified themselves as being from the United States and as fluent speakers of English.

Materials and procedure

We report a subset of a larger set of experiments concerning the relationship between metaethical beliefs and the use of explanations that cite ethical norms. The full experiment consisted of four main parts: (1) explanation solicitation, (2) moral objectivity measures, (3) moral progress measures and general belief in a just world measure (GBJW), and (4) baseline check on beliefs about the morality of social changes. The ordering of the parts was partially counterbalanced, as detailed below.

Explanation solicitation

In the full experiment, participants were presented with a description of a social change and asked to explain it in a few sentences (e.g., “Why was slavery abolished?”). The changes included the abolition of slavery, women’s suffrage, and the potential legalization of same-sex marriage. Given our present focus on the relationship between different metaethical beliefs, we do not report findings concerning explanation here (see Uttich et al. in prep).

Moral objectivity: Disagreement measure

Participants’ views concerning moral objectivity were examined in two different ways. The first involved an adaptation of the “disagreement” method used by both Goodwin and Darley (2008) and Sarkissian et al. (2011). Participants read vignettes where either a person similar to themselves (i.e., from their same time and culture) or a person from another time period (e.g., the eighteenth century) disagreed with an imagined friend of the participant about whether a social fact was morally problematic. The relevant social fact was always matched with that for which participants had been asked to provide an explanation. An example from the slavery condition involving the current time and place is presented below:

Imagine a person named Allison, a fairly ordinary student from your town who enjoys watching sports and hanging out with friends. Consider Allison’s views concerning the moral status of the following social institution: Slavery.

Allison thinks that slavery is not morally wrong.

This scenario was matched with one involving a judgment from a different time or place:

Imagine the social world of the United States in the eighteenth century. Most people in this time and place view slavery as morally acceptable. The existence of slavery is seen by many as part of the natural social order, slavery is permitted by the law and the slave trade is at its peak, and someone who owns many slaves is esteemed as admirable.

An individual, Jessica, from this society (eighteenth-century United States), regards slavery as not morally wrong.¹

In both cases, participants were then presented with a friend who disagreed:

Imagine that one of your friends thinks that slavery is morally wrong. Given that these individuals (Allison [Jessica] and your friend) 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 friend and Allison [Jessica] have different judgments about this case, at least one of them must be wrong.”

Participants rated their agreement with this statement on a 1–7-point scale with 1 corresponding to “definitely disagree,” 7 to “definitely agree,” and 4 to “neither agree nor disagree.” Each participant saw one current and one historical vignette, with order counterbalanced across participants.

Moral objectivity: Truth-value measure

Participants’ beliefs about moral objectivity were also examined using a method adapted from Goodwin and Darley (2008). Participants were asked whether statements about the moral permissibility of the social facts are true, false, or an opinion. The question prompt was adjusted from the original multiple-choice format used by Goodwin and Darley to a 7-point Likert scale to make it more comparable to the disagreement measure. Thus, participants rated their agreement with statements concerning the moral permissibility of a social practice (e.g., “slavery is not morally wrong”) on a 1–7-point scale with 1 being “is best described as true or false,” 7 “is best described as an opinion,” and 4 “is equally well described as either true/false or as an

opinion.” Participants answered questions concerning moral beliefs for all three historical facts (demise of slavery, women’s suffrage, legalization of same-sex marriage), with the historical fact for which they rated explanations presented first.

Moral progress and belief in a just world measures

Participants rated their agreement with eighteen statements intended to measure their metaethical commitments concerning moral progress and belief in a just world. Twelve items were constructed to measure participants’ beliefs in moral progress. The statements examined two dimensions of this belief: whether they concerned something *concrete* (i.e., moral progress with respect to a *particular* social practice or area of social life) or *abstract* (i.e., moral progress in *general*), and whether progress was described as a *tendency* or as *inevitable*. There were three questions for each of the four possible combinations (e.g., three concrete questions about tendency, three abstract questions about tendency, and so on). Participants also evaluated six statements concerning belief in a just world, taken from the GBJW (Dalbert et al. 1987). All eighteen statements are included in Table 10.1.

Participants rated the statements on a 1–7-point scale with 1 being “definitely disagree,” 7 “definitely agree,” and 4 “neither agree nor disagree.” The order of all moral progress and GBJW statements was randomized.

Baseline check

Participants were also asked for their personal views on whether the three social changes were good or bad. For example, for the slavery fact participants were presented with the following statement:

The demise of slavery was a good thing.

Participants rated their agreement with this statement on a 1–7-point scale with 1 being “definitely disagree,” 7 “definitely agree,” and 4 “neither agree nor disagree.” All three social facts were rated. The social fact related to the explanation for which each participant had been prompted was always presented first.

Table 10.1 Means and factor loadings for statements of moral progress and belief in a just world. Items with an asterisk were reverse coded.

Statement	Mean (SD)	Factor 1 ("Abstract progress")	Factor 2 ("GBJW")	Factor 3 ("Concrete inevitability")	Factor 4 ("Concrete progress")
<i>Abstract Tendency</i>					
Human history tends to move in the direction of a more perfect moral world.	3.84 (1.60)	0.757	0.212	- 0.001	0.189
As time goes on, humanity does NOT generally become more morally advanced.*	4.21 (1.68)	0.707	0.011	0.061	0.077
Over time we will move toward a more morally just world.	4.22 (1.60)	0.782	0.184	0.093	0.272
<i>Abstract Inevitability</i>					
The moral advancement of humanity is NOT inevitable.*	3.66 (1.76)	0.686	0.238	0.227	- 0.068
It is inevitable that on average our morals will be better than those of our distant ancestors.	3.94 (1.77)	0.743	0.173	0.274	0.091
An increase in moral justice in the world is inevitable.	4.05 (1.69)	0.729	0.340	0.227	0.072
<i>Concrete Tendency</i>					
Over time there is moral progress concerning slavery.	5.54 (1.45)	0.159	0.146	0.179	0.813
Over time there is moral progress concerning voting rights.	5.41 (1.35)	0.250	0.066	0.249	0.733
Over time there is moral progress concerning marriage rights.	4.79 (1.60)	0.586	- 0.001	0.290	0.354

(Continued)

Table 10.1 (Continued)

Statement	Mean (SD)	Factor 1 ("Abstract progress")	Factor 2 ("GBJW")	Factor 3 ("Concrete inevitability")	Factor 4 ("Concrete progress")
<i>Concrete Inevitability</i>					
The demise of slavery was inevitable.	5.14 (1.78)	0.262	0.093	0.652	0.162
The extension of the right to vote to women was inevitable.	5.43 (1.66)	0.097	0.186	0.776	0.258
The legalization of same-sex marriage is inevitable.	5.14 (1.82)	0.189	0.008	0.774	0.059
<i>General Belief in a Just World</i>					
I think basically the world is a just place.	3.77 (1.62)	0.170	0.563	0.116	- 0.109
I believe that, by and large, people get what they deserve.	4.16 (1.62)	0.093	0.705	- 0.017	0.037
I am confident that justice always prevails over injustice.	3.64 (1.75)	0.293	0.723	0.121	- 0.067
I am convinced that in the long run people will be compensated for injustices.	3.95 (1.75)	0.036	0.775	0.108	0.128
I firmly believe that injustices in all areas of life (e.g., professional, family, politic) are the exception rather than the rule.	3.82 (1.52)	0.082	0.620	0.055	0.122
I think people try to be fair when making important decisions.	4.69 (1.46)	0.159	0.607	- 0.006	0.271

Counterbalancing

Participants either provided an explanation first (part 1) and then completed the two moral objectivity measures (part 2) and the moral progress measures and GBJW measures (part 3), with the order of parts 2 and 3 counterbalanced, or they first completed the moral objectivity measures (part 2) and the moral progress and GBJW measures (part 3), with order counterbalanced, followed by explanations (part 1). Participants always completed the baseline check on social facts (part 4) last.

Results

We begin by reporting the data for each set of questions individually, and then consider the relationship between different metaethical commitments.

Individual measures

Baseline check measures

The baseline check confirmed our assumptions about participants' own attitudes toward the moral claims in question. The average ratings were 6.70 of 7 ($SD = 0.95$) for the demise of slavery, 6.63 ($SD = 1.00$) for women's suffrage, and 5.15 ($SD = 2.20$) for same-sex marriage.

Moral objectivism: Disagreement

The first measure of objectivism concerned participants' responses to disagreement between a friend and an individual in a current or historical period. Overall, participants provided higher ratings for the current scenario ($M = 4.84$, $SD = 1.95$) than for the historical scenario ($M = 4.62$, $SD = 1.97$), indicating greater objectivism in the latter case and consistent with Sarkissian et al.'s findings. To analyze the data statistically, we performed a repeated-measures ANOVA with time period (current vs. historical) as a within-subjects factor and social fact (slavery, women's suffrage, same-sex marriage) as a between-subject factor. This revealed two significant effects: a main effect of time period, $F(1,381) = 13.17$, $p < 0.01$, with more objectivist responses for the current vignette than for the historical vignette, and a main effect of social

fact, $F(2,381) = 36.35, p < 0.01$, with responses that were more objectivist for slavery ($M = 4.99, SD = 1.90$) and women's suffrage ($M = 4.90, SD = 1.69$) than for same-sex marriage ($M = 4.30, SD = 1.95$).

Because the correlation between participants' current (C) and historical (H) ratings was very high ($r = 0.817, p < 0.01$), we consider the average rating (C+H) for each participant ($M = 4.72, SD = 1.87$) in most subsequent analyses.

Moral objectivism: Truth value

Our second measure of moral objectivism was the "true, false, or opinion?" (TFO) measure adapted from Goodwin and Darley (2008). The average rating for the TFO measure was 4.31 ($SD = 2.15$), with lower scores indicating greater moral objectivism. This measure varied as a function of social fact, $F(2,382) = 53.65, p < 0.01$, with the most objectivist responses for slavery ($M = 3.71, SD = 2.58$), followed by women's suffrage ($M = 4.31, SD = 2.51$) and same-sex marriage ($M = 4.91, SD = 2.30$).² Nonetheless, participants' ratings across the three social facts were highly related ($\alpha = 0.82$). In subsequent analyses, we therefore focus on a given participant's average TFO rating across the three social facts. To facilitate the interpretation of correlations across measures, we report a *reversed* average (8 minus each participant's score) such that higher numbers correspond to greater objectivism, as in the C+H measure.

Moral progress and belief in a just world measures

We analyzed ratings across the 18 statements with a factor analysis employing principal components analysis as an extraction method and a varimax rotation. This analysis resulted in four factors with eigenvalues over one, accounting for a total of 59.4 percent of the variance. Table 10.1 reports the average rating for each statement as well as the factor loadings for each statement in the rotated component matrix, and suggests that the four factors can be characterized as follows: abstract progress (34.2% of variance), GBJW (11.5%), concrete inevitability (8.1%), and concrete tendency (5.6%). It's worth noting that beliefs about moral progress were indeed differentiated from GBJW, and that the dimension of abstract versus concrete appeared to be psychologically meaningful while the distinction between tendency and

inevitability emerged only for the concrete items, where participants may have been able to look back over time at the specific issues we considered to identify both general trends and temporary setbacks.³ In subsequent analyses we examine correlations between these four factors and our two measures of moral objectivism.

Relationships between metaethical commitment measures

Table 10.2 reports the correlations between our two measures of moral objectivism (C+H and TFO) as well as the four factors extracted from the factor analysis on moral progress items (abstract progress, GBJW, concrete inevitability, concrete tendency). There are several notable results.

First, while the correlation between the C+H ratings and the TFO ratings was significant ($r = 0.271$, $p < 0.01$), it was low enough to suggest that each

Table 10.2 Correlations between metaethical measures

	C+H	Current	Historical	C-H	Avg TFO
C+H: Current/ Historical disagreement	1	0.953**	0.954**	-0.020	0.271**
Current	0.953**	1	0.817**	0.285**	0.258**
Historical	0.954**	0.817**	1	-0.320**	0.258**
C-H: Difference score	-0.020	0.285**	-0.320**	1	-0.005
Avg TFO: True, False, or Opinion?	0.271**	0.258**	0.258**	-0.005	1
Factor 1: "Abstract progress"	0.063	0.059	0.061	-0.003	0.127*
Factor 2: "GBJW"	0.042	0.018	0.061	-0.072	-0.116*
Factor 3: "Concrete inevitability"	-0.040	-0.009	-0.066	0.094	-0.063
Factor 4: "Concrete tendency"	0.239**	0.258**	0.199**	0.093	-0.056

* <0.05 , ** <0.01 .

measure captures some important and unique variance in beliefs about moral objectivism, perhaps roughly capturing relativism and non-cognitivism, respectively. To further investigate the possible relationships between measures, we also considered whether TFO might be related to the *difference* between C and H ratings (C-H), which can be conceptualized as a measure of the extent to which a participant is influenced by sociocultural factors in evaluating the truth of a moral claim. One might therefore expect a significant negative correlation between TFO and C-H, but in fact the relationship was very close to zero. Coupled with the high correlation between judgments on the C and H questions, and the fact that C and H had very similar relationships to other variables, this suggests that varying the sociocultural context for a belief can indeed affect judgments concerning disagreement, but that the effect is more like a shift in the absolute value of participants' judgments than the recruitment or application of different moral commitments.

Second, while both the C+H and TFO ratings were related to moral progress and GBJW, they had unique profiles in terms of the specific factors with which they correlated. The C+H measure was correlated with the concrete tendency factor ($r = 0.239, p < 0.01$), while the TFO measure was positively correlated with the abstract progress factor ($r = 0.127, p < 0.05$) and negatively correlated with the GBJW factor ($r = -0.116, p < 0.05$). Although these correlations were small, they suggest systematic relationships between measures, and more surprisingly, *non-overlapping* relationships, providing further evidence that judgments of disagreement (C+H) and judgments concerning whether moral claims have a truth value (TFO) reflect different facets of folk metaethical commitments.

Finally, it's worth considering why C+H and TFO had these distinct profiles. We speculate that the dimension of concrete versus abstract evaluation can partially explain these results. Specifically, C+H and the concrete tendency factor were positively associated and involved particular moral claims (e.g., about slavery) rather than abstract claims, while TFO and the abstract progress factor were positively associated and involved judgments that were more explicitly metaethical in that they concerned the status of particular moral ideas (i.e., whether there is moral progress in general and whether particular claims have a truth value). However, this speculation does not explain why the C+H measure was not also associated with the concrete tendency factor, nor does it explain the negative association between TFO and the GBJW factor.

General discussion

Our results suggest that metaethical beliefs are varied and complex, with significant but modest relationships across different sets of beliefs. Our results also reinforce some of the conclusions from prior research. Like Goodwin and Darley (2008, 2012), we find significant variation in objectivism across individuals, and also that judgments reflect greater objectivism for some social facts (slavery) than for others (same-sex marriage), perhaps echoing their findings on the role of consensus, and also consistent with the strength of participants' attitudes concerning each social fact. Like Sarkissian et al. (2011), we find evidence that measures that highlight different perspectives seem to increase non-objectivist responses, as our "historical" vignette generated less objectivist responses than the matched "current" vignette, although the responses were strongly correlated. Our findings therefore support the need to consider the characteristics of both participants and measures in drawing conclusions about metaethical beliefs.

Beyond illuminating variation between individuals, our findings shed light on the coherence and variability of metaethical beliefs *within* individuals. Correlations between our measures of metaethical beliefs suggest two conclusions: that the metaethical concepts we investigate have some common elements, but also that there is only partial coherence in the corresponding beliefs. Our two separate measures of moral objectivity (C+H and TFO) were significantly correlated, but only weakly so. The correlation was weak despite modifications from Goodwin and Darley (2008) and Sarkissian et al. (2011) to make the measures more comparable: both involved judgments on 7-point scales and referred to the same moral claims. Analyses of the relationship between these two measures and the four factors concerning moral progress and GBJW suggest that moral objectivism is related to these ideas, but the two measures of objectivism had unique patterns of association. If participants have strong, stable, and consistent metaethical commitments, why might responses to metaethical questions be so weakly related?

We first consider methodological and conceptual answers to this question. One possibility is that we observe weak associations between metaethical commitments as an artifact of our methods of measurement. This idea is consistent with a suggestion by Sarkissian et al. (2011), who argue that when

forced to address radically different perspectives, people who *appeared* to have completely objectivist commitments reveal some underlying, relativist intuitions. It follows that methods for soliciting commitments might themselves account for substantial variance in responses. We suspect there is some truth to this idea, and our particular measures certainly have limitations. Nonetheless, it's worth repeating that while disagreements in current and historical contexts (C and H) involved different absolute ratings, they were very highly correlated and had matching patterns of association with our other measures. Moreover, the difference between these ratings (C-H)—that is the extent to which context shifted judgments—was not reliably associated with any measures. One interpretation is that people's *absolute* judgments may be quite easy to manipulate, but that the relationships between metaethical commitments, while weak, may be more stable.

Another possibility is that the metaethical commitments we investigated do not in fact correspond to coherent and unified sets of beliefs. Thus, participants' judgments could be inconsistent across measures because the philosophical constructs we aim to assess are themselves diverse or incoherent. For example, we expected a stronger relationship between moral objectivism and belief in moral progress, but such a relationship is not logically required—one can, for example, be a relativist and endorse moral progress, or an objectivist and deny it. We also expected our two measures of moral objectivism to be more strongly associated given their association within the philosophical literature and the fact that prior research has simply combined both measures (Goodwin and Darley 2008; Wright et al. in press), but it is logically possible, if unusual, to be (for example) a non-cognitivist universalist (e.g., Hare 1952).

While we suspect that conceptual dissociations between metaethical commitments partially explain our results, and that the findings are doubtless influenced by our particular methods of measurement, our results also point to three possible (and mutually consistent) proposals concerning the psychology of metaethical belief.

First, as suggested by Wright et al. (in press), it could be that objectivism in the moral domain is tempered by the need to tolerate and effectively interact with others who hold divergent beliefs. On this view, the apparent incoherence in participants' metaethical commitments serves an important psychosocial function, and we would expect the observed relationship

between the prevalence of a particular moral belief and an objectivist stance toward it.

Second, it could be that people do not hold a single “intuitive theory” of metaethics, but instead hold multiple theories with some moral content. For example, intuitive theories could be organized around general principles (such as fairness vs. justice), moral patients (such as humans vs. non-humans), or particular practices (such as slavery vs. marriage). This idea can help make sense of reliable relationships between metaethical commitments and other beliefs (e.g., descriptive beliefs about consensus, explanatory beliefs), attitudes (e.g., tolerance), and behaviors (e.g., charitable giving) despite only modest associations across different metaethical beliefs. This proposal builds on prior research positing “intuitive theories” across a wide range of domains, where such theories embody somewhat coherent but not full articulated bodies of belief (e.g., Carey 1985; Shtulman 2010; Thagard 1989). In the moral domain, for example, Lombrozo (2009) investigated the relationship between deontological versus consequentialist commitments and found evidence of a systematic but imperfect correspondence across more abstract and explicit versus scenario-based measures. With explicit articulation and examination, as typically occurs with philosophical training, different metaethical commitments could potentially become more reliably associated.

Finally, it could be that categories that make sense *a priori* philosophically play a relatively minor role in driving peoples’ responses, with a much greater role for (arguably philosophically irrelevant) properties, such as whether the question prompts are abstract or concrete. Both our factor analysis—which suggested that the dimension of “abstract” versus “concrete” was more psychologically significant than that between tendency and inevitability—and the patterns of correlations across measures support the importance of this dimension. Along these lines, Nichols and Knobe (2007) found that concrete vignettes about free will elicited compatibilist responses, while abstract vignettes elicited incompatibilist responses. More generally, research on Construal Level Theory suggests that level of abstraction can have important consequences for cognition (Trope and Liberman 2010). This final point should give pause to the assumption that folk morality will have any clean correspondence to extant philosophical categories. Instead, a more bottom-up, data-driven approach to understand folk moral commitments may be more successful.

Further research aimed directly at measuring the nature of metaethical commitments will aid in distinguishing these possibilities and further clarify the status and coherence of folk metaethical commitments. If such commitments don't correspond to philosophical distinctions that can be motivated a priori, which dimensions of moral experience do they track, and why? These are important questions for future research.

Notes

- * Author's Note: Kevin Uttich, University of California—Berkeley, George Tsai, University of California—Berkeley and University of Hawaii, and Tania Lombrozo, University of California—Berkeley. Corresponding author: Kevin Uttich, Email: uttich@berkeley.edu, 3210 Tolman Hall, Berkeley, CA, 94720. We thank Nicholas Gwynne, Michael Pacer, Kathie Pham, Jennifer Cole Wright, Hagop Sarkissian and the Moral Psychology group at Berkeley and the Concept and Cognition lab for feedback and data collection assistance. This research was supported by research funds from the McDonnell Foundation.
- 1 Information on the norms of the time period was added to the historical scenario to ensure that participants were aware of the relevant norms and understood that the scenario takes place before the change in the social fact.
 - 2 We obtained similar results in a separate experiment which used Goodwin and Darley's original multiple-choice format rather than a Likert scale: 18 "true" responses (6%), 102 "false" responses (35%), and 162 "opinion" responses (56%) out of 288 total responses (96 participants \times 3 social facts).
 - 3 We thank Jennifer Cole Wright for suggesting this interpretation for why the concrete items may have shown a differentiation between "tendency" and "inevitability" while the abstract items did not.

References

- Carey, S. (1985). *Conceptual Change in Childhood*. Cambridge, MA: MIT Press.
- Cohen, J. (1997). The arc of the moral universe. *Philosophy and Public Affairs*, 26(2), 91–134.

- Dalbert, C., Montada, L., and Schmitt, M. (1987). Glaube an eine gerechte Welt als Motiv: Validierungskorrelate zweier Skalen [Belief in a just world as motive: Validity correlates of two scales]. *Psychologische Beiträge*, 29, 596–615.
- Forsyth, D. R. (1980). A taxonomy of ethical ideologies. *Journal of Personality and Social Psychology*, 39, 175–84.
- Furnham, A. (2003). Belief in a just world: Research progress over the past decade. *Personality and Individual Differences*, 34, 795–817.
- Goodwin, G. P., and Darley, J. M. (2008). The psychology of meta-ethics: Exploring objectivism. *Cognition*, 106, 1339–66.
- (2010). The perceived objectivity of ethical beliefs: Psychological findings and implications for public policy. *Review of Philosophy and Psychology*, 1, 161–88.
- (2012). Why are some moral beliefs seen as more objective than others? *Journal of Experimental Social Psychology*, 48, 250–6.
- Hare, R. M. (1952). *The Language of Morals*. New York: Oxford University Press.
- Harman, G. (1975). Moral relativism defended. *Philosophical Review*, 84, 3–22.
- King, M. L., Jr. (1986). *A Testament of Hope: The Essential Writings of Martin Luther King, Jr.* In J. Washington (ed.). San Francisco: Harper and Row.
- Knobe, J., Buckwalter, W., Nichols, S., Robbins, P., Sarkissian, H., and Sommers, T. (2012). Experimental philosophy. *Annual Review of Psychology*, 63, 81–99.
- Lerner, M. (1980). *The Belief in a just World: A Fundamental Delusion*. New York: Plenum.
- Lombrozo, T. (2009). The role of moral commitments in moral judgment. *Cognitive Science*, 33, 273–86.
- Nichols, S. (2004). After objectivity: An empirical study of moral judgment. *Philosophical Psychology*, 17, 3–26.
- Nichols, S., and Knobe, J. (2007). Moral responsibility and determinism: The cognitive science of folk intuitions. *Nous*, 41, 663–85.
- Sarkissian, H., Park, J., Tien, D., Wright, J. C., and Knobe, J. (2011). Folk moral relativism. *Mind & Language*, 26, 482–505.
- Shtulman, A. (2010). Theories of god: Explanatory coherence in a non-scientific domain. In S. Ohlsson and R. Catrambone (eds), *Proceedings of the 32nd Annual Conference of the Cognitive Science Society*. Austin, TX: Cognitive Science Society, pp. 1295–300.
- Sinnott-Armstrong, W., Moral skepticism. *The Stanford Encyclopedia of Philosophy (Summer 2009 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/sum2009/entries/skepticism-moral/>.
- Thagard, P. (1989). Explanatory coherence. *Behavioral and Brain Sciences*, 12, 435–502.

- Trope, Y., and Liberman, N. (2010). Construal-level theory of psychological distance. *Psychological review*, 117(2), 440.
- Uttich, K., Tsai G., and Lombrozo, T. (2013). *Ethical Explanations, Moral Objectivism, and Moral Progress*. Manuscript submitted for publication.
- Wright, J. C., Grandjean, P. T., and McWhite, C. B. (2013). The meta-ethical grounding of our moral beliefs: Evidence for meta-ethical pluralism, *Philosophical Psychology*, 26(3), 336–61.
- Wright, J. C., McWhite, C. B., and Grandjean, P. T. (in press). The cognitive mechanisms of intolerance: Do our meta-ethical commitments matter? In J. Knobe, T. Lombrozo, and S. Nichols (eds), *Oxford Studies in Experimental Philosophy*, Vol. 1, New York, NY: Oxford University Press.
- Young, L., and Durwin, A. (2013). Moral realism as moral motivation: The impact of meta-ethics on everyday decision-making. *Journal of Experimental Social Psychology*, 49(2013), 302–6. doi:10.1016/j.jesp.2012.11.013