



The psychology of meta-ethics: Exploring objectivism [☆]

Geoffrey P. Goodwin ^{*}, John M. Darley

Department of Psychology, Princeton University, Princeton, NJ 08544, USA

Received 26 October 2006; revised 29 May 2007; accepted 22 June 2007

Abstract

How do lay individuals think about the objectivity of their ethical beliefs? Do they regard them as factual and objective, or as more subjective and opinion-based, and what might predict such differences? In three experiments, we set out a methodology for assessing the perceived objectivity of ethical beliefs, and use it to document several novel findings. Experiment 1 showed that individuals tend to regard ethical statements as clearly more objective than social conventions and tastes, and almost as objective as scientific facts. Yet, there was considerable variation in objectivism, both across different ethical statements, and across individuals. The extent to which individuals treat ethical beliefs as objective was predicted by the way they grounded their ethical systems. Groundings which emphasize the religious, pragmatic, and self-identity underpinnings of ethical belief each independently predicted greater ethical objectivity. Experiment 2 replicated and extended these findings with a refined measure of ethical objectivism. Experiment 3 demonstrated the robustness of the religious grounding of ethics, and differentiates it from mere religious belief and from political orientation. The results shed light on the nature of ethical belief, and have implications for the resolution of ethical disputes.

© 2007 Elsevier B.V. All rights reserved.

[☆] The first author thanks the Woodrow Wilson School, Princeton University for financial support. For their helpful comments, we thank Adam Alter, Sam Glucksberg, Adele Goldberg, Dena Gromet, Phil Johnson-Laird, Dan Osherson, Joe Simmons, Erika Sloan.

^{*} Corresponding author. Tel.: +1 609 258 9498; fax: +1 609 258 1113.

E-mail address: ggoodwin@princeton.edu (G.P. Goodwin).

Keywords: Moral cognition; Moral reasoning; Ethical objectivism; Meta-ethics; Ethical disagreement

1. Introduction

Debates about ethical issues commonly arise in everyday life. One sort of debate concerns the correct means to achieve an agreed upon end – for instance, is low taxation an efficient way to achieve a just and productive society? Another sort of debate concerns ethical ends themselves – for instance, is terminating the life of a terminally ill and pain-ridden individual morally defensible? The first sort of debate could in principle be resolved empirically. But the second sort of debate arguably cannot be resolved empirically. Here, instead, the focus is on foundational ethical principles. These sorts of disagreement are potentially irresolvable and are likely to be psychologically complex. And they are the focus of this paper.

Deep ethical disagreements are interesting partly because they give individuals occasion to think about how they would defend or justify their ethical beliefs. Reflective individuals might even be inclined to think about the epistemic status of their ethical beliefs, and what sort of ground they might have. Philosophers, of course, have spent a good deal of time on such “meta-ethical” issues. Some philosophers have argued that there are no moral facts, and that morality is not objective (e.g., Ayer, 1936; Blackburn, 1984; Hare, 1952; Harman, 1975; Mackie, 1977; Williams, 1985), whereas others have argued for the opposite position (e.g., Brink, 1986; Kant, 1959; Nagel, 1970; Railton, 1986; Smith, 1994; Sturgeon, 1985). This debate is real and not settled in the philosophical community. Indeed, within philosophy, “there are no dominant views” (Smith, 1994, p. 4).

But, how do ordinary individuals perceive and think about meta-ethics? Do they regard their ethical beliefs as factual and objective, or as more subjective and preferential? Curiously, this question has been largely unexplored. Most psychological investigations of morality to date have been concerned with questions of practical ethics, that is, with questions about the ethical beliefs and practices that individuals abide by (e.g., Baron & Spranca, 1997; Darley & Shultz, 1990; Haidt, 2001; Kohlberg, 1969, 1981; Maio & Olson, 1998; Piaget, 1965; Tetlock, 2003). The psychology of meta-ethics has been explored tangentially in the child development literature, which has focused on whether and at what age children are capable of distinguishing conventional from ethical rules. The evidence from this literature is extremely controversial (see e.g., Gabennesch, 1990a, 1990b; Helwig, Tisak, & Turiel, 1990; Shantz, 1982; Shweder, 1990; Tisak & Turiel, 1988; Turiel, 1978).

Moreover, psychological research that has specifically focused on meta-ethics, has not addressed questions concerning ethical objectivism. Instead, it has focused on the distinction between ethical universalism and ethical relativism – i.e., whether individuals treat their ethical beliefs as applying to all people, and all cultures (Nichols & Folds-Bennett, 2003). Participants in these studies are asked whether a particular moral belief they hold is shared by, or applicable to, all people or all cultures (e.g., Nichols & Folds-Bennett, 2003; Turiel, 1978), or whether ethics generally is depen-

dent on the individual or on culture (e.g., Forsyth, 1980, 1981; Forsyth & Berger, 1982). However, the question of whether ethical standards should apply to all cultures is a question about the *scope* of ethical standards, and is independent of the question of whether such standards and beliefs are objectively or subjectively true (Sayre-McCord, 1986; Snare, 1992; Williams, 1972). Our interest centers on this second question, which concerns the *source* of such beliefs or standards – whether they derive their truth (or warrant) independently of human minds (i.e., objectively), or whether instead, their truth is entirely mind-dependent or subjective (Sayre-McCord, 1986).¹

It is possible that the majority of people do not think about meta-ethics at all. However, we start from the assumption that individuals can be induced to think about meta-ethics. Certain civic or religious commitments may call for reflection on one's basic values and the source of those values. Indeed, one way that the question of how there could be objective moral facts might be answered is to view ethical statements as having a religious foundation – for instance, they are the word of God. As some have argued, those who ground their ethical beliefs in the notion of a divine being, are likely to view the source of morality as external and objective (e.g., Hunter, 1991). Conversely, those who do not, are more likely to conceive of ethics as internal and subjective. Hunter (1991) views this distinction as one between “orthodox” and “progressivist” ideologies, and regards it as a more stubborn moral impasse than that between conservatives and liberals (pp. 127–128). Grounding ethics in religion does not appeal to everyone, of course, and it is by no means the only alternative for one who wishes to be an ethical objectivist. In sum, there may be a rich variety of meta-ethical positions that lay individuals report.

Accordingly, the three main questions guiding the present research were: (1) Do individuals tend to be objectivists or subjectivists concerning ethics? (2) How do they regard the objectivity of ethical statements alongside other sorts of statements such as scientific facts or statements of taste? (3) What factors predict individual differences in degree of objectivism?

To answer these questions, our method was to investigate individuals' responses to situations in which another individual ostensibly disagreed with them about a particular ethical issue. Our motivation for this strategy was twofold. First, we considered it inadequate simply to ask people whether they were objectivists or subjectivists with regard to a particular ethical issue. Put baldly like this, the question is so ambiguous that it could admit multiple differing interpretations, and we would have no control over, or knowledge of, the precise interpretation that any participant adopted. Second, although there are a variety of ways that philosophers have distinguished objectivism and subjectivism, one simple and respectable formulation is as follows: if an individual takes a particular ethical claim to be true, and regards sit-

¹ This is not to deny that those who believe that an ethical statement is applicable to all people are more likely to consider the statement as objectively true, i.e., ethical objectivism and universalism are likely to be correlated psychologically. But, in theory they are distinct. One could, for instance, be an objectivist relativist (believing that it is objectively true that what is ethical varies by person or culture), or a subjectivist universalist (believing that although certain ethical standards are not true in any objective sense, they nevertheless apply to all persons and cultures).

uations of ethical disagreement as necessarily implying that at least one party is *mistaken*, then they are an objectivist (with respect to that statement), whereas if they instead allow that neither party need be mistaken, then they are a subjectivist (Smith, 1994; Snare, 1992).

With this in mind, we implemented the following general methodology across three experiments. We first presented individuals with a range of ethical and non-ethical statements in order to gauge how strongly they agreed or disagreed with them, and whether they regarded them as truths or mere opinions. We then told them that some individuals who we had previously tested (or in Experiment 3, hypothetical individuals) disagreed with them about certain statements. The key question concerned how our participants would regard such individuals – are they mistaken (indicating objectivism), or is it possible that neither party need be mistaken (indicating some form of subjectivism)?

Previous research suggests that participants will experience some tension in trying to gauge the objectivity of their ethical beliefs. On the one hand, we know that people can reliably distinguish ethical from social conventional violations (e.g., Turiel, 1978, 1983). There is also evidence that ethical values are akin to truisms in that they are widely shared and rarely questioned (Maio & Olson, 1998). On the other hand however, people's ethical values differ widely (see e.g., Schwartz, 1992), and individuals are liable to know that there is widespread disagreement about many ethical matters. Accordingly, we predicted that individuals would treat ethical statements as more objective than statements of taste or social convention, although not quite as objective as statements of plain or scientific fact. Moreover, we predicted that this pattern would obtain controlling for how strongly they agreed with the statements in the first place.

We also investigated the ways in which our participants *grounded* their ethical systems. Philosophically speaking, ethical objectivism can stem from multiple sources or groundings. Religious groundings tend to be objective (Hunter, 1991). Similarly, groundings which emphasize the *intrinsic* rightness or goodness of core ethical beliefs tend also to be objective. However, groundings which stress the *instrumental* utility of certain ethical beliefs are non-objective. These background observations allowed us to predict the following three factors to link to greater objectivism: citing a *divine being* as providing the foundation for one's ethical system, viewing the holding of certain ethical beliefs as important and universal constituents of being a *good person*, and viewing certain ethical beliefs as *self-evidently* true. However, we predicted that a pragmatic or instrumental justification – that society could not survive unless its citizens held these beliefs – would not predict greater objectivism.

2. Experiment 1

2.1. Method

2.1.1. Participants

Fifty undergraduate students (22 male, 28 female) from Princeton University participated for course credit.

2.1.2. Design, materials, and procedure

The participants acted as their own controls, and completed two separate parts of the experiment. In the first part, they rated their level of agreement or disagreement with 26 statements (on a six-point scale ranging from 1: strongly disagree to 6: strongly agree), and whether they thought the statement was true, false, or an opinion. The statements were chosen to be prototypical instances of four main categories: ethical statements, statements of social convention, statements of artistic or aesthetic taste, and plainly factual or scientific statements.² The statements were selected on the basis of pilot testing as those that tended to produce either relatively strong agreement or disagreement. An illustrative factual statement was: *Boston (MA) is further north than Los Angeles (CA)*; an example ethical statement was: *Robbing a bank in order to pay for an expensive holiday is a morally bad action*; an example social convention statement was: *Wearing pajamas and bath robe to a seminar meeting is wrong behavior*; and an example taste statement was: *Frank Sinatra was a better singer than is Michael Bolton*. The full set of 26 statements is shown in [Appendix A](#). Halfway through the experiment, statements 2 and 5 were changed, as shown in [Appendix A](#) (reasons for this change are discussed in Section 2.2). Each participant received the statements in the order shown in [Appendix A](#). They were instructed as follows:

We are interested in how people think about a range of issues. The first part of the experiment asks you to rate your agreement with 26 statements (on a scale from 1 to 6), and to indicate your opinion about the status of each statement – whether it is true, false, or an opinion. Please read each statement carefully, and give each question your full consideration. Accept only the information given and try not to introduce additional assumptions that go beyond the information as stated. That is, try to interpret the statements in as “normal” and non-exceptional a way as possible. If an event is described, assume that it occurs or occurred in the U.S.A. The second part of the experiment will ask several follow-up questions regarding the statements.

The instruction not to introduce additional assumptions and to interpret the statements in a “normal” way was included because the second phase of the experiment confronted participants with someone who disagreed with them. We wanted to curtail the range of possible interpretations of each statement, thereby reducing the ambiguity surrounding possible disagreement (i.e., to forestall the possible interpretation that a person who disagreed with the statement was simply thinking of a different situation or set of circumstances; see [Turiel, Hildenbrandt, & Wainryb, 1991](#)).

Participants marked their level of agreement on the six-point scale for each statement, and they then answered the question:

How would you regard the previous statement? Circle the number.

² We are aware that some ethical philosophers regard ethical statements as factual. Here we use the term *factual* to denote a relatively broad category of statements that are part of general knowledge, and which may be considered statements of empirical or scientific truth, and which are clearly not ethical.

- (1) True statement.
- (2) False statement.
- (3) An opinion or attitude.

After they had made both judgments for each of the 26 statements, participants performed a task for an unrelated experiment, allowing the experimenter to examine their responses and prepare the next phase of the experiment. During this time, the experimenter located five statements that the participants were to respond to in the second phase of the experiment: two ethical statements, and a single social convention, taste, and factual statement that the participant had indicated relatively strong agreement or disagreement with (the exact procedure is described later in Section 2.1.2). These statements were then entered by hand into an open-ended response sheet shown in [Appendix B](#).

The purpose of this second phase of the experiment was to examine participants' reactions to another person who disagreed with them. Prior to completing this stage, participants were instructed as follows:

Earlier you rated your agreement with a set of statements. 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. In what follows, you will be asked to indicate how you interpret disagreement with your own attitudes. Please give each question your full consideration. Accept only the information given and try not to introduce additional assumptions that go beyond the problem as stated. Remember also that the other people who rated these statements were instructed in the same way that you were – i.e., they were instructed to interpret the statements in a normal and non-exceptional way, and not to introduce additional assumptions.

The principal question asked how participants interpret the information that another person disagrees with them. Each participant had to select one of the following options:

- (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.

As previously discussed, some ethical philosophers (e.g., [Smith, 1994](#); [Snare, 1992](#)) have often taken one of the hallmarks of ethical objectivism to be the implication that, in cases of genuine moral disagreement (i.e., a disagreement about moral ends, rather than a disagreement about the best means to pursue a particular end), at least one of the parties must be mistaken. We combined the truth versus opinion ratings at the first stage of the experiment with the responses to disagreement at the second stage to create a simple scale of objectivism. We distinguish three sorts of response in terms of their level of objectivism. The most objective response (which

we term *fully objective*) is to regard a particular belief as true (or false), and to regard someone who disagrees with that belief as surely mistaken.³ An *intermediately objective* response is to regard a particular belief as true (or false), but to see no need for either party to be mistaken if another person disagrees with that belief. Alternatively, a second type of *intermediately objective* response is to regard a particular belief as an opinion, but to regard a disagreeing other as surely mistaken. In terms of ethical beliefs, the first of these intermediate positions is consistent with philosophical subjectivism – a moral subjectivist regards their ethical beliefs as *true*, but only in a mind-dependent way, i.e., because of certain mental states that they (or perhaps the members of their group) have, and not because of something external to their own minds. Other people who have different minds can hold different ethical beliefs without either party necessarily being mistaken (see Sayre-McCord, 1986). Ethical truth on this view is thus a kind of personal or subjective truth. Although participants who respond in this way are unlikely to be self-conscious philosophical subjectivists, this response is philosophically defensible, and does at least capture the basic notion of a personal ethical truth. The second intermediate position seems more inconsistent, and may indicate oscillation (or possibly confusion) regarding a statement's objectivity. Yet, although quite different, both of these responses share a basic tension between trying to imbue certain beliefs (ethical or otherwise) with objective and subjective aspects, and so we classify them together as intermediately objective. Finally, the *least objective* response is to regard a particular belief as an opinion, and to see no need for either party to be mistaken if another person disagrees.

The procedure of the second phase of the experiment was as follows. Having inspected the participants' responses in the first phase of the experiment, the experimenter selected a set of five statements that the participant had expressed either strong agreement (responses 5, 6), or disagreement (responses 1, 2) with. The set always included two ethical statements, one social convention statement, one statement of taste, and one factual statement. In choosing which statements were presented, we followed a procedure which prioritized keeping the strength of agreement scores as constant as possible across participants. This meant that participants did not always respond to the same statements as each other, since their agreement scores for any particular statement were not identical. However, the procedure kept the statement sets as constant as possible given the differences in agreement scores. The percentage of statements used for each category is shown in Appendix C. Participants were asked to fill out the questions presented in Appendix B for each of the five statements – the order of the statements was randomly determined for each new participant. Our main prediction was that on the three point scale of objectivism that we have outlined, there should be a decreasing trend of objectivist responses across factual, ethical, social convention, and taste statements. In other words, we predicted that individuals would treat ethical statements as more objective than statements of taste or social convention on the one hand, although not quite as objective as factual statements.

³ An alternative and similarly objective response (both here and for the intermediately objective responses) is to say that the other person may in fact be correct and you yourself mistaken – this indicates a view that the issue is one where mistakes are possible. This response, however, was exceedingly rare in our data.

Finally, at the end of the experiment, we asked another set of questions about the way individuals ground or justify their ethical positions. The participants could select as many of the groundings shown in [Appendix B](#) as applied to them. We predicted that individuals who ground their ethical systems in the notion of a divine being, or in the notion of a moral self-identity, or in the “self-evidence” of their ethical beliefs, ought to be more objective about ethics than individuals who do not ground their ethical systems in such a way. Moreover, we suspected that those individuals who viewed ethics as tightly dependent on religion, i.e., those who could not conceive of right or wrong acts without the existence of a divine being, would be particularly objective about ethics, although this idea was not tested until the second experiment. We did not predict that grounding one’s ethical system in its pragmatic benefits would be associated with greater objectivism.

2.2. Results

Basic descriptive data for the first stage of the experiment are presented in [Table 1](#). It indicates the mean agreement–disagreement ratings for each of the 26 statements, and the overall frequencies of true, false, and opinion ratings.

Considering the ethical statements, it is noticeable that the assignment of truth to ethical statements varies considerably with the content of the statement. Participants generally agreed (on a six-point scale) with the goodness of anonymous donations (5.42), the badness of opening gunfire on a crowd (5.79), or of robbing a bank (5.77), and the wrongness of conscious racial discrimination (5.86) or of cheating on a lifeguard exam (5.72). But they varied considerably in how likely they were to regard these statements as true: 36%, 68%, 61%, 54%, and 58%, respectively. Perhaps more strikingly, although participants generally agreed (albeit not as strongly) with the permissibility of abortion (4.12), assisted death (4.36), and stem cell research (4.58) in the way we described them, they were highly reluctant to assign truth to statements expressing this agreement: 2%, 8%, and 2%, respectively. In other words, meta-ethical judgments about the truth of ethical claims appear to be highly sensitive to the content of the claims in question (i.e., robbery vs. abortion), and not merely to whether the claims are generally agreeable. Only 13 out of 50 participants applied the same category (truth vs. opinion) to the eight ethical statements they rated in the first part of the experiment.⁴ The remaining 37 varied their assignment of truth/falsity versus opinion in some way.

We now turn to consideration of the data from the second and more important phase of the experiment. The primary question concerned how people would respond to the new knowledge that somebody else we had tested disagreed with them. To recapitulate our main predictions, we expected that on this measure there would be a decreasing tendency to view statements as objectively true across the following categories: factual, ethical, social convention, and taste.

⁴ For this analysis, we examined ratings of truth versus opinion for statements participants agreed with, and ratings of falsity versus opinion for statements participants disagreed with.

Table 1
Mean agreement ratings, and the overall percentages of true, false, and opinion responses for the 26 statements used in Experiment 1

Statement Category	Content	Mean agreement scores (s.error)	Percentage of “true” responses	Percentage of “false” responses	Percentage of “opinion” responses	Percentage of “other” responses
Fact	Geography	5.78 (.12)	98		2	
	Evolution	5.28 (.19)	68	4	28	
	Earth	5.58 (.15)	92		8	
	Exercise	5.68 (.09)	96		4	
	Mars*	1.54 (.18)	26	68	6	
	Mean	5.56 (.15)	84	6	10	
Ethics	Donate	5.42 (.10)	36		64	
	Gunfire	5.79 (.12)	68		32	
	Robbery	5.77 (.09)	61		39	
	Discrimination	5.86 (.06)	54		46	
	Cheating	5.72 (.09)	58		42	
	Abortion	4.12 (.25)	2	6	92	
	Euthanasia	4.36 (.20)	8	2	90	
	Stem cells	4.58 (.22)	2	4	94	
	Testimony*	1.90 (.19)	2	42	56	
	Mean	5.23 (.15)	37	2	62	
Convention	First name	4.58 (.14)	40		60	
	Pajamas	4.14 (.19)	18		82	
	Drive left	5.59 (.10)	86		12	
	Red light	5.54 (.12)	80		20	
	Talking*	1.70 (.14)	2	50	48	
	Mean	5.03 (.14)	55		44	
Taste	Writers	5.16 (.16)	16		84	
	Musicians	5.33 (.29)	16		79	
	Singers	4.87 (.14)	3		97	
	Speakers	5.52 (.10)	12		88	
	Schindler’s	5.16 (.15)	6		90	2
	Music*	3.28 (.17)	0	2	98	
	Painters	3.56 (.14)		2	98	
	News	4.82 (.16)	12		88	
	B. Mind	3.74 (.21)			100	
Mean	4.65 (.17)	7		91		

Note. Percentages do not always sum to 100 owing to some missing data. To compute the category averages, items which tended to produce disagreement (marked with an asterisk) were reverse coded. Missing cell values represent percentages of 0.

We were wary of the fact that some individuals might interpret ethical disagreement as indicating that the disagreeing other person might have been thinking of extraordinary extenuating circumstances, or that they had misread the question, or that they may not have understood the words used in a conventional way, and so on (see Turiel et al., 1991). The list of such possible caveats to interpreting the prima facie ethical disagreement is large. However, for each disagreement, we asked people to further comment on the source of the disagreement that they were con-

fronted with. When individuals did not interpret the disagreement in a bona fide way, as in the examples just mentioned, we excluded such data from all foregoing analyses.⁵ In fact, only seven out of a total of 102 responses were excluded on these grounds. The overall trend for objectivist responses is shown in Fig. 1.

Overall, considering the mean scores on the three-point scale of objectivism, the effect of statement content on degree of objectivism was highly significant in an ANOVA ($F(3, 102) = 39.54, p < .001$).⁶ Planned contrasts indicated that all of the adjacent comparisons were highly significant. Factual statements ($M = 2.91$) were treated as more objective than ethical statements ($M = 2.56; F(1, 34) = 9.62, p < .01$). Ethical statements were treated as more objective than statements of social convention ($M = 2.00, F(1, 34) = 9.62, p < .01$), which in turn were treated as more objective than statements of taste or preference ($M = 1.34, F(1, 34) = 15.17, p < .001$). For the ethical statements, the most common response (50 out of 100) was fully objective (3 on the scale of objectivism).⁷ The next most common response was intermediately objective (28 out of 100; 11 of these were “true (or) false, but neither party need be mistaken”; and 17 were “opinion, but the other party is surely mistaken”), eleven responses out of 100 occupied the least objective position and the remaining 11 responses could not be categorized because participants chose to respond *other* to the question regarding disagreement.

The four statement categories differed in objectivism, as we have just demonstrated, but they also differed in strength of agreement ratings, $F(3, 46) = 22.16, p < .001$. In order to control for this, we re-computed the ANOVA on four new within-subjects variables that were created by subtracting the strength of agreement ratings from the objectivism ratings (having rescaled the objectivism scores so that they were on the same 1–6 scale as the strength scores).⁸ The resulting ANOVA was again highly reliable ($F(3, 102) = 39.57, p < .001$). On this analysis, planned contrasts indicated that factual statements ($M = -.11$) were treated as more objective than ethical statements ($M = -.83; F(1, 34) = 10.44, p < .01$). Ethical statements were not treated more objectively than statements of social convention although

⁵ For instance, the following response was excluded for the gunfire version of question 2: “A difference in perception of a situation in which gunfire was opened on a crowded city street. I was thinking gunfire from terrorists/criminals; other person may have thought gunfire from police officers to catch a criminal”. Responses in which participants said they needed more information about the context of the events were also excluded for this reason. However, responses in which participants said that perhaps the disagreeing other was operating with a different sense of “morally wrong/bad” were not excluded. Statements 2 and 5 were changed midway through the experiment (see earlier), because their original versions were more likely to allow disagreements of this non-bona fide sort.

⁶ There were some missing data, owing to the fact that some responses to disagreement were not interpretable. This meant that only 35 subjects were involved in this ANOVA, but the effect remained and was strengthened when each of the statement category means were substituted for the missing data.

⁷ Only one of these 50 fully objective responses was based on the participant having said that: “It could be that you are mistaken (i.e., *the participant themselves*), and the other person is correct”, the remaining 49 were based on the participant having said that: “The other person is surely mistaken”.

⁸ This method presented itself as the simplest and best method of controlling for strength of agreement scores. There is no simple ANOVA procedure for controlling for a covariate (i.e., strength of agreement) which differs at each level of the repeated-measures variable (i.e., category of claim).

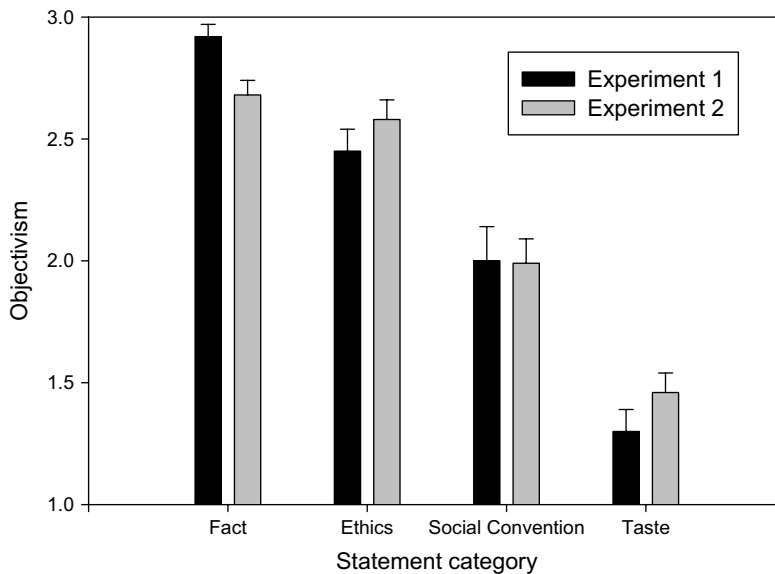


Fig. 1. The level of objectivism shown for the four types of statement in Experiments 1 and 2.

the trend was in that direction ($M = -1.34$; $F(1, 34) = 2.46$, $p < .13$). And statements of social convention were treated more objectively than statements of taste ($M = -3.00$, $F(1, 34) = 26.38$, $p < .001$).

We were interested in whether the ways individuals grounded their ethical systems might predict different levels of moral objectivism. For each of the four groundings that we presented (see [Appendix B](#)), a binary variable coded whether participants cited that grounding or not. In order to test for the independent effects of each of the four grounding variables, they were simultaneously entered as predictors of ethical objectivism in a multiple regression analysis. The resulting model was significant, $F(4, 44) = 3.30$, $R^2 = .23$, $p < .02$. As predicted, individuals who grounded their ethical beliefs in the notion of a divine being were more objective than those who did not, $\beta = .27$, $sr^2 = .25$, $p < .04$, one-tailed. Similarly, those who grounded their ethical beliefs in their importance to a moral self-identity (“every good person on earth holds these beliefs”) were more objective than those who did not, $\beta = .25$, $sr^2 = .23$, $p < .05$, one-tailed. And surprisingly, individuals who cited a pragmatic grounding for their ethical beliefs (“society could not survive without its citizens holding these beliefs”) were marginally more objective about those beliefs than those who did not cite such a reason, $\beta = .22$, $sr^2 = .22$, $p < .06$, one-tailed. Contrary to the prediction, individuals who claimed that their moral beliefs were self-evident were not more objective than those who did not.

The different groundings we examined were relatively independent. The only reliable correlation between them was that grounding ethics in a divine being was somewhat related to grounding in the notion of a moral self-identity, $r(50) = .31$, $p < .03$. Yet, the more groundings an individual cited (of the three that were shown to play some role in predicting objectivism: divine being, moral self-identity, pragmatism),

the more objective about ethics they tended to be. This relationship was very clear-cut. Individuals who cited none of the three grounding reasons were the least objective about ethical statements (1.5). They were less objective than those who cited one grounding (2.42), who in turn were less objective than those who cited two groundings (2.65), who themselves were less objective than those who cited three groundings (2.81). The number of groundings an individual cited was a highly reliable predictor of objectivism in a hierarchical regression analysis in which strength of agreement was controlled for by entering it in the first block. Strength of agreement did not predict a significant proportion of variance in objectivism, $R^2 = .02$, $F(1, 47) = 1.08$, $p > .3$. After adding number of groundings into the regression, $R^2 = .29$, $F(2, 46) = 9.24$ ($p < .001$). The change in R^2 produced by adding number of groundings was highly reliable, $F(1, 46) = 17.23$ ($p < .001$).

Within the category of ethical statements, the final composite measure of objectivism significantly correlated with strength of agreement, $r(49) = .33$, $p < .02$.⁹ One might therefore worry that participants were using the truth versus opinion distinction to represent degrees of certainty or confidence they had in each statement, rather than to mark an epistemic distinction between them. This objection cannot explain the fact that participants agreed very strongly with the taste statements that we gave them, but did not treat them as objective (see [Tables 1 and 2](#)). They were able to differentiate between the strength of their attitude (or certainty) towards these statements and the statements' objectivity. This objection is further refuted by the results of Experiment 2, which we present next.

Experiment 2 was designed to replicate the findings of Experiment 1 using a different methodology. In the first stage of the experiment, rather than asking participants to assign truth, falsity or opinion to the various statements, we instead asked them a simple question: can there be a correct answer as to whether the statement in question is true? This question is very clearly concerned with the nature of the statement itself, and cannot be interpreted as asking for an assessment of confidence.

3. Experiment 2

3.1. Method

3.1.1. Participants

Seventy-one undergraduate students (29 male, 42 female) from Princeton University participated for course credit, and five students (3 male, 2 female) from Princeton's theological seminary participated for a payment of \$10.¹⁰

⁹ This correlation was also reliable for the three other statement categories.

¹⁰ We included the seminarians in order to increase the number of participants who grounded their ethics in the notion of a divine being. Circumstances precluded our including as many seminarians as we had planned, although indeed, four out of the five did ground their ethics in the notion of a divine being, compared with 18 out of 71 of the remaining participants. However, the seminarians who grounded their ethical beliefs in this way were not more objective about ethics than the other students who did so – in fact, they were practically identical on this measure ($M = 2.75$, seminarians; $M = 2.74$, non-seminarians, n.s.).

Table 2
Mean agreement ratings, and the overall frequencies of correct answer and no correct answer responses for the 26 statements used in Experiment 2

Statement category	Content	Mean agreement scores (s. error)	Percentage of “correct answer” responses	Percentage of “no correct answer” responses
Fact	Geography	5.53 (.12)	97	3
	Evolution	5.43 (.12)	83	17
	Earth	5.64 (.09)	85	15
	Exercise	5.53 (.08)	84	16
	Mars*	1.79 (.17)	96	4
	Mean	5.47 (.12)	89	11
Ethics	Donate	5.43 (.08)	51	47
	Robbery	5.29 (.16)	83	17
	Discrim.	5.62 (.08)	79	21
	Cheating	5.63 (.08)	83	17
	Abortion	4.03 (.21)	22	78
	Euthanasia	4.38 (.14)	17	83
	Stem cells	4.84 (.15)	24	76
	False testimony*	2.14 (.16)	66	34
Mean	5.01 (.13)	53	47	
Convention	First name	4.08 (.15)	43	57
	Pajamas	3.88 (.16)	36	65
	Drive left	5.39 (.11)	83	17
	Red light	5.29 (.11)	86	15
	Talking*	1.83 (.11)	74	26
	Mean	4.76 (.13)	64	36
Taste	Writers	4.96 (.12)	25	75
	Singers	4.68 (.13)	17	83
	Speakers	5.37 (.10)	26	74
	Schindler’s	4.75 (.12)	16	84
	Music	3.32 (.16)	4	96
	Painters	3.50 (.12)	4	96
	News	4.84 (.13)	34	66
	B. Mind	3.86 (.17)	9	91
Mean	4.41 (.13)	17	83	

Note. Percentages do not always sum to 100 owing to some missing data. To compute the category averages, items which tended to produce disagreement (marked with an asterisk) were reverse coded.

3.1.2. Design, materials, and procedure

As in Experiment 1, participants acted as their own controls and completed two separate parts of the experiment. The materials and procedure for the both parts were almost identical to those in Experiment 1. In the first part, participants rated their agreement with the 26 statements shown in [Appendix A](#) on a six-point scale (rating the replaced versions for items 2 and 5). However, instead of judging whether they thought each statement was true, false, or an opinion, they were instead asked: “According to you, can there be a correct answer as to whether this statement is true?”, and were only given the options *yes* and *no* to choose from. In the second part,

the only difference was that after 18 participants, instead of assigning statement 6: “Boston (Massachusetts) is further north than Los Angeles (California)”, as the factual statement, we chose to assign statement 12 instead: “Homo sapiens evolved from more primitive primate species”, which we still expected large agreement with.¹¹

3.2. Results

We observed a broadly similar pattern of results as occurred in Experiment 1. Descriptive data for the 26 statements are shown in Table 2. We observed the same effect of the content of the ethical statements on individuals’ tendency to regard them as objective. For instance, although participants tended to agree with the wrongness of robbing a bank (5.29), the goodness of anonymous giving (5.43), and the permissibility of assisted death (4.38), they were far more willing to say that there was a correct answer as to the wrongness of robbery (83%), than they were to say the same about the goodness of giving (52%), and highly unlikely to claim a correct answer as to the permissibility of assisted death (17%). Across the eight ethical statements judged in the first part of the experiment, only 9 out of the 76 participants made the same response to the correct answer query for all eight statements. Thus, our participants’ meta-ethical positions were again highly dependent on the specific ethical content that they were asked to judge.

We now turn to the data from the second phase, in which participants were asked to respond to disagreement. In the present experiment, unlike Experiment 1, there were no instances where the participants thought that the disagreeing other was conceiving of a different set of circumstances than they themselves were, so no responses were excluded on this basis. Objectivism scores were computed by combining responses to the “correct answer” and “mistake” questions in the same way as in Experiment 1, to create a three-point scale of objectivism, ranging from 1 (least objective) to 3 (most objective). We again predicted that on the objectivity measure, ethical statements would fall between factual statements on the one hand and statements of social convention or taste on the other. This was largely what we found, as is shown in Fig. 1, with the exception that in this experiment, ethical statements were treated just as objectively as factual statements. Overall, statement content exerted a reliable effect on objectivism, as demonstrated in an ANOVA ($F(3,198) = 58.18$, $p < .001$). Factual statements were treated the most objectively (2.70), although planned contrasts showed that ethical statements were in fact treated no less objectively (2.60; $F(1,66) = 1.21$, $p > .2$).¹² However, ethical statements were treated more

¹¹ However, this statement was only presented in the second phase of the experiment to those participants who agreed with it. If participants did not agree with this statement, we presented an alternative factual statement.

¹² As noted in the main text, for the majority of participants we used a more controversial factual statement in this experiment: “Homo sapiens evolved from more primitive primate species”, than the uncontroversial geographical statement used in the last experiment (and for the first participants in this experiment). Within the present experiment, participants who received the evolution statement were marginally less objective about it than participants who received the geographical statement were about that statement (2.64 vs. 2.87, $t(31.57)$, unequal variances) = 1.99, $p < .06$).

objectively than statements of social convention (1.97; $F(1,66) = 25.12, p < .001$), which in turn were treated more objectively than statements of taste (1.43; $F(1,66) = 19.34, p < .001$).

Concerning the ethical statements, 107 of the 152 responses (70%) were fully objective.¹³ Sixteen of the 152 responses (11%) were intermediately objective (12 of which were “correct answer, neither party need be mistaken”; and four of which were “no correct answer, other party is surely mistaken”), and 21 of the 152 responses (14%) occupied the least objective position, and the remaining eight responses (5%) were not able to classified.

As in Experiment 1, the four statement categories differed not only in objectivism, but also in strength of agreement ratings, $F(3,73) = 21.15 (p < .001)$.¹⁴ We controlled for this in the same way as in Experiment 1, by rescaling the objectivism scores so that they were on the same scale as the strength of agreement scores, and then subtracting the strength of agreement scores from them. The resulting ANOVA was still highly significant ($F(3,198) = 154.42, p < .001$), and the planned contrast results were unchanged.

We next investigated the same predictors of objectivism as those tested in Experiment 1. First, we entered these four variables into a multiple regression analysis to test for their independent effects. This model significantly predicted variance in ethical objectivism, $F(4,64) = 3.69, R^2 = .19, p < .01$. Replicating Experiment 1, individuals who grounded their ethical systems in the notion of a divine being were reliably more objective than those who did not, $\beta = .20, sr^2 = .19, p < .05$, one-tailed. Similarly, individuals who grounded their ethical systems in the notion of a moral self-identity were more objective about ethical statements than those who did not, $\beta = .32, sr^2 = .26, p < .02$, one-tailed. Unlike Experiment 1, individuals who grounded their ethical systems in their pragmatic consequences were not more objective than those who did not, $\beta = .15, sr^2 = .12, p > .14$, one-tailed. However, this is partly due to shared variance between the predictors: taken alone, the pragmatic grounding reliably predicts objectivism, $r = .29, p < .01$. The self-evident grounding was again not reliable.

The groundings were again relatively independent of each other. The only reliable correlation between them was that grounding one’s ethical system in the notion of a moral self-identity was correlated with grounding one’s ethical system in its pragmatic consequences, $r(71) = .52, p < .001$. As in Experiment 1, the more groundings an individual cited – of those that were predictive of objectivism (divine being, moral self-identity, pragmatic consequences) – the more objective they tended to be. Individuals who cited none of these groundings were the least objective (1.83). They were less objective than those who cited one grounding (2.58), who were in turn less objective than those who cited two (2.77) or three groundings (2.94). A hierarchical regres-

¹³ In all of these cases, the participants said that the other party is surely mistaken (rather than that they themselves might be mistaken and the other person correct).

¹⁴ Objectivism was also significantly correlated with strength of agreement *within* the ethical statements, as was also seen in Experiment 1, $r(74) = .53, (p < .001)$, and marginally within the social convention statements, $r(70) = .21, p < .08$, although not for the factual and taste statements.

sion analysis in which strength of agreement was entered first showed that the number of groundings cited was a reliable predictor of objectivism. In this experiment, strength of agreement predicted a significant proportion of the variance in objectivism, $R^2 = .26$, $F(1, 71) = 24.28$, $p < .001$. Adding number of groundings to the regression equation, accounted for a higher proportion of the variance in objectivism, $R^2 = .37$, $F(2, 70) = 20.80$, $p < .001$, and this change in R^2 was highly reliable, $F(1, 70) = 13.16$, $p < .01$.

3.3. Discussion: Experiments 1 and 2

The findings from Experiments 1 and 2 are highly consistent and highlight several important facets of lay meta-ethics. First, how objectively individuals treat their ethical beliefs depends to a large extent on the content of those beliefs. Lay meta-ethical systems are not monolithic, but rather, are sensitive to specific ethical content. Second, individuals seem to treat core ethical beliefs as being almost as objective as scientific or plainly factual beliefs, and reliably more objective than beliefs about social convention or taste. Third, there was variation amongst individuals in terms of how objectively they treated their ethical beliefs. Individuals who grounded their ethical systems in the notion of a divine being, or in the notion of a moral self-identity, or in the pragmatic consequences of holding those beliefs, all tended to be more objective than those who did not. Moreover, these groundings functioned independently and additively. The more such groundings individuals cited the more objective they tended to be about their ethical beliefs.

The finding that individuals who grounded their ethical systems in the notion of a divine being tended to be more objective than those who did not led us to pursue one further hypothesis concerning religious belief. Grounding one's ethics in religious belief is the most obvious way that one could be an objectivist about ethics (although not the only way), and it may give rise to a particularly strong form of objectivism. To investigate this, at the end of both experiments we had asked participants: "According to you, is it possible for there to be right and wrong acts, without the existence of God?". We predicted that participants who did not believe that there could be right and wrong acts without a God, and who also grounded their ethical systems in a God, ought to be most objective about their ethical beliefs. We collapsed the data from Experiments 1 and 2 to create three groups. The first group (unreligious grounding), consisted of 84 participants who did not cite a divine being as part of the grounding for their ethical beliefs, and were the least objective (2.43). The second group (religious grounding, but ethics not dependent on religion), consisted of 27 participants who cited a divine being, but also thought that there could be right and wrong acts without the existence of such a being, and were somewhat more objective (2.65). The third group (religious grounding, ethics dependent on religion), consisted of 12 participants who cited a divine being, and thought either that there could not be right and wrong acts without the existence of such a being, or were not sure of their position on that issue. This finally group was almost maximally objective (2.95). The extent to which people saw ethics as dependent on religion exerted a

reliable effect on objectivism, $F(2, 120) = 3.85$, $p < .03$. Moreover, this variable exerted a reliable effect on objectivism, even controlling for how strongly participants in each group held their beliefs. In a hierarchical regression analysis regressing ethical objectivism on strength of agreement and the dependence of ethics on religion (coding unreligious as 0, religious but ethics independent of religion as 1, and religious and ethics dependent on religion as 2), when strength of agreement was entered first, $R^2 = .19$, $F(1, 121) = 28.21$, $p < .001$. Adding the dependence of ethics on religion variable to this analysis accounted for more variance, $R^2 = .22$, $F(2, 120) = 16.41$, $p < .001$, and the increase in variance accounted for was reliable, $F(1, 120) = 3.93$, $p < .05$.

Experiments 1 and 2 used different methods to measure objectivism in the first stage of the experiment. Experiment 1 asked participants whether they thought various beliefs were true (or false), or alternatively, opinions, whereas Experiment 2 asked whether various statements could have a correct answer or not. A priori, this second question seemed less susceptible to being interpreted as asking for an indication of confidence. Converging with this general point, the question in Experiment 2 also produced more interpretable results for the ethical statements in terms of the intermediately objective responses. In Experiment 1, 39% of such responses were of the form “true, but neither party need be mistaken”, with the remaining 61% being “opinion, but the other party is surely mistaken”. Earlier, we remarked that this second response is harder to justify, and may indicate oscillation between objective and subjective views of ethics (or possibly confusion), whereas the first is consistent with ethical subjectivism. In Experiment 2 however, the respective percentages switched markedly – 75% of the intermediately objective responses were of the form “correct answer, but neither party need be mistaken” (which, despite the different question, can also be regarded as a philosophically defensible form of subjectivism), whereas only 25% of such responses were “no correct answer, but other party is surely mistaken”. We take this as converging evidence for the greater felicity of the question asked in Experiment 2. Notably, a similar proportional increase in this more interpretable intermediate response was observed for all three remaining statement categories from Experiment 1 to 2 (facts: 50–80%; social conventions: 55–73%; taste: 25–57%).

4. Experiment 3

Experiment 3 aimed to further extend these findings by investigating the role of religiosity and political attitudes. It also generalized our initial findings on the grounding of ethical beliefs by using a different methodology. We were interested in testing two key hypotheses. First, does mere religious belief predict ethical objectivism, or is it the link between ethics and religious belief that is important, i.e., the grounding of ethics in religion? Second, is the difference between those who ground ethics in religion versus those who do not, subsumed by more general political or cultural differences – specifically differences between liberal and conservative attitudes (see e.g., Haidt & Graham, *in press*)?

4.1. Method

4.1.1. Participants

Two hundred and forty-seven Princeton University students participated in the experiment for the chance to win a \$50 prize. Participants were solicited via an email which invited them to take part in a short web-survey.

4.1.2. Design, materials, and procedure

Participants read three short scenarios which described three separate ethical decisions. One scenario concerned the decision of someone to donate 10% of their annual income to charity; a second concerned the decision of a woman to have an abortion due to the financial strain a child would cause; and a third concerned the decision of someone to provide a false alibi for a friend who is facing murder charges. The full description of each scenario is shown in [Appendix D](#). Participants rated the extent to which they agreed with claims about the moral goodness, moral permissibility, and moral wrongness of each action, respectively (all scales in this experiment had six points). For each scenario, after making this initial agreement rating, participants were asked two further questions – first, the extent to which they thought there could be a correct answer as to whether each statement was true, and second, what they would conclude about someone who hypothetically disagreed with them on a scale that ranged from “neither of us need be mistaken” to “other person is clearly mistaken”. Participants also rated how they grounded their moral beliefs, using the same groundings as in the previous experiments.¹⁵ However, in this experiment, we asked participants to make ratings of the extent to which each of the four groundings provided support for their moral beliefs, rather than simply checking those which applied. Additionally, we asked participants to indicate their degree of belief in the existence of a supreme being or God (or supreme beings/Gods) on a scale ranging from “not at all” to “very strongly”.

The design involved two levels of counter-balancing. For half the participants, the three moral scenarios were presented first, in a randomly determined order, followed by the grounding and religious belief questions. For the other half of the participants, the reverse order was used. Within each of these two orders, the grounding question preceded the religious belief question for half the participants, whereas the reverse order was used for the remaining half. At the very end of the experiment, participants indicated their general political attitudes on a six-point scale ranging from “very conservative” to “very liberal”.

4.2. Results and discussion

The data from two participants were removed because they failed to provide responses to the ethical questions. To aid interpretation, we dichotomized scores

¹⁵ We made one change to the moral self-identity grounding, so that for this experiment it read: “Holding these beliefs is essential to what it means to be a good person”.

on all scales into those above and below the mid-point. No predictive power was lost in doing this. For each of the three ethical scenarios, ethical objectivism scores were computed as they were in the previous experiments. In this experiment, the inter-correlations between objectivism scores for the three separate ethical scenarios were low, so we treated each scenario separately. For none of the three scenarios (donating money, false alibi, abortion) did the counter-balancing of whether the ethical grounding and religious belief questions came before or after the ethics questions exert any effect on objectivism scores (p values $>.5$, $.18$, $.97$, respectively).

Our main analysis of interest was to regress objectivism scores on the four different groundings and on the religious belief and political attitudes questions. The results confirm the main conclusions of the previous experiments, and help answer the questions posed at the outset of the study. Whether or not participants grounded their ethical systems in the notion of a divine being was the only variable to reliably predict objectivism for all three scenarios. Grounding ethics in the notion of a divine being was a marginally reliable predictor of objectivism about the goodness of anonymous donations, and was the only predictor to approach significance ($\beta = .12$, $sr^2 = .09$, $p < .08$, one-tailed).¹⁶ Similarly, it was the only reliable predictor of objectivism about the wrongness of providing a false alibi ($\beta = .17$, $sr^2 = .12$, $p < .03$, one-tailed). And grounding ethics in God was a reliable predictor of the impermissibility of abortion ($\beta = .21$, $sr^2 = .16$, $p < .01$, one-tailed), alongside mere belief in God ($\beta = .25$, $sr^2 = .19$, $p < .01$, one-tailed), and grounding ethics in the self-evidence of moral beliefs ($\beta = .16$, $sr^2 = .15$, $p < .01$, one-tailed). Political attitudes did not correlate with objectivism for any of the three scenarios.

These results confirm the findings of the previous two experiments, that grounding one's ethical beliefs in the notion of a divine being predicts greater objectivism. This variable appears to have greater predictive power than does mere religious belief, although in some cases (i.e., abortion), grounding ethics in God and mere belief in God do explain independent portions of variance. Grounding ethics in God also has greater predictive power than do simple political attitudes.

5. General discussion

Most current theorizing on the psychology of ethics has focused on identifying the actual ethical beliefs and practices that individuals subscribe to (see e.g., Baron & Spranca, 1997; Darley & Shultz, 1990; Haidt, 2001; Kohlberg, 1969, 1981; Maio & Olson, 1998; Piaget, 1965; Tetlock, 2003). Although very important, this sort of research does not answer questions about how individuals think more generally about their ethical beliefs – what sort of status they accord them, and in particular, how they think about them when they are challenged. In this paper, we develop a method that distinguishes ethical objectivists (i.e., individuals who take their ethical beliefs to express true facts

¹⁶ Due to a high correlation with mere belief in God, $r(241) = .64$, $p < .001$, grounding in God was the only reliable predictor of objectivism concerning the goodness of anonymous donations when belief in God was removed from this analysis ($\beta = .18$, $sr^2 = .12$, $p < .01$, one-tailed).

about the world) from ethical subjectivists (i.e., individuals who take their ethical beliefs to be mind-dependent, and to express nothing more than facts about human psychology). According to this method, objectivism is measured by combining participants' judgments about the *status* of a particular belief that they hold (truth vs. opinion: Experiment 1; possibility of a correct answer: Experiments 2 and 3), with their judgments about the *mistakenness* of someone who disagrees with that belief.

The findings across all three experiments were very consistent. The first major finding was that individuals were not particularly consistent in their meta-ethical positions about various ethical beliefs, and were instead highly influenced by the content of the beliefs in question. This finding suggests that unlike the meta-ethical systems of philosophers, which tend to be uniform in their treatment of a range of ethical beliefs, ordinary individuals' meta-ethical systems are highly nuanced. This finding is reminiscent of work in the deductive reasoning literature, which has shown that just as our participants did not make judgments as ethical philosophers would have, ordinary individuals do not reason as logicians. They are heavily influenced by content and not just by logical form. Perhaps the most famous demonstration of this in reasoning is in the effect of content in the Wason selection task (e.g., Cheng & Holyoak, 1989; Cosmides, 1989; Griggs & Cox, 1982; Johnson-Laird, Legrenzi, & Legrenzi, 1972), but the effect has also been shown repeatedly in the effects of belief on syllogistic reasoning (e.g., Klauer, Musch, & Naumer, 2000; Newstead, Pollard, Evans, & Allen, 1992), and in the pragmatic modulation of conditionals (Johnson-Laird & Byrne, 2002).

In the present context, the effects of content are likely to be multi-faceted. However, our data suggested one obvious cause – individuals' ratings of objectivism are influenced by how strongly they hold a particular ethical belief in the first place. This finding is arguably quite peculiar because there are clear instances where strength of agreement is disconnected from objectivism. For instance, our participants tended to strongly agree with the taste statements we presented, but viewed them as much less objective than each of the other three sorts of statements. Conversely, while one might be very uncertain whether Mars is bigger or smaller than Venus, one can still be fully confident that there is an objective fact about the matter. Although our participants clearly understood the difference between strength of agreement ratings and objectivism ratings, and decoupled them in some cases, the link between strength of agreement and objectivism was still strong: individuals were more objective about those ethical statements that they strongly agreed with. This finding was not confined to the ethical statements. But the fact that it occurs for ethical statements shows that, unlike philosophers of meta-ethics lay individuals are objectivists about some ethical beliefs but not others.

Content effects in reasoning have often been attributed to the reliance on semantic rather than purely syntactic processes. One version of this explanation is that reasoners build a specific mental model of the assertions they reason from, and are thus pulled away from an abstract syntactic analysis based on the form of assertions (see e.g., Johnson-Laird & Byrne, 2002; Newstead et al., 1992). A similar process might be hypothesized in the present context. Individuals envisage the concrete ethical scenario they are presented with, and form a model of the relevant intentions and motivations at work. Their judgments of the objectivity of various moral claims are

based on this particular and concrete model. Hence, they tend to treat each ethical judgment as unique, and any requirements of judgmental consistency across ethical scenarios are not considered (see relatedly, Kahneman & Lovallo, 1993).

The second major finding was that ethical beliefs were treated almost as objectively as scientific or factual beliefs, and decidedly more objectively than social conventions or tastes. Individuals seem to identify a strong objective component to their core ethical beliefs, and thus treat them as categorically different from social conventions. Arguably, many of our participants viewed their ethical beliefs as true in a mind-independent way. Such a view need not make ethical beliefs impervious to disagreement, but it does establish them as far less contestable than social conventions which can be considered somewhat arbitrary constructions. Moreover, it implies that there is an objectively true fact of the matter concerning whether ethical beliefs are true. This interpretation is vulnerable to the criticism that we only tested a very small range of beliefs, and that one could easily construct counter-examples in which a social convention or taste would be treated more objectively than an ethical belief. Even within our own data, there are conventions that are thought to be true statements and to have a correct answer by the majority of participants. Participants tended to think that it was true that driving round a blind corner on the wrong side of the road, and driving through a red light were wrong actions (86%, 80%, respectively, Experiment 1), and that there could be a correct answer as to whether these statements were true (83%, 86%, respectively, Experiment 2). We did not ask participants to respond to a disagreeing other for these statements, but it is likely that they would have tended to be highly objectivist about these sorts of conventional statements. However, this sort of conventional statement is not a prototypical social convention, and is instead what has been referred to as a “high social utility rule” – i.e., a coordination rule that is essentially arbitrary, but that it is very important to observe once it has been set (Miller & Bersoff, 1988; Turiel et al., 1991). Transgressions of this sort of rule have been shown to be treated more seriously than transgressions of more prototypical social conventions (Miller & Bersoff, 1988). Moreover, the statements we presented to participants in the second phase of our experiments were deliberately chosen to be prototypical instances of each of the four different categories (factual, ethical, social conventional, taste). Hence, relevant counter-examples notwithstanding, our results support the claim that ethical statements are treated almost as objectively as facts, and more objectively than social conventions and tastes.

The third major finding was that the way in which individuals grounded their ethical systems predicted how objective they were about their ethical beliefs. Those individuals who grounded their ethical systems in the notion of a divine being, in the concept of a moral self-identity, or in the pragmatic consequences of a society not adhering to that system, all tended (to some extent) to be more objective than those who did not. These groundings were largely independent of each other, and the more groundings an individual cited, the more objective they tended to be, suggesting multiple routes to ethical objectivism.

The most robust of these predictors of objectivism was the religious grounding. Individuals who grounded their ethical beliefs in the notion of a divine being were more objective than those who did not across all three experiments. Moreover, the

subset of individuals who very tightly connected ethics with religion by not admitting the possibility of right or wrong acts without the existence of a God, were almost maximally objective. Experiment 3 further emphasized the importance of the link between ethics and religion, by showing that grounding ethics in God was a more powerful predictor of ethical objectivism than mere belief in God, and more powerful than political conservatism–liberalism. This null effect for political attitude is consonant with Hunter's (1991) depiction of orthodox and progressivist ideologies as cutting across simple liberal and conservative divisions (p. 118).

It was a surprise that in Experiments 1 and 2, individuals who grounded their ethical beliefs in their pragmatic consequences for society tended to be more objective than those who did not. A pragmatic grounding makes reference to the external consequences of ethical beliefs, rather than to their intrinsic truth or standing, and thus ought not to be objective. Respondents may have been reasoning that at the societal level, the pragmatic consequences of certain ethical beliefs loom large enough to imbue such beliefs with an objective character.

Several potential criticisms of these findings about ethical groundings can be raised. We have described these findings as though the groundings that we assessed gave rise to the differing levels of objectivism. Of course however, this sort of causal account cannot be supported by our existing data. It may instead be that getting people to think about ethical disagreements first, gave rise to their endorsing different sorts of ethical groundings. This possibility strikes us as unlikely, but it would require a more elaborate study to disentangle the direction of causation here. It may also be that a third variable explains both phenomena. A second criticism is that the amount of variance in the objectivism ratings that is explained by the different groundings is quite low. Evidently, the groundings that we have explored only go so far in explaining objectivism. Future research could aim to clarify which other variables are relevant to this explanation.

A question not pursued in the present research is whether meta-ethical differences in objectivism are associated with different ethical beliefs, behaviors, and practices. Philosophers have long noted that there is no logical connection between meta-ethics and practical ethics. Yet there may be some psychological connection. One hypothesis is that moral objectivism tends to be associated with a more strident ethical philosophy, for instance, one which is more inclined to endorse optional duties, such as say, the duty to give 10% of one's income to charity (see e.g., Singer, 1972). Accordingly, one potential benefit of holding an objectivist view of ethics is that it may lead to a more morally committed life. A different possible corollary concerns ethical dialogue. Some philosophers have argued that an objective view of ethics should make one more responsive to dialogue and more willing to listen to opposing sides of an ethical debate (e.g., Snare, 1992). According to this idea, since an objectivist believes that there is a truth to be discovered about ethics, she should therefore feel that ethical dialogue might expose her to new information which could change her mind. On the other hand, a subjectivist has his ethical preferences and can simply stick with them regardless of what new information comes to hand, since there simply is no truth to be uncovered about them. In contrast to this rather idealized sketch, we suspect that the opposite hypothesis is more likely to be true, and that objectivists are in

fact less willing to engage in ethical dialogue and to listen to opposing views than subjectivists. Culture wars that are fought over fundamental ethical values may become more intractable to the extent that each side of the debate harbors an objective view of the truth of its own beliefs.

Our aim in this paper was to develop a method for assessing and predicting ethical objectivism. The present results represent only a first step towards a psychological understanding of this construct, and there is a need for further empirical investigation. Ethical disagreement and ethical conduct more generally may be elucidated by better understanding its nature and determinants.

Appendix A

The 26 statements that were used in Experiments 1 and 2 in the order they were presented to participants

No.	Statement category	Statement
1	Ethical	Anonymously donating a significant proportion of one's income to charity is a morally good action
2	Ethical	Opening gunfire on a crowded city street is a morally bad action/replaced by: robbing a bank in order to pay for an expensive holiday is a morally bad action
3	Convention	Calling teachers by their first name, without being given permission to do so, in a school that calls them "Mr." or "Mrs." is wrong behavior
4	Taste	Shakespeare was a better writer than is Dan Brown (author of "The Da Vinci Code")
5	Taste	Miles Davis was a better musician than is Britney Spears/replaced by: Frank Sinatra was a better singer than is Michael Bolton
6	Factual	Boston (Massachusetts) is further north than Los Angeles (California)
7	Ethical	Consciously discriminating against someone on the basis of race is morally wrong
8	Ethical	Cheating on a knowledge section of a lifeguard exam, to obtain a job for which one is not qualified is morally wrong
9	Convention	Wearing pajamas and bath robe to a seminar meeting is wrong behavior
10	Taste	Bill Clinton is a better public speaker than George W. Bush
11	Taste	<i>Schindler's List</i> is a better film than <i>Police Academy</i>
12	Factual	Homo sapiens evolved from more primitive primate species

(continued on next page)

Appendix A (*continued*)

No.	Statement category	Statement
13	Ethical	Before the 3rd month of pregnancy, abortion for any reason (of the mother's) is morally permissible
14	Ethical	Assisting in the death of a terminally ill friend who is in terrible pain, and who wants to die, is morally permissible
15	Convention	Driving round a blind corner on the left hand side of the road (in the USA) is a wrong action
16	Taste	Classical music is better than rock music
17	Taste	Da Vinci was a better painter than was Monet
18	Factual	The earth is not at the center of the known universe
19	Ethical	Scientific research on embryonic human stem cells that are the product of in vitro fertilization is morally permissible
20	Ethical	Providing false testimony in court about the whereabouts of a friend who is being charged with murder (i.e., to protect that friend by offering an alibi) is morally permissible
21	Convention	Driving through a red light at a busy intersection because you are late for work is a wrong action
22	Convention	Talking loudly and constantly to the person next to you during a lecture is a permissible action
23	Taste	CNN provides better news coverage than does FoxNews
24	Taste	<i>A Beautiful Mind</i> is a better film than <i>The Matrix</i>
25	Factual	Frequent aerobic exercising (i.e., running, swimming, cycling) usually helps people to lose weight.
26	Factual	Mars is the smallest planet in the solar system

Appendix B

The follow-up and final questions used in Experiments 1 and 2.

B.1. Follow up questions

Take claim ____: _____.

You circled __ on the scale (1–6) which means that you strongly **agreed/disagreed** with this statement. A person who we tested, strongly **agrees/disagrees** with this statement, which means that he or she sharply disagrees with you.

- What would you conclude about this disagreement? We are interested in what you would privately think about this – the question is **not** about what you would be willing to say to this other person. Please circle the number.
 - (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 – please explain _____
- Give us your thoughts about why it is that there is disagreement. What could be its source? _____

B.2. Final questions

- This is a general question about what we could call the “grounding”, or “justification” for moral statements. People firmly believe certain moral statements (such as “do not commit murder”). Check as many of the following statements that you consider support for your moral beliefs. That is, that provide the reasons that you hold the particular set of moral beliefs that you do.
 - they are ordained by a supreme being;
 - every good person on earth, regardless of culture, holds these beliefs;
 - a society could not survive without its citizens holding these beliefs;
 - their truth is self-evident.
- According to you, is it possible for there to be right and wrong acts, without the existence of God?
 - (1) yes;
 - (2) no;
 - (3) not sure.

Appendix C

The extent to which the different statement types were presented to participants in the second phase of Experiments 1 and 2

	Experiment 1		Experiment 2	
Factual	6. Boston location	92%	12. Homo sapiens	67%
	18. Earth orbit	8%	6. Boston location	20%
			18. Earth orbit	12%
Ethical			25. Exercise	1%
	8. Cheating on exam	84%	8. Cheating on exam	97%

(continued on next page)

Appendix C (*continued*)

	Experiment 1		Experiment 2	
	2. Robbery	60%	2. Robbery	89%
	2. Open fire on crowd	36%	7. Discrimination	13%
	7. Discrimination	10%		
	20. False testimony	8%		
	21. Assisted death	2%		
Conventional	3. First name	38%	9. Pajamas to seminar	46%
	9. Pajamas to seminar	36%	3. First name	21%
	22. Talking in class	26%	22. Talking in class	33%
Taste	10. Bill Clinton	38%	10. Bill Clinton	50%
	5. Frank Sinatra	30%	5. Frank Sinatra	47%
	5. Miles Davis	24%	4. Shakespeare	1%
	4. Shakespeare	4%	17. Da Vinci	1%
	16. Classical music	2%		
	11. Schindler's List	2%		

Note. The percentages sum to 100 for each category of statement except the ethical category where they sum to 200 because two statements were presented. This may not be exact in some cases owing to rounding error.

Appendix D

The three ethical scenarios presented in Experiment 3

1. Donation

In the past, John has saved 10% of his income for vacations. But, after some deliberation, he decides that this money could be put to better use. He decides to change his savings plan so that he instead donates this saved income to charity, which he does so anonymously. Rate the extent to which you agree with the claim that John's actions are morally good.

2. False alibi

One of Megan's best friends is being charged with murder. Megan is convinced that he is innocent, although she does not know what he was doing on the night of the alleged murder. Without having been asked, Megan provides a false alibi to the police for her friend, claiming that she was with him on the night of the alleged murder. Rate the extent to which you agree with the claim that Megan's actions are morally wrong.

3. Abortion

Eve is 2 months pregnant. Despite Eve's wanting to have the child, she does not know who the father is, and after considering her financial situation, she considers that having the child would be too big a burden. She decides instead to have an abortion. Rate the extent to which you agree with the claim that Eve's choice is morally permissible.

References

- Ayer, A. J. (1936). *Language, truth, and logic*. London: Gollancz.
- Baron, J., & Spranca, M. (1997). Protected values. *Organizational Behavior and Human Decision Processes*, 70, 1–16.
- Blackburn, S. (1984). *Spreading the word*. Oxford: Oxford University Press.
- Brink, D. (1986). Externalist moral realism. *Southern Journal of Philosophy*, 24(Suppl.), 23–42.
- Cheng, P. W., & Holyoak, K. J. (1989). Pragmatic reasoning schemas. *Cognitive Psychology*, 17, 397–416.
- Cosmides, L. (1989). The logic of social exchange: Has natural selection shaped how humans reason? Studies with the Wason selection task. *Cognition*, 31, 187–276.
- Darley, J. M., & Shultz, T. R. (1990). Moral rules: Their content and acquisition. *Annual Review of Psychology*, 41, 525–556.
- Forsyth, D. R. (1980). A taxonomy of ethical ideologies. *Journal of Personality and Social Psychology*, 39, 175–184.
- Forsyth, D. R. (1981). Moral judgment: The influence of ethical ideology. *Personality and Social Psychology Bulletin*, 7, 218–223.
- Forsyth, D. R., & Berger, R. E. (1982). The effects of ethical ideology on moral behavior. *Journal of Social Psychology*, 117, 53–56.
- Gabennesch, H. (1990a). The perception of social conventionality by children and adults. *Child Development*, 61, 2047–2059.
- Gabennesch, H. (1990b). Recognizing conventionality: Reply to Shweder and Helwig et al. *Child Development*, 61, 2079–2084.
- Griggs, R. A., & Cox, J. R. (1982). The elusive thematic-materials effect in Wason's selection task. *British Journal of Psychology*, 73, 407–420.
- Haidt, J. (2001). The emotional dog and its rational tail: A social intuitionist approach to moral judgment. *Psychological Review*, 108, 814–834.
- Haidt, J. & Graham, J. (in press). When morality opposes justice: Conservatives have moral intuitions that liberals may not recognize. *Social Justice Research*.
- Hare, R. M. (1952). *The language of morals*. Oxford: Oxford University Press.
- Harman, G. (1975). Moral relativism defended. *Philosophical Review*, 84, 3–22.
- Helwig, C. C., Tisak, M. S., & Turiel, E. (1990). Children's social reasoning in context: Reply to Gabennesch. *Child Development*, 61, 2068–2078.
- Hunter, J. (1991). *Culture wars: The struggle to define America*. New York: Basic Books.
- Johnson-Laird, P. N., & Byrne, R. M. J. (2002). Conditionals: A theory of meaning, pragmatics, and inference. *Psychological Review*, 109, 646–678.
- Johnson-Laird, P. N., Legrenzi, P., & Legrenzi, M. S. (1972). Reasoning and a sense of reality. *British Journal of Psychology*, 63, 395–400.
- Kahneman, D., & Lovallo, D. (1993). A cognitive perspective on risk taking. *Management Science*, 39, 17–31.
- Kant, I. (1959). *Foundations of the metaphysics of morals* (L.W. Beck Trans.). Indianapolis, IN: Bobbs-Merrill (Original work published 1786).
- Klauer, K. C., Musch, J., & Naumer, B. (2000). On belief bias in syllogistic reasoning. *Psychological Review*, 107, 852–884.
- Kohlberg, L. (1969). Stage and sequence: The cognitive-developmental approach to socialization. In D. A. Goslin (Ed.), *Handbook of socialization theory and research*. New York: Rand McNally.
- Kohlberg, L. (1981). *Essays on moral development (Vol. 1). The philosophy of moral development: Moral stages and the idea of justice*. San Francisco: Harper & Row.
- Mackie, J. L. (1977). *Ethics: Inventing right and wrong*. New York: Penguin.
- Maio, G. R., & Olson, J. M. (1998). Values as truisms: Evidence and implications. *Journal of Personality and Social Psychology*, 74, 294–311.
- Miller, J. B., & Bersoff, D. M. (1988). When do American children and adults reason in social conventional terms?. *Developmental Psychology* 24, 366–375.
- Nagel, T. (1970). *The possibility of altruism*. Princeton: Princeton University Press.

- Newstead, S. E., Pollard, P., Evans, J. St. B. T., & Allen, J. L. (1992). The source of belief bias effects in syllogistic reasoning. *Cognition*, *45*, 257–284.
- Nichols, S., & Folds-Bennett, T. (2003). Are children moral objectivists? Children's judgments about moral and response-dependent properties. *Cognition*, *90*, B23–B32.
- Piaget, J. (1965). *The moral judgment of the child*. New York: Free Press.
- Railton, P. (1986). Moral realism. *Philosophical Review*, *95*, 163–207.
- Sayre-McCord, G. (1986). The many moral realisms. *The Southern Journal of Philosophy*, *24*(Suppl.), 1–22.
- Schwartz, S. H. (1992). Universals in the content and structure of values: Theoretical advances and empirical tests in 20 countries. In M. P. Zanna (Ed.), *Advances in experimental social psychology* (Vol. 25, pp. 1–65). New York: Academic Press.
- Shantz, C. U. (1982). Children's understanding of social rules and the social context. In F. C. Serifica (Ed.), *Social cognitive development in context* (pp. 167–198). New York: Guilford.
- Shweder, R. A. (1990). In defense of moral realism: Reply to Gabennesch. *Child Development*, *61*, 2060–2067.
- Singer, P. (1972). Famine, affluence, and morality. *Philosophy and Public Affairs*, *1*, 229–243.
- Smith, M. (1994). *The moral problem*. Oxford: Blackwell.
- Snare, F. (1992). *The nature of moral thinking*. London: Routledge.
- Sturgeon, N. (1985). Moral explanations. In D. Copp & D. Zimmerman (Eds.), *Morality reason and truth* (pp. 49–78). Totowa, NJ: Rowman & Allenheld.
- Tetlock, P. E. (2003). Thinking the unthinkable: Sacred values and taboo cognitions. *Trends in Cognitive Sciences*, *7*, 320–324.
- Tisak, M. S., & Turiel, E. (1988). Variation in seriousness of transgressions and children's moral and conventional concepts. *Developmental Psychology*, *24*, 352–357.
- Turiel, E. (1978). Social regulations and domains of social concepts. In W. Damon (Ed.), *New directions for child development. Vol. 1. Social cognition* (pp. 45–74). New York: Gardner.
- Turiel, E. (1983). *The development of social knowledge*. Cambridge, UK: Cambridge University press.
- Turiel, E., Hildenbrandt, C., & Wainryb, C. (1991). Judging social issues: Difficulties, inconsistencies, and consistencies. *Monographs of the Society for Research in Child Development*, *56*, v–103.
- Williams, B. (1972). *Morality: An introduction to ethics*. New York: Harper & Row.
- Williams, B. (1985). *Ethics and the limits of philosophy*. Harvard: Harvard University Press.