

The Perceived Objectivity of Ethical Beliefs: Psychological Findings and Implications for Public Policy

Geoffrey P. Goodwin · John M. Darley

Published online: 17 December 2009
© Springer Science+Business Media B.V. 2009

Abstract Ethical disputes arise over differences in the content of the ethical beliefs people hold on either side of an issue. One person may believe that it is wrong to have an abortion for financial reasons, whereas another may believe it to be permissible. But, the magnitude and difficulty of such disputes may also depend on other properties of the ethical beliefs in question—in particular, how objective they are perceived to be. As a psychological property of moral belief, objectivity is relatively unexplored, and we argue that it merits more attention. We review recent psychological evidence which demonstrates that individuals differ in the extent to which they perceive ethical beliefs to be objective, that some ethical beliefs are perceived to be more objective than others, and that both these sources of variance are somewhat systematic. This evidence also shows that differences in perceptions of objectivity underpin quite different psychological reactions to ethical disagreement. Apart from reviewing this evidence, our aim in this paper is to draw attention to unanswered psychological questions about moral objectivity, and to discuss the relevance of moral objectivity to two issues of public policy.

1 Introduction

A central question in the study of ethics is the extent to which a belief about some ethical¹ matter represents a true, objective fact about the world, as opposed to a

¹For the purposes of this paper we use the terms “moral” and “ethical” interchangeably.

G. P. Goodwin (✉)
Department of Psychology, University of Pennsylvania, 3720 Walnut St., Solomon Lab Bldg.,
Philadelphia, PA 19146, USA
e-mail: ggoodwin@psych.upenn.edu

J. M. Darley
Department of Psychology, Princeton University, Princeton, NJ, USA

merely subjective preference. As we will review below, this “meta-ethical” question has occasioned considerable disagreement amongst philosophers. However our review will not be oriented to advancing the consideration of the usual issues among philosophers. Instead we propose to investigate whether there is agreement—or disagreement—on these issues among ordinary members of a culture. Comparatively little is known about this, the “psychology of meta-ethics”. How do individuals without philosophical training think of their ethical beliefs? Do they think of them as facts or preferences? How do they react to ethical disagreement? What consequences stem from holding a particular view of an ethical belief? In this paper, we review work conducted by ourselves and other researchers which provides some answers to these questions.

We first discuss some of the philosophical background to the study of ethical objectivity, in order to discover which meta-ethical issues usefully could be explored within the belief systems of philosophically naïve individuals; but then we turn our attention to discussing existing psychological findings about meta-ethics. Towards the end of the paper, we discuss several implications this work has for issues of public interest, including the credibility of criminal laws and voting behavior.

2 Philosophical Background

Philosophers have been inclined to take different positions regarding the question of whether ethical beliefs are objective. Some philosophers have argued that there are no true moral facts, or that morality is not objective. On this view, moral beliefs are irreducibly subjective, akin to reports of subjective (or inter-subjective) experiences, or to expressions of personal preference or taste (e.g. Ayer 1936; Blackburn 1984; Greene 2002; Harman 1975; Mackie 1977; Williams 1985). Many other philosophers, however, have argued for a more objective position (e.g., Brink 1986; Kant 1959; Nagel 1970; Railton 1986; Smith 1994; Sturgeon 1985).

A complication in thinking about this issue is that questions about the objectivity of ethical beliefs have been made in a variety of different ways by moral philosophers. Philosophical debate in meta-ethics is organized by marking out a set of closely related yet distinguishable polarities (see Sayre-McCord 1986; Snare 1992). Ethical “objectivism”, “realism”, “cognitivism”, and “universalism” are related, yet different philosophical positions that lie at one end of a scale that tracks the broad concept of objectivity, whereas “subjectivism”, “anti-realism”, “non-cognitivism”, and “relativism” are the respective opposites of these positions. Precision about these terms is important, philosophically, although there is far from complete concordance about their meaning.

To proceed, we will first define how we use these terms. According to our usage, the distinction between universalism and relativism refers to a difference in the *scope* of a particular ethical claim, i.e., who it applies to. If I say that it is wrong for you to break a promise in a particular situation, then, if I am a universalist, I am also committed to the view that it is wrong for anyone else in the same situation to break the same promise. The set of features that need to be “the same” in this formulation is hard to specify, but it clearly does not include aspects of personal identity, i.e., if it is wrong for me to break the promise, it cannot be okay for you to break it, merely in

virtue of you being a different person than me. The distinction between universalism and relativism is usually made at the level of culture, i.e., in terms of whether something that is wrong in one culture is also wrong in another. But it can also be made at the interpersonal level, as above, i.e., in terms of whether something that is wrong for one person to do is also wrong for another person to do (see e.g., Sayre-McCord 1986; Williams 1972). Whichever way it is framed, this question is correctly interpreted as one about what is *actually* right or wrong (permissible, impermissible, good, bad, etc.), and is not simply a descriptive or sociological query about what the accepted views happen to be in different regions or eras (see Snare 1992).

The distinction between relativist and universalist ethics has been explored to some extent by psychologists (see Gabennesch 1990a for a review; and Forsyth 1980). However, other distinctions in meta-ethics are arguably more fundamental, and have remained largely unexplored by empirical psychologists. One issue concerns whether moral claims have a truth value. Cognitivism is the view that at least some moral propositions are capable of having a truth value (Sayre-McCord 1986; Snare 1992), whereas non-cognitivism, or emotivism, is the view that moral propositions are devoid of cognitive status—that is, they are not capable of having a truth value, and serve simply to express attitudes or emotions (Ayer 1936; Stevenson 1937). Some modern non-cognitivists (e.g., Blackburn 1993, 1998; Gibbard 1990, 2003) do not deny that moral claims have a truth value, although they do deny that a certain kind of truth applies to moral claims, i.e., truth that corresponds to the objective facts in the world (see further, Sinnott-Armstrong 2006).

On the other hand, both objectivists and subjectivists agree that moral claims can be literally true, but they diverge on what makes such claims true (Sayre-McCord 1986). Subjectivists believe that the truth of moral claims depends on the subjective states of individuals (Sayre-McCord 1986), such that the claim: “theft is wrong”, is true only when its speaker has the appropriate attitude of disapproval towards theft. This view can again be framed at the level of individuals—what is ethically true for an individual depends on the contents of their mental states; or at the level of a culture (so-called, “inter-subjectivism”)—what is ethically true for a group depends on some aggregation over the mental states of the group members. Objectivists, on the other hand, hold that the truth conditions of moral claims are mind-independent in the sense that a moral claim can be true without reference to the subjective states of the individual making the judgment, and without reference to the conventions of any group of people who are making the moral judgment (Sayre-McCord 1986).² Hence, these two meta-ethical distinctions intersect as Fig. 1 illustrates.

Two points need clarification. First, meta-ethical subjectivism as defined by Sayre-McCord (1986) and Snare (1992), is different from non-cognitivism in the following sense—the non-cognitivist takes moral claims to be mere *expressions* of attitudes, while the subjectivist takes them to be *reports* about subjective states or attitudes (and hence they are capable of being true or false). Thus, although both meta-ethical subjectivists and emotivists disagree with objectivists, the difference between them is akin to the distinction between someone who in response to a sharp jab says “that hurts” (a subjectivist, cognitivist position which is capable of having a

² This way of describing objectivism is itself not uncontroversial, however (see Putnam 1987).

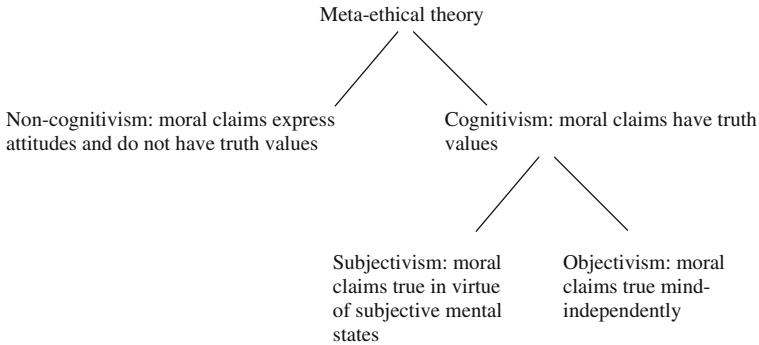


Fig. 1 Intersection of cognitivism, non-cognitivism and objectivism, subjectivism

truth value), as opposed to “ouch” (an emotivist, non-cognitivist position, which is not capable of having a truth value).

Second, although similar to relativism, meta-ethical subjectivism is different in the sense that subjectivism makes a claim about the *truth conditions* of moral claims (see Sayre-McCord 1986), whereas relativism makes a claim about the *scope of application* of ethical principles—who they are true for, rather than what makes them true. So for instance, one could be an objective relativist, by saying that what is ethically right for person A is different from what is right for person B, and that this is true independent of any subjective attitudes that persons A and B might have about the matter (see Snare 1992). In this case the *scope* of the claim about what is ethically right is relative, because it does not apply to everyone, but its *truth conditions* are purportedly objective.

Similarly one could be a subjectivist and a universalist, by claiming that although the *truth conditions* of moral claims are subjective mental states, it so happens that there is an underlying convergence of these subjective states (even if perhaps not consciously realized), and so moral claims *apply* universally. This view would be challenged by the fact of obvious ethical disagreements in practice, but one could counter-argue that such disagreements ultimately rest on disagreements over factual, non-ethical matters (for instance, about the meaning or consequences of a particular act), rather than on fundamental matters of ethical principle; or alternatively that disagreements arise over peripheral but not core ethical principles, and that the core ethical principles are universally shared (though subjective).

This picture is necessarily a simplification of meta-ethical theory. Other meta-ethical distinctions exist, such as between realists and anti-realists, and philosophers have often taken surprising meta-ethical positions. J. L. Mackie, for instance, developed a form of cognitivist subjectivism, which held that ethical beliefs have truth values, but that they are all false (Mackie 1977).³ It bears mention too that the terminology used here is by no means universally agreed upon. Williams (1972) for instance, regards non-cognitivism as a form of subjectivism, in contrast to both Snare (1992) and Sayre-McCord (1986).

³ Whether this falsehood is supposed to be a contingent or necessary matter is not entirely clear in Mackie (1977).

This dense philosophical landscape may seem impenetrable for psychological investigation. And indeed, it would be a somewhat pointless exercise to ask naïve participants to produce fine distinctions between sophisticated meta-ethical views. However, we argue that people do have a set of underlying intuitions about the objectivity or subjectivity of their ethical beliefs, and that these intuitions exist without prior explicit theorizing. That is, in addition to having a wide range of basic, or what philosophers refer to as “first order” moral beliefs, people also have “second order” moral beliefs. We surmise that these second order beliefs have a set of interesting and unique properties that cannot be captured simply by asking people how strongly they agree with a particular ethical claim, i.e., by asking them about their first order beliefs. But, if we want to understand the underlying psychology of ethical objectivity, we need ways to pose questions about the topic that are understandable to human participants without any philosophical training—that is, we need questions that are both philosophically and psychologically meaningful.

To that end, our empirical investigations have used two kinds of questions to distinguish objective and subjective positions about ethical beliefs. First, we ask questions that pertain to whether people take their ethical beliefs to be objectively true statements of fact, or alternatively, subjective preferences or attitudes. Hence in terms of Fig. 1, the contrast relevant for this paper is best approximated as that between the far bottom right, objectivist position, and either of the remaining two positions—either philosophical subjectivism or non-cognitivism. (This distinction reflects the way Williams (1972) distinguishes between subjectivist and objectivist positions.) Second, we ask participants whether in a situation of moral disagreement between two (or more) parties, one party or other must be mistaken. An objectivist must reply that at least one party is mistaken, whereas a subjectivist need not make this claim (Snare 1992).

3 Psychological Literature

Psychologists have tended not to be interested in this sort of question. And indeed, they have tended to focus very little on meta-ethical questions in general. We review what little research there is on meta-ethical questions here before reporting on our own investigations.

The psychological literature on moral judgment, at least during the twentieth century, was heavily influenced by the developmental work of Piaget (1965) and Kohlberg (1969, 1981). Both were strongly cognitivist in orientation, and sought to understand the reasoning processes that underlie moral judgment, which puts them at odds with more modern affective approaches to moral judgment (see later). Kohlberg proposed an ambitious stage theory of moral development. His dominant research method was to present vignettes portraying various moral dilemmas, which involved a conflict between acts that either complied with sociolegal rules, or else violated these rules in order to serve human needs or welfare (Rosen 1980, p.69). Children’s moral development was assessed by the reasoning processes they exhibited in reaching a conclusion, rather than by the specific conclusions they reached (see Rosen 1980). Kohlberg’s classificatory scheme comprised three overall stages each with two sub-stages into which any individual’s mode of moral

reasoning could be placed (Kohlberg 1963, 1976). These three levels, in order of increasing moral sophistication, were called the *concrete individual perspective*, the *member-of-society perspective*, and the *prior-to-society perspective* (see Rosen 1980). This model was both controversial and influential (Rosen 1980; Darley and Shultz 1990).

Kohlberg's model is concentrated on normative moral judgments, and offers few resources for exploring questions about meta-ethics. Although he preferred an "active organism" model of moral development, in which a person constructs and then reconstructs their moral knowledge (Rosen 1980), there is little evidence that Kohlberg (himself a moral realist and objectivist), regarded the pursuit of meta-ethical questions as a part of this active, constructive process.

The highest stage—stage 6 of Kohlberg's scheme—is described in terms of the recruitment of universal ethical principles (Rosen 1980, p. 80). The moral agent at this stage "is completely de-centered from society's expectations and bases his resolutions to ethical conflicts upon universal principles of justice which are prescriptively consistent without exception" (Rosen 1980, p.80). This characterization implies that the seeking of just solutions to moral problems, though not necessarily straightforward, is ultimately a resolvable problem (cf. Nagel 1979). Moral problems have objectively correct answers, and the moral agent of stage 6 can discover what these answers are. This is a strong position, because not even all objectivists agree that all moral dilemmas are indeed resolvable.⁴

More tellingly, Kohlberg's scheme characterizes the pursuit of meta-ethical questions as a developmental aberration. Having erected his scheme, Kohlberg and his collaborators observed a phenomenon which they had not previously paid much attention to. Moral reasoners (typically boys) who had approached stage 4 in their moral development appeared to be regressing to an earlier, egoistic stage of moral development (see Rosen 1980, pp. 89–90). This stage, characterized by relativistic and skeptical moral thought, was interpreted as a "rebellious moratorium", in preparation for a restoration and deepening of the higher levels of moral thinking. Initially this posed a challenge for Kohlberg's linear, deterministic theory, which did not permit backwards movement through the stages. However, theoretical resolution was achieved by classifying this phase as a transitional stage 4B, in which one "is on the threshold of a breakthrough to principled moral reasoning which, although he does not yet realize it, will supply him with the new standard for which he is striving" (Rosen 1980, p.92).

This is the only place in Kohlberg's scheme that allows for moral doubt or skepticism. Meta-ethical reflection plays no role in the higher levels (5 and 6) of his scheme, and is instead confined to a stage of adolescent rebellion. But, the possibilities are greater than this. Unlike Kohlberg's ethical skeptics (who tend also to be egoists in his mind), not all ethical skeptics are egoists. A full characterization of moral thinking needs to incorporate the possibility of a mature, and not amoral, ethical skepticism.

⁴ Thomas Nagel, an ethical objectivist (see Nagel 1997) argues that there may be situations where various moral values are incommensurate, and there is decisive support for two or more incompatible courses of action (Nagel 1979). Although we may look for a "single scale on which all these apparently disparate considerations can be measured, added, and balanced" (Nagel 1979, p.131), there may be no such scale.

Meta-ethical questions have been investigated in the child development literature. A range of studies have explored children's perceptions of social conventions, and in particular, whether or not children are able to distinguish social and ethical conventions (see review by Gabennesch 1990a). This research has shown that both children and adults can distinguish ethical prescriptions from social conventions in terms of seriousness (Shantz 1982), alterability (Tisak and Turiel 1988), and dependence on cultural norms (Turiel 1978; Nichols and Folds-Bennett 2003; although both Nichols (2002, 2004) and Kelly et al. (2007) call the distinction between ethical and conventional prescriptions into question. Participants in Turiel's studies were asked questions such as whether a transgression would still be wrong if there was no rule to prohibit it, or whether it would still be wrong in another culture that has no prohibition on it. Such questions are undeniably meta-ethical, but they concern the *scope* of application of an ethical principle, i.e., whether it is universal or relative, not the nature of its *truth conditions*. While it is likely true in practice that universalists are more likely to be objectivists than are relativists, there is no logical requirement that this be the case.

Consider the question whether it would still be wrong to commit a transgression even if there were no rule to prohibit it. Answering "yes" to this question does not automatically make one an objectivist. A "yes" response implies that the warrant for the principle lies deeper than the mere fact that a rule has been enacted—the principle is taken to be more than a mere convention, as Turiel and others have argued (although this interpretation is controversial, see Gabennesch 1990a, b; Shweder 1990; Helwig et al. 1990). But, this deeper warrant could be either subjective or objective—it could arise from the subjective mental states of members of a community (subjective), or it could arise from an external source, such as God (objective), or from some other objective source. Indicating that the existence of a rule is not sufficient to ground moral commitments does not commit a person to being either a subjectivist or an objectivist.

Similarly, consider the question whether it would be wrong for another culture to allow something that in our culture is prohibited, i.e., whether a particular moral claim is wrong "for some people" or "for real" (Nichols and Folds-Bennett 2003). A complication exists with this question in that "for some people" and "for real" are not true opposites—the question is thus somewhat ambiguous as to whether it is asking about universalism or objectivism. On one reading of this question, "for real" is equivalent to "for all people", and thus answering "for real" indicates a universalist position, but not necessarily an objectivist one—a subjectivist might want to apply a moral claim universally, without being committed to its having an objective foundation. On another reading of this question, the question is purely descriptive, and asks about the presence of conflicting attitudes within society. One final interpretation is that the question does ask about objectivity—the phrase "for some people" denotes a kind of "non-realness" precisely because it is contrasted with "for real". On this reading, the question is similar to the ones we have asked our participants. But, without further specification the question is ambiguous.

This is not intended as a decisive criticism of the approach adopted by Nichols and Folds-Bennett (2003). They were interested in the moral responses of 4–6 year old children, and the question they asked, despite its complications, seems simpler and thus more appropriate for children in that age group, who are not likely to be

troubled by the above considerations. Their chief finding was that children treat moral transgressions as universally wrong, but do not treat conventional transgressions in this way.

Attempts to investigate adult meta-ethical views have also tended to focus on perceptions of the universality of ethical claims. In a series of studies, Donelson Forsyth and his collaborators (Forsyth 1980, 1981; Forsyth and Berger 1982; Forsyth and Pope 1984) explored what he terms “ethical ideologies”, which are based on the intersection of two variables, *relativism* and *idealism*, as measured by his Ethics Position Questionnaire (Forsyth 1980). Forsyth’s resulting taxonomy was found to predict the severity of moral judgments (Forsyth 1981; Forsyth and Berger 1982), and although it was found not to predict moral behavior in a setting where the temptation to cheat was presented, it did predict differences in participants’ reactions to their own moral transgressions (Forsyth and Berger 1982).

Forsyth’s questions regarding relativism included items such as the following:

14. Different types of moralities cannot be compared as to ‘rightness’.
15. Questions of what is ethical for everyone can never be resolved since what is moral or immoral is up to the individual.
16. Moral standards are simply personal rules which indicate how a person should behave, and are not to be applied in making judgments of others.

Such items perhaps do capture meta-ethical beliefs, but they go no further than the child development literature in this regard. The principal distinction is again between relativism and universalism. And although questions of the universality of ethical claims are undoubtedly within the scope of meta-ethics, they do not directly map onto questions about the objectivity of ethical claims.

The second dimension of Forsyth’s (1980, 1981) taxonomy is puzzling. In his Ethics Position Questionnaire (Forsyth 1980), ten statements (which participants rate their level of agreement with), are used to assess the construct of “idealism”. Two such examples are:

1. A person should make certain that their actions never intentionally harm another even to a small degree.
5. One should not perform an action which might in any way threaten the dignity and welfare of another individual.

Eight further statements assess ethical idealism in a similar way. The odd characteristic of these statements is that they seem to assess beliefs about the way the world works rather than any sort of ethical ideology. Many small acts of intentional harm occur daily, often without any ethical component: a dentist inflicts intentional harm on a patient when probing sensitive teeth (see e.g., Bucciarelli et al. 2008). The possibility of avoiding all of these minor harms is quite remote. It would thus seem more accurate to think of this “idealism” scale as characterized by naivety at one end and world wisdom at the other. And it is unclear how this dimension assesses an important component of ethical thought.

The bottom line from the preceding considerations is that questions about universalism, although likely to produce answers that correlate highly with questions about objectivity, are not identical to those questions. Our own research has concentrated on questions about objectivity: to what extent do people regard their

moral beliefs as objective statements of fact, as opposed to mere subjective preferences? To that end, in a recent paper (Goodwin and Darley 2008), we asked participants two different sorts of question in order to get at the perceived objectivity of their ethical (and other sorts of) beliefs: whether they thought that there could be a correct answer as to whether a particular ethical belief was true (or alternatively, in Experiment 1, whether they thought that a particular belief that they held was true, as opposed to being an opinion or attitude); and whether they thought that a person who disagreed with them about the belief was mistaken, or whether instead neither party need be mistaken in the face of disagreement. Our strategy was to average responses to these questions together to create a composite index of ethical objectivity (for further details, see Goodwin and Darley 2008).

These three questions track some of the central philosophical thinking on the topic of ethical objectivity. Philosophers have been concerned with whether ethical statements possess truth values, and with whether correct answers exist to ethical questions.⁵ And one common philosophical criterion for whether a belief is objective or not is whether disagreement over it implies that one of the disagreeing parties is mistaken (see e.g., Smith 1994; Snare 1992). Moreover, both pilot testing and intuition suggested that the questions we asked are readily understandable, and that experimental participants are able to produce answers to them.

Before embarking on this research, it was an open question how people would perceive the objectivity of their ethical beliefs. Psychological research on “naïve realism” has shown that people typically treat their own perceptions and preferences as fundamentally “correct” (Ross and Ward 1995). Extending this perspective to the ethical domain, it follows that most people should be highly objective about their ethical beliefs (and in fact, any sort of belief). However, other considerations suggest the opposite prediction: that people will self-consciously appreciate the socially constructed nature of ethical beliefs, and thus be moral subjectivists. Chief among these considerations is the observation made by several social commentators that we live in a more morally contested and possibly relativistic society than we once did (e.g., Hunter 1991). These broad social trends may thus be reflected in the perception that moral issues are ultimately subjective matters.

Our initial research yielded the following three central findings (Goodwin and Darley 2008), which are contrary to both of these perspectives. We concentrated primarily on canonical examples of moral transgressions that involved violations of justice or the infliction of harm (see Haidt 2007). The results of several experiments indicated that people take beliefs about the wrongness of such transgressions to be quite objective—almost as objective as beliefs about matters of everyday or scientific fact, and more objective than beliefs about social conventions (i.e., appropriate dress wear and manners) and beliefs about matters of taste. That is, they tend to indicate that such ethical beliefs are true, that there is a correct answer as to their truth, and that in the face of disagreement, one or other party is mistaken. This

⁵ A note about our use of the term “belief”: although it is the standard philosophical position, we do not intend the term “belief” to imply that a person assents to the truth of the proposition they believe (see e.g., Lacey 1996)—that would presuppose an answer to the question we want to investigate. Rather, we intend a weaker, more psychological reading of the term belief, which denotes an attitude of agreement or assent, without a further commitment to truth (see e.g., Reber 1996).

finding was robust when controlling for how strongly individuals agreed with beliefs from each of these classes.

However, individuals differed considerably in the degree of objectivity they ascribed to such beliefs. The majority of individuals were quite objective, but there were some individuals who were not. One pertinent predictor of ethical objectivity was whether a person thought that ethical beliefs are ultimately grounded by a religious deity—those individuals who thought this was the case tended to be more objective than those who did not. Moral objectivity was particularly strong among individuals who did not think that right and wrong acts are possible without the existence of a divine being.

There was also some indication from these initial experiments that the specific content of ethical beliefs makes a difference in terms of how objective such beliefs are perceived to be. For instance, people tended to treat a statement about the wrongness of consciously discriminating against another person on the basis of race as more objective than a statement about the goodness of anonymously donating 10% of one's income to charity. This finding contrasts with a presupposition that runs strongly through some philosophical writing—that ethical beliefs as a whole are objective or subjective, and that one's meta-ethical view should apply en masse to one's entire set of ethical beliefs, in a top-down, deductive fashion—what Sinnott-Armstrong (2009) refers to as the *uniformity assumption*.⁶ Clearly, our participants did not accord with this description.

This initial research thus partially supported a naïve realist account of ethical beliefs—people generally found ethical beliefs about canonical moral transgressions to be highly objective. But there are two features of our data which cannot be accounted for by a simple naïve realist account. First, not everyone treats their ethical beliefs as objective—some participants were highly subjective even about canonical moral transgressions. Second, not all ethical beliefs are treated equally, and there appears to be considerable variation in perceived objectivity, depending on the kind of ethical belief in question.

One caveat from this first investigation is that we only surveyed samples of undergraduate and graduate student participants from a single American university. Certainly this is a limitation. But, given the nature of our findings we do not consider it to be a critical one, because we suspect that in a broader sample, at least two of our main results would be amplified. A broader sample of Western participants is likely to include many who have been much less exposed to skeptical arguments about the nature of morality and religious belief than the students in our samples. We would thus expect that perceptions of moral objectivity would be stronger than in our sample, and that the religious grounding of ethics would also be a stronger predictor of ethical objectivity. It would, of course, be of interest to investigate this, and also to investigate whether perceptions of objectivity are stronger in Western cultures than in others (see Mackie 1977, for discussion of how notions of objectivity permeate Western moral discourse).

In more recent research we have followed up three questions left unanswered by our initial investigations: (1) Why are some ethical beliefs treated as more objective

⁶ Although see Gill (2008, 2009) for philosophical critique of this uniformity assumption, and for a defense of meta-ethical variability.

than others? (2) Why are some people more objective about their ethical beliefs than others? (3) Does holding an objectivist or subjectivist position about an ethical belief have any important psychological or behavioral consequences?

4 Why Are Some Ethical Beliefs Treated More Objectively Than Others?

The initial evidence in our 2008 paper showed that people do not treat all ethical beliefs as being equally objective. Why is this? What might predict whether a moral belief is treated as highly objective, or not? In unpublished work, we have investigated three specific hypotheses (Goodwin and Darley 2009).

The first hypothesis is that people will be more objective about transgressions that involve the infliction of direct personal harm or injustice. Jon Haidt and his colleagues (e.g., Haidt and Graham 2007; Graham et al. 2009) have argued that such transgressions are the only ones consistently regarded as immoral in Western societies. Other actions that do not cause physical harm or injustice, but that provoke disgust or contempt, are regarded as moral violations only by more conservative (Haidt and Graham 2007; Graham et al. 2009) and lower class Westerners (Haidt et al. 1993). Other research has indicated that only harm and justice based transgressions are considered to be moral violations in all cultures (Shweder 2003; Shweder et al. 1997; although some cultures judge as immoral at least some actions that generate emotions of contempt or disgust). Taking a less empirical, more philosophical perspective, Royzman et al. (2009), summarize the arguments that harm is the fundamental basis of moral judgments. Each of these perspectives suggests that people will be highly objective only about harm or justice-driven transgressions.⁷

The second hypothesis is that people will regard the negativity of moral transgressions (i.e., their wrongness or badness) as more objective than the positivity of morally exemplary acts (i.e., their rightness or goodness). One motivation for this hypothesis stems from the fact that legal codes are primarily concerned with prohibiting rather than recommending action. This may, in part, reflect an important underlying moral principle, i.e., that clear and objective moral principles can be formulated about what is bad or wrong to do, but not about what is good or right.

The third hypothesis is that people will be more likely to view a moral belief as objective to the extent that they think other people hold that belief—that is, they are influenced by the degree of consensus they perceive to hold about the belief.

To investigate these hypotheses, we developed an experiment in which participants were asked to rate the objectivity of 12 different moral beliefs, as well as some factual beliefs and some beliefs about social conventions. The moral beliefs fell into four different classes. One class consisted of beliefs about the wrongness of causing harm or injustice—assault, cheating, and the provision of a false alibi. The second class consisted of beliefs about the wrongness of symbolic harms—desecrating a 9/11 memorial, performing a Nazi salute to a Jewish sporting audience, and flag burning. The third class consisted of beliefs about the goodness of morally exemplary acts—donating money to charity, performing a risky swim rescue, and performing onerous energy

⁷ Note, however, that other arguments exist which call into question whether harm is the fundamental basis of morality (see e.g., Kelly et al. 2007).

conservation activities. And the fourth class consisted of beliefs about the permissibility of highly contested value of life issues—abortion, and two different instances of euthanasia. The contrast between the first (harm, injustice) and second (symbolic harm) classes allows an assessment of whether direct harm or injustice predicts greater objectivity. The contrast between the first two classes (transgressions) and the third class (morally exemplary acts) allows an assessment of whether people ascribe greater objectivity to moral transgressions than they do to morally exemplary acts. And the contrast between the first three classes and the fourth class (contested value of life issues) allows an assessment of whether people are more objective about beliefs that they suspect are widely shared. To investigate this question more precisely, we also asked participants to estimate the percentage of US citizens they thought would agree with each belief.

The three items within each category hung together well, both in terms of the patterns of means and inter-correlations (see Goodwin and Darley 2009). And, importantly, all of the ensuing analyses controlled for how strongly participants agreed with each belief. As in our previous studies, participants were highly objective about the wrongness of inflicting harm and injustice. However, they were no more objective about such transgressions than they were about transgressions which inflicted much more symbolic wrongs. This result was surprising. It does not rule out direct, physical harm or injustice as important predictors of perceived objectivity, but it does show that their presence is not necessary for an ethical belief to be treated as highly objective. It is of course possible that our participants thought that the symbolic wrongs would inflict emotional harm of equivalent magnitude to the harms caused by the directly harmful or unjust actions. This would be useful to investigate. However, even if this did turn out to be true, it would not undercut our main conclusion: our results would still show that the nature of the harm or injustice that a moral transgression causes need not be direct or physical, in order for the wrongness of that act to be perceived as highly objective.

Participants were also more objective about the wrongness of both kinds of transgression (direct harms or injustices, as well as symbolic harms) than they were about the goodness of the morally exemplary actions. This difference was highly significant for each of the items we investigated (there was no overlap between any of the means from the two different categories: good vs. wrong). This suggests that wrongness is indeed perceived to be a more objective moral property than goodness, even when controlling for how strongly participants agreed with the beliefs in the two categories. This finding, however, might have occurred for two different reasons. On the one hand, it could be that, as we suggest, the valence of the behavior (whether it is positive or negative) is what accounts for the difference in perceived objectivity. But, it might also be that the crucial difference is one between statements of value (goodness, badness) and statements of norms (rightness, wrongness). Our finding could be explained by the idea that people take claims about rightness and wrongness to be more objective than claims about goodness and badness, which is why they treated *morally wrong* as a more objective category than *morally good*. This seems implausible for positively valenced items, i.e., it seems unlikely to us that people would treat rightness as a more objective category than goodness, at least when it is predicated of the same action. The claim that some action is morally right makes it more obligatory than the claim that it is merely morally good (see e.g., Gert 2005, p. 321–322), and we thus suspect that a claim about rightness is less likely to be perceived as objective than a claim about goodness. But, we cannot as confidently rule out this explanation in the

negative domain—perhaps people do indeed treat wrongness as more objective than badness.⁸ We are thus currently conducting an investigation to determine whether valence or moral language is the chief driver of the effects described in the present study.⁹

Finally, our participants treated the goodness of morally exemplary actions as more objective than the permissibility (or impermissibility, depending on how they responded) of highly contested value of life choices. In fact, they were no more objective about the wrongness of social conventions than they were about the permissibility (or impermissibility) of the value of life choices. Further analyses, which used participants' estimated consensus responses for each moral belief, showed that such estimates were very reliable predictors of perceived objectivity (our results demonstrated an across-items correlation between objectivity and perceived consensus of $r=.84$), over and above how strongly participants agreed with each of the beliefs in question. Because they are correlational, these data are of course consistent with an alternative interpretation according to which perceived objectivity inflates consensus estimates, i.e., it is objectivity that drives consensus judgments rather than the reverse. Our current data is not sufficient to rule out this hypothesis, and indeed it may well be true. Perhaps both directions of causation operate. At any rate, in order to draw a causal inference from consensus assessments to objectivity, a study which manipulates perceived consensus estimates is called for.

These results suggest three reasons why some ethical beliefs are treated as more objective than others. First, although harm and injustice are likely to be important predictors, their occurrence is not necessary in order for a moral belief to be perceived as highly objective. Second, people seem to treat beliefs about the wrongness of moral transgressions as more objective than beliefs about the goodness of morally exemplary actions. And third, perceived consensus appears to predict objectivity assessments. These conclusions are preliminary at this point; all three results invite follow-up investigations, which we are currently carrying out.

5 Systematic Individual Differences in Ethical Objectivity

The previous study established that there is systematic variation in objectivity assessments across different kinds of ethical belief. In our previous studies, we had also shown that there were some systematic individual differences in objectivity.

⁸ We thank Walter Sinnott-Armstrong for this point. Gert (2005, p. 322) makes the point, which we agree with, that the terms “morally good” and “morally right” tend to refer to different kinds of action, whereas “morally bad” and “morally wrong” are used more interchangeably. We interpret this point as counting in favor of our valence hypothesis regarding negative moral actions—it suggests to us that there will not be a substantial difference in judgments of objectivity regarding the categories *morally bad* and *morally wrong*. Gert himself, however, takes wrongness to be a more objective property than badness (see p. 325). For positive moral actions, Gert's point about the non- interchangeability between the terms “good” and “right” suggests that in order to clearly test whether goodness is considered less objective than rightness, these moral properties may need to be predicated of *different* actions.

⁹ A further possible explanation of our finding, suggested by an anonymous reviewer, is that people may think that it is more likely that people would do morally good (or right) things for the wrong reasons, than that they would do morally bad (or wrong) things for the right reasons. This is an intriguing possibility, and worth investigating.

Those who grounded their ethical beliefs in religious ideology were more objective than those who did not. This finding is correlational, and it would be desirable to show that there is a causal influence at work here. This might be done experimentally by manipulating the salience of religious ideas, and showing a boost in objectivity scores (among religious participants) when religious ideas are made salient. While this finding would be important, other individual differences may also predict judgments of objectivity, and these have been the focus of our most recent investigations. In particular, we have been interested in whether certain cognitive and personality variables predict judgments of moral objectivity.

One striking observation from our previous work concerned the kinds of explanations participants gave for why disagreement may have arisen about an ethical issue. In those studies (Goodwin and Darley 2008), we asked participants to explain why disagreement may have arisen in each case. The differences in the kinds of explanations people offered were striking. Some participants, typically the objectivist responders, seemed to be less inclined to *explain* why the disagreement may have arisen. Rather, they tended to make blanket responses which either reiterated their own belief, which expressed disbelief that another person could have disagreed with them, or which ascribed some moral defect to the disagreeing other person. For instance, in response to a disagreement over whether robbery to pay for a vacation is wrong, they would say things about the person who disagreed with them like:

The other person obviously has no values, and no respect for the property of others!

or:

The other person has a warped view of what is a moral action.

By contrast, those who responded more subjectively seemed more interested in explaining the disagreement. They would say things like:

Values are probably the source of disagreement. A hedonist would feel morally obliged to steal for the vacation. A compassionate person would not.

or:

The source of disagreement is most likely how the other person and I derive our morals. It could be that he is a radical hedonist while I take a more utilitarian approach.

The subjectivists' responses seemed to place more emphasis on explaining the source of the disagreement in terms of competing moral values, whereas the objectivist responses tended to explain the disagreement in terms of the sheer presence or absence of moral values at all. And indeed, as rated by an independent coder, there were reliable differences in the quality and kind of explanations that subjectivists and objectivists gave (Goodwin 2009).

These results suggest that part of what may underlie objectivist responding to an ethical disagreement is a tendency *not* to think of some of the alternative reasons for doubting or disagreeing with the belief in question. If a person actively considers some of the reasons for doubting an ethical belief they hold, they may be inclined to be less objective about that particular belief. To investigate this idea, we presented

participants with a task that measures disjunctive reasoning ability, i.e., the tendency to actively unpack alternative possibilities when reasoning (Goodwin 2009). We used the following “five blocks” task originally devised by Levesque (1986, 1989) and investigated further by Toplak and Stanovich (2002) and Over et al. (in press):

There are five blocks in a stack, where the second one from the top is green and the fourth is not green. Is a green block definitely on top of a non-green block?

Participants have to choose between three response options: yes, no, cannot tell. The most intuitively plausible response is to say “cannot tell”, but this is incorrect. To infer the correct solution you need to consider the alternative possibilities for the indeterminate third block—you need to unpack the two distinct alternatives for this block and examine their consequences. The correct answer is “yes”: if the third block is green, then since it is directly on top of the fourth non-green block, the condition is fulfilled; alternatively, if the third block is not green, then since it is directly underneath the second green block, the condition is also fulfilled. Thus, whichever color the third block happens to be, there is definitely a green block directly on top of a non-green block.

Toplak and Stanovich (2002) examined this task along with a range of other tasks which are aimed at assessing “disjunctive thinking”. These tasks correlate only moderately with indicators of cognitive ability such as SAT scores, which led Toplak and Stanovich to describe them as measures of a dispositional tendency towards disjunctive thinking rather than as cognitive ability measures. In our initial experiment (Goodwin 2009), we used only the five block task because it has no moral content, and because it seemed to capture the difference in the kind of thinking that underpinned ethically objectivist and subjectivist responses in our free response task. Our hypothesis was that participants who performed the five blocks task correctly would tend to be less objective about their ethical beliefs than participants who did not. Such individuals should possess a dispositional tendency to think disjunctively, and thus, in a moral context, they should be more inclined to actively consider the reasons why another person might disagree with their ethical beliefs. However, they should not be more objective about matters of taste, where disagreement can exist at the level of brute preferences rather than reasons.

We also presented participants with a range of other measures that we suspected might predict objectivity: Frederick’s (2005) Cognitive Reflection Test, which is a simple, three measure instrument which measures a person’s tendency to over-ride an immediate, intuitive response, and which correlates highly with general intelligence; Cacioppo et al.’s (1984) Need for Cognition scale, which measures a person’s “tendency to engage in and enjoy effortful cognitive endeavors”; and Kruglanski et al.’s (1993; Webster and Kruglanski 1994) Need for Closure scale, which measures a person’s “need for an answer on some topic, any answer as opposed to confusion and ambiguity”.

The study demonstrated that only the disjunctive thinking measure reliably predicted ethical objectivity. Those participants who performed this task correctly were significantly less objective on ethical items than those who performed it incorrectly (we observed a correlation between performance on this problem and ethical objectivity of $r=-.36$). But they were no more objective than the incorrect

responders on the items about matters of taste. In contrast, and somewhat to our surprise, none of the other measures reliably predicted ethical objectivity. This initial result lends credence to the idea that a propensity towards disjunctive thinking inclines people towards a more subjective view of their ethical beliefs.

A follow-up study presented participants with a wider array of disjunctive reasoning tasks that were drawn from Toplak and Stanovich (2002). One of these tasks was isomorphic to the three blocks task presented earlier and the other was closely related to it. The average performance on these items correlated negatively with objectivity assessments, replicating the previous finding. We again measured Need for Cognition, Need for Closure, and Cognitive Reflection, and these measures were not correlated with objectivity. The experiment excluded one extra “mindset” variable. When answering the objectivity questions, half of the participants were instructed to adopt an “experiential” mindset—they were instructed to respond with their “first, natural response” and that the experiment was only interested in investigating people’s “gut level reactions” (see Epstein 1991, 1994). In contrast, the remaining participants were instructed to be “as rational and analytic as possible” (Epstein 1991, 1994). This manipulation did not produce a difference in the overall level of objectivity scores, but it did produce a marked difference in the pattern of observed correlations. The correlation between the disjunctive thinking measure and objectivity was much stronger in the “experiential” condition ($r = -.59$), and was not reliable at all in the “rational” condition ($r = .05$). We surmise that the experiential instructions reflect the normal way in which participants respond to these sorts of question, which is why this condition replicated the strong correlation observed in our first experiment. However, this result was not predicted, and the exact reason for the difference between the two conditions is not clear. One possibility is that the instruction to think rationally pushed participants who would not normally consider the alternative reasons for disagreeing with an ethical belief, to do so. However, the overall mean objectivity score in the rational condition was not lower than it was in the experiential condition, which does not support this explanation.

Regardless, these correlational results suggest that there is an underlying cognitive variable—perhaps the tendency to reason disjunctively—which underpins the way people make assessments of ethical objectivity. Our current working hypothesis is that individuals who are inclined to think disjunctively on a problem-solving task, are also more inclined to generate alternative reasons why another person might doubt or disagree with an ethical belief that they themselves hold. Sometimes these alternative reasons may have persuasive appeal—thus, on average, this kind of thinking is likely to lower assessments of ethical objectivity. Thus, in our current research, we are experimentally inducing participants to first think of alternative reasons for doubting or disagreeing with an ethical belief before asking them to respond to questions about the objectivity of that belief—predicting a drop in perceived objectivity under these conditions, compared with a control condition.

6 Consequences of Taking an Objective Position

One of our underlying motivations for studying the psychology of meta-ethics was to examine the sorts of consequences that stem from holding a particular meta-

ethical view. Do objectivists and subjectivists differ in how they respond to ethical disagreement? Philosophers have often written as though, in the face of ethical disagreement, objectivists ought to be more open to listening to the arguments of the disagreeing other party (e.g., Snare 1992). If you are an objectivist, you believe that there is a true fact of the matter to be discovered about any ethical belief. As a consequence, when disagreement arises, you should be interested to hear whether the disagreeing other party has any convincing reasons that might change your mind. Even if you are highly confident about the belief in question, you should still be interested to hear an opposing view, in case you have overlooked some pertinent piece of information. But, if you are a subjectivist, you believe that there is no fact of the matter to be discovered, and that the disagreement is essentially a clash of brute preferences. So you should be less motivated to listen to the arguments of the other party.

As a psychological hypothesis, however, this prediction seems implausible, particularly when examining the open-ended responses that objectivists and subjectivists gave when explaining why disagreement may have arisen (see earlier; objectivist: “the other person has no values”, subjectivist: “values are probably the source of disagreement”). Nevertheless, this way of framing the issue does pose an interesting psychological question: when confronted with ethical disagreement, do ethical objectivists respond in a more “open” or a more “closed” fashion than subjectivists? To examine this question, one of our earlier studies asked participants three different questions to investigate their responses to ethical disagreement. One question asked them to indicate how comfortable they would be to have the person who disagreed with them as a room-mate (a standard social psychological measure of social distance), a second question asked about the personality attributions they would make about the disagreeing other person (personality attributions), and a third question asked whether they thought it would be possible for them to give up their agreement with their original belief (self-reported rigidity of view). Each of these three questions was asked for each of the ethical disagreements that participants were presented with.

Across all three variables, objectivist responders responded in a more “closed” way in the face of ethical disagreement (Goodwin and Darley 2009). They were less comfortable than subjectivist responders with a person who disagreed with them about an ethical issue, even when controlling for how strongly they agreed with the ethical claim in question. They were also more likely to say that the disagreeing other person was “not a moral person”. And they were less likely to indicate that it was possible that they could give up their agreement with the belief in question. A follow-up study replicated these effects using slightly different measures. And in each case, ethical objectivity predicted more closed responses in the face of disagreement, controlling for how strongly participants agreed with the beliefs in question (Goodwin and Darley 2009).

These results indicate that there appear to be important judgmental and possibly behavioral consequences that stem from holding a particular meta-ethical view. Moreover, lay meta-ethical views seem to have predictive effects over and above how strongly people agree with a particular belief. Although strength of agreement and objectivity are correlated—as people agree with an ethical belief more strongly, they tend to be regard it as being more objective—they are not perfectly correlated,

and they have independent predictive effects. Follow up work in this domain should aim to test how meta-ethical views affect real interactions in which an ethical disagreement occurs.

7 Objectivity and Current Theories of Moral Cognition

What is the relationship between studies of moral objectivity and current theorizing in moral cognition? A good deal of current theorizing is concerned with the extent to which emotions relate to moral judgments. One kind of theory concentrates on the emotions that moral violations produce. For instance, Rozin and his collaborators have proposed a theory which argues that characteristic emotions are produced by different sorts of moral violation (Rozin et al. 1999). According to this theory, anger typically arises from violations of personal autonomy (i.e., infractions involving harm, injustice, violations of rights or restrictions in freedom), contempt arises from violations of a community ethic (i.e., violations of social duties and obligations, disrespect for social standing and authority, and disloyalty), and disgust arises from violations of a divinity ethic (i.e., violations of the natural order, of the body, and of purity and sanctity). These authors' experiments demonstrated support for these linkages—when presented with violations of the three different sorts, people tended to select and generate the emotions that 'fit', according to the theory.

Other theories have concentrated on the way in which emotions themselves produce moral judgments. Jon Haidt and his collaborators have performed an intriguing series of studies which demonstrate that incidental emotions can exert large effects on moral judgment. In one study, participants were hypnotized to feel a pang of disgust whenever they read a non-moral word such as "often" (Wheatley and Haidt 2006). They later read a moral vignette which described a moral transgression. All participants were hypnotized, but the crucial variable was whether the targeted word occurred in the story, or not. When it did, participants made harsher moral evaluations of the transgression. The incidental occurrence of moral emotions thus appears to affect moral judgment. Similarly, in more recent studies, Haidt and his collaborators have shown in different ways how the incidental experience of disgust affects moral judgments. For instance, Schnall et al. (2008) have shown how disgust that is elicited by the surrounding environment, or by remembered experiences, produces an increase in the severity of moral judgments about unrelated matters.

Other evidence demonstrates the causal role of emotions in different ways. For instance, Haidt has documented the phenomenon of 'moral dumbfounding', whereby people have strong moral reactions to certain transgressions such as incest, but are not able to explain their reaction in any satisfactory way—'it just feels wrong' being a typical response (e.g., Haidt 2001). The interpretation of this finding was that the moral judgment is primarily caused by an emotional reaction, rather than a reason. Similarly, Haidt has documented that for these kinds of violations, an individuals' emotional reaction to the scenario is a more powerful predictor of their eventual moral judgments than is their assessment of the harm that the action produces (Haidt et al. 1993).

There is also some evidence that self-directed moral emotions are an important cause of moral decisions. Monin and Miller (2001) have documented a phenomenon

which they refer to as ‘moral licensing’. In their studies, when participants were able to affirm their non-prejudiced moral identities in an initial phase of the experiment, they were subsequently more likely to express politically incorrect opinions, and to make hypothetical choices which accorded with those opinions. The explanation for these effects was that a person’s moral behavior is driven in the moment by self-perception of their own moral identity, which arguably has an important emotional component. In a more recent study, Sachdeva et al. (2009) corroborated this finding, and also demonstrated evidence of the inverse effect. Participants who had written about a past instance of their own moral behavior were subsequently less likely to donate money to charity than those had not. But, participants who had written about a past instance of their own immoral behavior were subsequently more likely to donate money to charity. Both effects suggest that self-directed moral emotions can exert a causal influence on moral behavior.

Finally, more recent research has relied on fMRI scanning to document the nature of the emotional responses that may underlie moral judgment (e.g., Greene et al. 2001, 2004; Moll et al. 2003; Sanfey et al. 2003). The upshot of this research has been to show that brain areas that are associated with emotional responding are activated in response to a variety of moral situations. For instance, Greene et al. (2001) showed that moral dilemmas that are thought to activate moral emotions, do in fact activate areas of the brain associated with emotional processing (the medial frontal gyrus, posterior cingulate gyrus, angular gyrus) in contrast with a set of control problems which activated areas associated with working memory (middle frontal gyrus, parietal lobes). Sanfey et al. (2003) showed that the anterior insula region of the brain—an area which is also known to be associated with emotional processing—is more strongly activated in response to unfair offers in the context of the Ultimatum Game,¹⁰ than in response to fair offers. More recent work has illustrated that cognitive processes interact with emotional processes in moral judgment. Greater recruitment of brain regions known to subserve abstract reasoning and cognitive control (including dorsolateral prefrontal cortex and anterior cingulate cortex) predicted longer response times on difficult personal dilemmas, and also more utilitarian responses on such dilemmas (Greene et al. 2004).

There is thus an accumulated body of evidence that illustrates that emotions are strongly associated with moral judgments and actions. But there is disagreement about the nature of this relationship. Prinz (2006) has usefully categorized the possible relationships that could be argued for. The least controversial position, which it is now almost impossible to disagree with, is that emotions co-occur with moral judgments, at least some of the time. Both fMRI and self-report data unequivocally support this contention. It is more difficult to know how often emotions co-occur with moral judgments, and there is little existing data which is relevant to this broader issue.

The most radical position, which Prinz adopts, is that emotions are constitutive of moral judgment—all moral judgments are underpinned by an affective response of some kind, even if this affective response is not activated at the time of judgment.

¹⁰ In this game, participants are first divided into pairs. One participant is given a small sum of money (say \$10), and makes an offer to the other participant which that person has the choice either to accept or reject.

There are some considerable attractions to this position as a philosophical theory, although its empirical support is indirect (see Joyce 2009; Prinz 2007).

A third position is that at least in some cases emotional reactions *cause* moral judgments. Haidt's experimental manipulations of moral emotions and the demonstration of the subsequent effects this has on moral judgment support this position. And indeed, this evidence has made it difficult to argue with the idea that emotions play at least some causal role in moral judgment.

However, there are major debates concerning the extent of emotional involvement in moral judgment, and there is still great scope for disagreement. Some theorists have argued for the dual role of cognitive and emotional processes in moral judgment (e.g., Greene et al. 2001, 2004; Moore et al. 2008; Bucciarelli et al. 2008; see also Cushman et al. 2006; Hauser et al. 2007), although there is no great agreement between these authors about the nature of the cognitive processes involved. Other theorists, however, including Jon Haidt and his collaborators, have maintained that emotions are the critical driver of moral judgments. Although allowing that cognitive override may sometimes affect moral judgments, these authors have maintained a commitment to a social intuitionist model, according to which: 'moral reasoning is an important part of moral life, but for most people, most of the time, most of the action is in the quick, automatic, affective evaluations they make of people and events' (Schnall et al. 2008, p. 1097).

This claim is very strong, and we take our findings concerning the perceived objectivity of ethical beliefs to be relevant to it in the following two ways. First, our data has shown that meta-judgments of moral objectivity seem best predicted by cognitive rather than affective factors. Second, it has also shown that objectivity judgments are divergent from first order moral judgments, i.e., they are divergent from strength of agreement ratings concerning moral beliefs. Both findings suggest that there are important dimensions of moral judgment that are distinct from the initial, affective reactions that play a role in producing first order moral judgments.

Our data also indicates that at least some of the predictors of objectivity judgments—i.e., the extent to which a person grounds their ethical beliefs in God, and the propensity for disjunctive thinking (i.e., the five blocks task)—are not wholly affective. Taken at face value, these measures are predominantly cultural and cognitive, respectively. Is it possible that they have more primitive, affective roots? This seems possible for the religious grounding of ethics, but not for the measure of disjunctive thinking. Perhaps some other, more affective measure, is a better predictor of individual differences in objectivity judgments, and subsumes the two measures we have found to be predictive. This, of course, cannot be ruled out entirely. But it strikes us as an implausible possibility, particularly since Need for Closure, which measures more affective and motivational components than the disjunctive thinking measure, had almost no predictive power in our experiments. Of course, affective reactions may explain part of the variance in objectivity assessments, and as yet, we have not yet measured them directly. Nevertheless, the current evidence suggests there are cognitive factors that underpin assessments of objectivity, and that these do not depend on affective responses.

Our research has also shown people's perceptions concerning the objectivity of an ethical belief diverge from how strongly they agree with that belief. This result is hard to account for on the view that a single affective response underpins all aspects

of moral judgment. Assuming that an underlying affective response causes a person's agreement or disagreement with a particular ethical belief (the social intuitionist claim), then if that affective response is the principal driver of all aspects of moral judgment, it should also cause a convergent judgment about objectivity. We did not observe this sort of convergence in judgment, however. Moreover, although more evidence is needed, judgments of objectivity look like they have important consequences. In our studies, they predicted discomfort with a disagreeing other person over and above, and in fact more strongly than did strength of agreement ratings. They also predicted the tendency to attribute a morally defective personality to a disagreeing other person. And they predicted a self-reported reluctance to give up an ethical belief. All three of these responses are likely to be important when individuals are faced with real ethical disagreements.

Thus, it seems that there is an important property of moral judgment, namely objectivity, which is not predicted by affective reactions, and yet which has important consequences. This claim may not be strictly contrary to Haidt's social intuitionist model, because it could be that affect is the prime cause of first order moral judgments, but not of second order, meta-ethical judgments—a state of affairs that would be surprising for the social intuitionist account, but not inconsistent with it. Nevertheless, we do wish to draw attention the multifaceted nature of moral judgment, and to claim that this ought not to be ignored. Moral judgment depends not just on the interplay between affect and cognition, but also on the interplay between first order and second order moral judgments. Further attention needs to be paid to each of these dimensions in order to better understand moral judgment in all of its complexity.

8 Behavioral Implications

Research on moral objectivity arguably has several implications for public behavior and policy. One implication is that judgments of objectivity may predict how satisfied a person is with the prevailing legal code with regard to particular issues, and a second implication is that judgments of objectivity may affect the extent to which moral values are prioritized in voting behavior. We deal with these two, admittedly speculative, possibilities in turn.

The relationship between law and morality is a notoriously vexed issue of jurisprudence. Our aim is not to comment directly on long-standing debates between legal positivist and natural law interpretations of the law, but rather to position our findings with respect to recent arguments that have highlighted the important relation between moral intuitions and the law. One such argument has claimed that the law influences behavior not just through its threat of sanction, but also through its power to influence and shape social norms (Robinson and Darley 1997). According to this argument, the law is capable of harnessing the powerful forces of social influence and stigmatization in order to produce compliant behavior. This is particularly important in cases where the morality (or immorality) of a particular behavior is not obvious, and is not yet the subject of social norms. However, in order to gain this power to shape norms and behavior, the law must itself possess general moral credibility. It can gain such credibility by accurately reflecting the intuitions of

justice that are held by the majority of the members of a society. If, on the other hand, the law does not have sufficient moral credibility, it will have a reduced power to influence social norms, and in some extreme cases pernicious consequences can result, including vigilantism and flouting of the law.

The Rodney King race riots in Los Angeles were a vivid example of this idea. Recent experimental evidence has also corroborated it. Nadler (2005) showed that when participants were presented with a piece of legislation (Study 1) or a court decision (Study 2) that they found to be unjust (by being too lenient), they reported a greater intention to engage in minor law-breaking activity (Study 1), and to disregard the instructions of a judge when acting as jurors in a mock-trial setting (Study 2). These results are intriguing, and suggest that if the law fails to properly reflect the moral standards of society, negative consequences may result.¹¹ It has yet to be demonstrated that law-breaking *behavior* (as opposed to the intention to commit such behavior) is made more likely as a result of perceiving the criminal justice system to be unjust, which limits the conclusions that can be drawn from this research. However, while this is an important question for future investigations to pursue, other evidence does point to the strong predictive links between intentions and behavior, particularly in cases where the behavior is under volitional control (see e.g., Ajzen and Fishbein 1975, 1977, 1980; Schifter and Ajzen 1985).

Nadler's (2005) flouting effects were demonstrated in cases where participants consider that a particular criminal liability is disproportionate to the transgression that occasioned it. But the effect may occur not only when the proportionate severity of transgression and punishment is in question, but also when the status of the transgression itself—whether it is in fact a transgression—is in question. This is particularly pertinent for a range of activities which are not obviously categorized as wrong in themselves (or 'malum in se' in legal terminology). Such activities include so-called 'victimless crimes' including prostitution, pornography, drug use, incest between consenting adults, and 'public order violations' such as flag burning, or desecrations of symbolic icons.

Whether a particular activity is considered wrong in itself as opposed to wrong instrumentally is a separate conceptual issue from whether it is considered wrong objectively, but there is likely to be considerable overlap in people's judgments. For our purposes, the important question is whether an activity that the law criminalizes and punishes is widely considered to be wrong in an objective sense, and conversely whether an activity that the law fails to criminalize and does not punish, is widely considered to be objectively permissible. Understanding how the public perceives the objective-subjective dimension of particular moral issues is an important part of understanding their moral intuitions about those issues—there is more to a moral intuition than simply whether a person agrees with a judgment of wrongness or permissibility. Our data has shown that perceptions of objectivity predict psychological reactions independently of how strongly a person agrees with a claim about permissibility or wrongness. Thus, we predict that if a person considers say, incest, to be an objective moral wrong, they are likely to be disappointed with a legal

¹¹ Consistent with this interpretation, some theorists have argued that the law has an 'expressive' function, by which it signals the social values of a society, particularly in cases where existing norms about a particular behavior are weak or undecided (e.g., Lessig 1996; Sunstein 1996).

system that does not both criminalize it and impose considerable sanctions upon its occurrence. And they should be more disappointed about this than a person who considers incest to be wrong only in a subjective sense. Conversely, if a person considers say, marijuana smoking to be wrong, but not in any objective sense, they may be disappointed with a legal system that does criminalize and impose harsh sanctions upon it, whereas a person would likely be more comfortable with this result if they consider marijuana smoking to be objectively wrong.

These are speculative claims, because no directly relevant data exist to support them directly. Nevertheless, there is good reason to suspect that researchers' assessments of moral intuitions neglect an important dimension if they do not take into account meta-ethical perceptions of objectivity and subjectivity. Such perceptions may affect people's reactions to a criminal code's decisions about both which acts to criminalize, and its decisions about sanctioning, which may in turn affect their generalized respect for the law, and their intention to follow it.

Perceptions of objectivity may also impact voting behavior. Whether moral beliefs and values are held to be objective facts is likely to influence how those beliefs and values are prioritized in the polling booth. Although there is considerable disagreement about the impact of moral values on voting behavior (see e.g., Fiorina et al. 2004; Burden 2004; Hillygus and Shields 2005; Langer and Cohen 2005), there is good evidence that they play at least some role in how people vote. John Jost (2006) has shown that voting in US presidential elections spanning back to 1972 is strongly predicted by political ideology—a psychological construct which includes moral values to a considerable extent. Somewhat earlier, Sears et al. (1980) demonstrated that long-term political and ideological commitments influence voting behavior more than do immediate considerations of self-interest. Differences in political ideology are obviously influenced by different factual beliefs about, for instance, the most efficient use of government resources, and the likely effects of redistribution of wealth—i.e., by beliefs about the best means to achieve certain ends. But they are also influenced by disagreements over the correct ends to pursue, and the relative importance of moral values such as fairness, equality, self-determination, self-responsibility, and the value of life.

Moral values can affect voting about value of life issues, such as abortion or gay marriage. But their effects extend well beyond these familiar examples. They also affect political opinions about economic, environmental, health-care, and foreign policy issues, to name just a few examples (see e.g., Rorty 1999). Perceptions about the objectivity of these opinions, and the values underlying them, may play a role in voting behavior in the following sense. Some values are considered to be 'sacred' or 'protected' (see e.g., Baron and Spranca 1997; Tetlock et al. 2000; Tetlock 2003), in that they are thought to be so important that they should not be compromised. Tetlock et al. (2000), demonstrated the existence of such 'taboo trade-offs'. They showed that the mere contemplation of a trade-off between a sacred value such as the right to vote, and money (i.e., a scenario in which votes were bought in an election) elicits both moral outrage, and a desire for moral 'cleansing'. Baron and Spranca (1997) showed that values which people claim ought never be compromised, are associated with several other psychology properties, including quantity insensitivity (people are insensitive to the amount of the compromise), agent relativity (the agent's participation in breaching the value is what matters most, not

the consequences of the breach), and moral obligatoriness (that is, they are treated as holding independently of whether people think that is the case).

Voting, of course, inherently involves compromise and trade-off—no one political candidate is likely to endorse the full range of values that you do. But a protected or sacred value is one that is unlikely to be compromised when voting. However, it is not entirely clear why some values are considered sacred or protected, and others are not. One factor that might predict this is objectivity (Baron and Spranca 1997, make a similar suggestion). Hence, we would predict that if a moral value is thought to be objective, it is more likely to be afforded a sacred or protected status—and thus, an individual will be unwilling to compromise on that value when deciding how to vote. Again, no data exist to support this contention, and so it is a task for future research to examine it.

9 Objectivity and Moralism

One final connection with public policy stems from an intriguing recent study which investigated how broadly individuals construe the moral universe. Lovett and Jordan (2005) investigated whether or not people view a range of everyday activities as imbued with a moral quality—a tendency they refer to as ‘moralism’. They demonstrated that people who voted for George Bush in the 2004 US Presidential election were more likely to moralize than were people who voted for John Kerry. That, is they were more likely to perceive activities such as sleeping in late, overeating, and taking on a challenging college course, as involving a moral aspect. One interpretation of this finding is that such voters were more inclined to see personal or private virtue as fundamentally a moral issue.

A tendency to moralize is not the same as a tendency to view moral values as objective. A person might see only a limited range of activities as involving a moral component, but perceive such values to be highly objective. Thus, an interesting psychological question arises as to whether individuals who tend to be objective about moral values are also more likely to see a larger range of activities as involving a moral component—i.e., to ‘moralize’. We suspect that this is the case, although again, this is a question for future research.

10 Conclusion

The psychological investigation of ethical objectivity is in its infancy, but important findings have been made. We have argued that there are systematic sources of variance with respect to ethical objectivity. Some individuals treat their ethical beliefs as more objective than do others, and this appears to be predicted partly by the religious grounding of ethics, and partly by a disposition towards disjunctive thinking. Other factors may yet be discovered which yield further insight into this individual difference. Moreover, some ethical beliefs are treated as more objective than others and this appears to be predicted by the valence of the action under consideration, and may also be predicted by social factors such as perceived consensus. Holding an objective view predicts a more ‘closed’ response to ethical

disagreement, which has implications for how real ethical disputes may be resolved. Many questions still remain to be answered, and we have described several which we think are worth pursuing next.

Acknowledgements We thank Adam Alter, Walter Sinnott-Armstrong, and an anonymous reviewer for helpful comments on an earlier draft.

References

- Ajzen, I., and M. Fishbein. 1975. *Belief, attitude, intention and behavior: An introduction to theory and research*. Reading: Addison-Wesley.
- Ajzen, I., and M. Fishbein. 1977. Attitude-behavior relations: a theoretical analysis and review of empirical research. *Psychological Bulletin* 84: 888–918.
- Ajzen, I., and M. Fishbein. 1980. *Understanding attitudes and predicting social behavior*. Englewood Cliffs: Prentice-Hall.
- Ayer, A.J. 1936. *Language, truth, and logic*. London: Gollancz.
- Baron, J., and M. Spranca. 1997. Protected values. *Organizational Behavior and Human Decision Processes* 70: 1–16.
- Blackburn, S. 1984. *Spreading the word*. Oxford: Oxford University Press.
- Blackburn, S. 1993. *Essays in quasi-realism*. New York: Oxford University Press.
- Blackburn, S. 1998. *Ruling passions*. New York: Oxford University Press.
- Brink, D. 1986. Externalist moral realism. *Southern Journal of Philosophy* 24(Suppl): 23–42.
- Bucciarelli, M., S. Khemlani, and P.N. Johnson-Laird. 2008. The psychology of moral reasoning. *Judgment and Decision Making* 3: 121–139.
- Burden, B.C. 2004. An alternative account of the 2004 Presidential election. *The Forum*, 2. <http://www.bepress.com/forum/>, last accessed on October 10, 2008.
- Cacioppo, J.T., R.E. Petty, and C.F. Kao. 1984. The efficient assessment of need for cognition. *Journal of Personality Assessment* 48: 306–307.
- Cushman, F.A., L. Young, and L.D. Hauser. 2006. The role of reasoning and intuition in moral judgments: testing three principles of harm. *Psychological Science* 17: 1082–1089.
- Darley, J.M., and T.R. Shultz. 1990. Moral rules: their content and acquisition. *Annual Review of Psychology* 41: 525–556.
- Epstein, S. 1991. Cognitive-experiential self-theory: An integrative theory of personality. In *The relational self: Theoretical convergences in psychoanalysis and social psychology*, ed. R.C. Curtis, 111–137. New York: Guilford.
- Epstein, S. 1994. Integration of the cognitive and the psychodynamic unconscious. *American Psychologist* 49: 709–724.
- Fiorina, M., C.M. Abrams, and J.C. Pope. 2004. *Culture war? The myth of a polarized America*. New York: Longman.
- 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., and W. R. Pope. 1984. Ethical ideology and judgments of social psychological research: Multidimensional analysis. *Journal of Personality and Social Psychology* 46: 1365–1375.
- Forsyth, D.R., and R.E. Berger. 1982. The effects of ethical ideology on moral behavior. *Journal of Social Psychology* 117: 53–56.
- Frederick, S. 2005. Cognitive reflection and decision making. *Journal of Economic Perspectives* 19: 24–42.
- 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.
- Gert, B. 2005. *Morality: Its nature and justification* (Rev. ed.). New York: Oxford University Press.
- Gibbard, A. 1990. *Wise choices, apt feelings*. Cambridge: Harvard University Press.
- Gibbard, A. 2003. *Thinking how to live*. Cambridge: Harvard University Press.

- Gill, M.B. 2008. Metaethical variability, incoherence, and error. In *Moral psychology, Vol. 2, The cognitive science of morality: Intuition and diversity*, ed. W. Sinnott-Armstrong, 387–401. Cambridge: MIT.
- Gill, M.B. 2009. Indeterminacy and variability in metaethics. *Philosophical Studies* 145: 215–234.
- Goodwin, G.P. 2009. *Individual differences in judgments of moral objectivity*. Manuscript in preparation.
- Goodwin, G.P., and J.M. Darley. 2008. The psychology of meta-ethics: exploring objectivism. *Cognition* 106: 1339–1366.
- Goodwin, G.P., and J.M. Darley. 2009. *Why are some moral beliefs treated more objectively than others?: Implications for ethical disagreement*. Manuscript in preparation.
- Graham, J., J. Haidt, and B. A. Nosek. 2009. Liberals and conservatives rely on different sets of moral foundations. *Journal of Personality and Social Psychology* 96: 1029–1046.
- Greene, J.D. 2002. *The terrible, horrible, no good, very bad, truth about morality and what to do about it*. Doctoral dissertation, Department of Philosophy, Princeton University.
- Greene, J.D., R.B. Sommerville, L.E. Nystrom, J.M. Darley, and J.D. Cohen. 2001. An fMRI investigation of emotional engagement in moral judgment. *Science* 293: 2105–2108.
- Greene, J.D., L.E. Nystrom, A.E. Engell, J.M. Darley, and J.D. Cohen. 2004. The neural bases of cognitive conflict and control in moral development. *Neuron* 44: 389–400.
- Haidt, J. 2001. The emotional dog and its rational tail: a social intuitionist approach to moral judgment. *Psychological Review* 108: 814–834.
- Haidt, J. 2007. The new synthesis in moral psychology. *Science* 316: 998–1002.
- Haidt, J., and J. Graham. 2007. When morality opposes justice: conservatives have moral intuitions that liberals may not recognize. *Social Justice Research* 20: 98–116.
- Haidt, J., S.H. Koller, and M.G. Dias. 1993. Affect, culture, and morality, or is it wrong to eat your dog? *Journal of Personality and Social Psychology* 65: 613–628.
- Harman, G. 1975. Moral relativism defended. *Philosophical Review* 84: 3–22.
- Hauser, M.D., F.A. Cushman, L. Young, R. Kang-Xing Jin, and J. Mikhail. 2007. A dissociation between moral judgments and justifications. *Mind and Language* 22: 1–21.
- Helwig, C.C., M.S. Tisak, and E. Turiel. 1990. Children's social reasoning in context: reply to Gabennesch. *Child Development* 61: 2068–2078.
- Hillygus, D.S., and T.G. Shields. 2005. Moral issues and voter decision making in the 2004 Presidential election. *Political Science and Politics* 38: 201–210.
- Hunter, J.D. 1991. *Culture wars: The struggle to define America*. New York: Basic Books.
- Jost, J.T. 2006. The end of the end of ideology. *American Psychologist* 61: 651–670.
- Joyce, R. 2009. Review: Jesse J. Prinz: The emotional construction of morals. *Mind* 118: 508–518.
- Kant, I. 1959. *Foundations of the metaphysics of morals* (L.W. Beck Trans.). Indianapolis: Bobbs-Merrill (Original work published 1786).
- Kelly, D., S. Stich, K.J. Haley, S.J. Eng, and D.M.T. Fessler. 2007. Harm, affect, and the moral/conventional distinction. *Mind and Language* 22: 117–131.
- Kohlberg, L. 1963. The development of children's orientation toward a moral order. I. Sequence in the development of moral thought. *Vita Humana* 1: 11–33.
- Kohlberg, L. 1969. Stage and sequence: The cognitive-developmental approach to socialization. In *Handbook of socialization theory and research*, ed. D.A. Goslin. New York: McNally.
- Kohlberg, L. 1976. Moral stages and moralization: The cognitive-developmental approach. *Moral development and behavior: Theory, research, and social issues*, ed. T. Lickona, 31–53. New York: Holt, Rinehart & Winston.
- 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.
- Kruglanski, A.W., D.M. Webster, and A. Klem. 1993. Motivated resistance and openness to persuasion in the presence or absence of prior information. *Journal of Personality and Social Psychology* 65: 861–876.
- Lacey, A.R. 1996. *A dictionary of philosophy*, 3rd ed. London: Routledge.
- Langer, G., and J. Cohen. 2005. Voters and values in the 2004 election. *Public Opinion Quarterly* 69: 744–759.
- Lessig, L. 1996. Social meaning and social norms. *University of Pennsylvania Law Review* 144: 2181–2189.
- Levesque, H.J. 1986. Making believers out of computers. *Artificial Intelligence* 30: 81–108.
- Levesque, H.J. 1989. Logic and the complexity of reasoning. In *Philosophical logic and artificial intelligence*, ed. R.H. Thomason, 73–107. Dordrecht: Kluwer Academic.
- Lovett, B.J., and A.H. Jordan. 2005. Moral values, moralism, and the 2004 Presidential election. *Analyses of Social Issues and Public Policy* 5: 165–175.
- Mackie, J.L. 1977. *Ethics: Inventing right and wrong*. New York: Penguin.

- Moll, J.R., R. de Oliveira-Souza, and P.J. Eslinger. 2003. Morals and the human brain: a working model. *Neuroreport* 14: 299–305.
- Monin, B., and D.T. Miller. 2001. Moral credentials and the expression of prejudice. *Journal of Personality and Social Psychology* 81: 31–43.
- Moore, A.B., B.A. Clark, and M.J. Kane. 2008. Who shalt not kill? Individual differences in working memory capacity, executive control, and moral judgment. *Psychological Science* 19: 549–557.
- Nadler, J. 2005. Flouting the law. *Texas Law Review* 83: 1399–1441.
- Nagel, T. 1970. *The possibility of altruism*. Princeton: Princeton University Press.
- Nagel, T. 1979. *Mortal questions*. Cambridge: Cambridge University Press.
- Nagel, T. 1997. *The last word*. New York: Oxford University Press.
- Nichols, S. 2002. Norms with feeling: toward a psychological account of moral judgment. *Cognition* 84: 223–236.
- Nichols, S. 2004. *Sentimental rules: On the natural foundations of moral judgment*. Oxford: Oxford University Press.
- Nichols, S., and T. Folds-Bennett. 2003. Are children moral objectivists? Children's judgments about moral and response-dependent properties. *Cognition* 90: B23–B32.
- Over, D., J. St. B. T. Evans, and S. Elqayam. in press. Conditionals and non-constructive reasoning. In *Cognition and conditionals: Probability and logic in human thought*, ed. M. Oaksford. Oxford: Oxford University Press.
- Piaget, J. 1965. *The moral judgment of the child*. New York: Free Press.
- Prinz, J.J. 2006. The emotional basis of moral judgments. *Philosophical Explorations* 9: 29–43.
- Prinz, J.J. 2007. *The emotional construction of morals*. Oxford: Oxford University Press.
- Putnam, H. 1987. *The many faces of realism*. LaSalle: Open Court.
- Railton, P. 1986. Moral realism. *Philosophical Review* 95: 163–207.
- Reber, A. 1996. *The penguin dictionary of psychology*, 2nd ed. London: Penguin.
- Robinson, P.H., and J.M. Darley. 1997. The utility of desert. *Northwestern University Law Review* 91: 453–499.
- Rorty, R. 1999. *Achieving our country: Leftist thought in twentieth-century America*. Cambridge: Harvard University Press.
- Rosen, H. 1980. *The development of sociomoral knowledge*. New York: Colombia University Press.
- Ross, L., and A. Ward. 1995. Psychological barriers to dispute resolution. *Advances in Experimental Social Psychology* 27: 255–304.
- Rozyman, E.B., R.F. Leeman, and J. Baron. 2009. *Unsentimental ethics: Towards a content specific account of the moral-conventional distinction*. *Cognition* 112: 159–174.
- Rozin, P., L. Lowery, S. Imada, and J. Haidt. 1999. The CAD triad hypothesis: a mapping between three moral emotions (contempt, anger, disgust) and three moral codes (community, autonomy, divinity). *Journal of Personality & Social Psychology* 76: 574–586.
- Sachdeva, S., R. Iliev, and D.L. Medin. 2009. Sinning saints and saintly sinners: the paradox of moral self-regulation. *Psychological Science* 20: 523–528.
- Sanfey, A.G., J.K. Rilling, J.A. Aronson, L.E. Nystrom, and J.D. Cohen. 2003. The neural basis of economic decision-making in the ultimatum game. *Science* 300: 1755–1758.
- Sayre-McCord, G. 1986. The many moral realisms. *The Southern Journal of Philosophy* 24(Suppl): 1–22.
- Schifter, D.E., and I. Ajzen. 1985. Intention, perceived control, and weight loss: an application of the theory of planned behavior. *Journal of Personality and Social Psychology* 49: 843–851.
- Schnall, S., J. Haidt, G.L. Clore, and A.H. Jordan. 2008. Disgust as embodied moral judgment. *Personality and Social Psychology Bulletin* 34: 1096–1109.
- Sears, D.O., R.L. Lau, T.T. Tyler, and H.M. Allen Jr. 1980. Self-interest vs. symbolic politics in policy attitudes and presidential voting. *The American Political Science Review* 74: 670–684.
- Shantz, C.U. 1982. Children's understanding of social rules and the social context. In *Social cognitive development in context*, ed. F.C. Serifica, 167–198. New York: Guilford.
- Shweder, R.A. 1990. In defense of moral realism: reply to Gabennesch. *Child Development* 61: 2060–2067.
- Shweder, R. A. 2003. *Why do men barbecue?: Recipes for cultural psychology*. Harvard University Press.
- Shweder, R.A., N.C. Much, M. Mahapatra, and L. Park. 1997. The “big three” of morality (autonomy, community, and divinity), and the “big three” explanations of suffering. In *Morality and health*, ed. A. Brandt, and P. Rozin, 119–169. New York: Routledge.
- Sinnott-Armstrong, W. 2006. *Moral skepticisms*. New York: Oxford University Press.
- Sinnott-Armstrong, W. 2009. Mixed-up meta-ethics. *Philosophical Issues* 19: 235–256.
- Smith, M. 1994. *The moral problem*. Oxford: Blackwell.

- Snare, F. 1992. *The nature of moral thinking*. London: Routledge.
- Stevenson, C.L. 1937. The emotive meaning of ethical terms. *Mind* 46: 14–31.
- Sturgeon, N. 1985. Moral explanations. In *Morality, reason and truth*, ed. D. Copp, and D. Zimmerman, 49–78. Totowa: Rowman & Allenheld.
- Sunstein, C. 1996. On the expressive function of law. *The University of Pennsylvania Law Review* 144: 2021–2053.
- Tetlock, P.E. 2003. Thinking the unthinkable: sacred values and taboo cognitions. *Trends in Cognitive Sciences* 7: 320–324.
- Tetlock, P.E., O.V. Kristel, B. Elson, M. Green, and J. Lerner. 2000. The psychology of the unthinkable: taboo trade-offs, forbidden base rates, and heretical counterfactuals. *Journal of Personality and Social Psychology* 78: 853–870.
- Tisak, M.S., and E. Turiel. 1988. Variation in seriousness of transgressions and children's moral and conventional concepts. *Developmental Psychology* 24: 352–357.
- Toplak, M.E., and K.E. Stanovich. 2002. The domain specificity and generality of disjunctive reasoning: Searching for a generalizable critical thinking skill. *Journal of Educational Psychology* 94: 197–209.
- Turiel, E. 1978. Social regulations and domains of social concepts. In *New directions for child development. Vol. 1. Social cognition*, ed. W. Damon, 45–74. New York: Gardner.
- Webster, D.M., and A.W. Kruglanski. 1994. Individual differences in need for cognitive closure. *Journal of Personality and Social Psychology* 67: 1049–1062.
- Wheatley, T., and J. Haidt. 2006. Hypnotic disgust makes moral judgments more severe. *Psychological Science* 16: 780–784.
- 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.