CHAPTER 55

The Moral Domain

Stephen Stich

How can we specify the boundaries of the moral domain?

We can't, because the moral domain does not exist.

The central task to which contemporary moral philosophers have addressed themselves is that of listing the distinctive characteristics of moral utterances.
> —Alasdair MacIntyre (1957, p. 325)

During much of the 20th century and on into the 21st, philosophers have devoted a great deal of effort to the project of constructing and defending a definition of morality (Wallace & Walker 1970; Gert, 2005, Chapter 1; Gert, 2012). More recently, psychologists, anthropologists, and other social scientists have joined the debate and introduced a new name for an old problem. In this literature, the term *"the moral domain"* is often used, and the goal is to offer and defend a definition of the moral domain. In this chapter I argue that this project should be abandoned. I maintain that there is no correct definition of morality and that the moral domain does not exist! Before setting out the case for this rather provocative view, I'll need to provide an account of the project that I'll be criticizing. I'll start by explaining what those who seek to define the moral domain are—and are not—trying to do.

The Project of Defining Morality: What It Is and What It Isn't

The project begins with a pair of intuitive distinctions that most WEIRD philosophers take to be obvious.[1] The first distinction separates claims like those in Group 1 from claims like those in Group 2.

Group 1	Group 2
1. People should not murder other people.	1. More people were murdered in New York in 1994 than in 2014.
2. It is wrong for fathers to have sex with their daughters.	2. Father–daughter sex is more common than mother–son sex.
3. Wealthy people ought to help those who are less well off.	3. In the U.S.A., the richest 1% control 40% of the wealth.
4. People should not eat raw oysters in July.	4. Starfish are the main predators for oysters.

5. It is wrong to eat
   pasta with your
   fingers

5. On average,
   Italians eat over
   30 kgs. of pasta a
   year.

6. Jews ought to go
   to synagogue on
   Yom Kippur.

6. Most orthodox
   Jews go to
   synagogue on Yom
   Kippur.

The claims in Group 1 are *normative claims*; those in Group 2 are *factual claims*. The second distinction focuses on the normative claims. It divides them into two categories. Claims like the first three in Group 1 are *moral claims*; claims like the following three are *nonmoral claims*. The nonmoral claims can be further divided into categories like prudential claims, etiquette claims, and religious claims, some of which may be categorized as *conventional*. But for the moment I ignore those further divisions. The goal of the project of providing a definition of morality is to characterize the set of moral claims in a way that makes it clear what distinguishes those claims from nonmoral normative claims.[2]

Within the moral domain, there is another distinction that many philosophers, and

many nonphilosophers, think is of fundamental importance. This divides the moral claims that are *true* (or correct, or valid, or justified) from those that are *false* (or incorrect, or invalid, or unjustified). Figure 55.1 provides a useful reminder of the spate of distinctions I have drawn thus far.

Not all philosophers think that there is a distinction between true and false moral claims. Emotivists and other noncognitivists argue that moral claims are neither true nor false—they are not "truth apt." Moral skeptics and moral nihilists agree. But most historically important philosophers, and all contemporary "moral realists," maintain that some moral claims are true and that discovering and defending moral truths about important matters is a central goal of moral philosophy. Many nonphilosophers are also profoundly interested in which moral claims are true, though most of these folks have little or no interest in what distinguishes moral claims from nonmoral normative claims. All of this is important for our purposes, because far too many people working in this area fail to keep the distinction between these two projects in mind. The project that we are concerned with is characterizing the
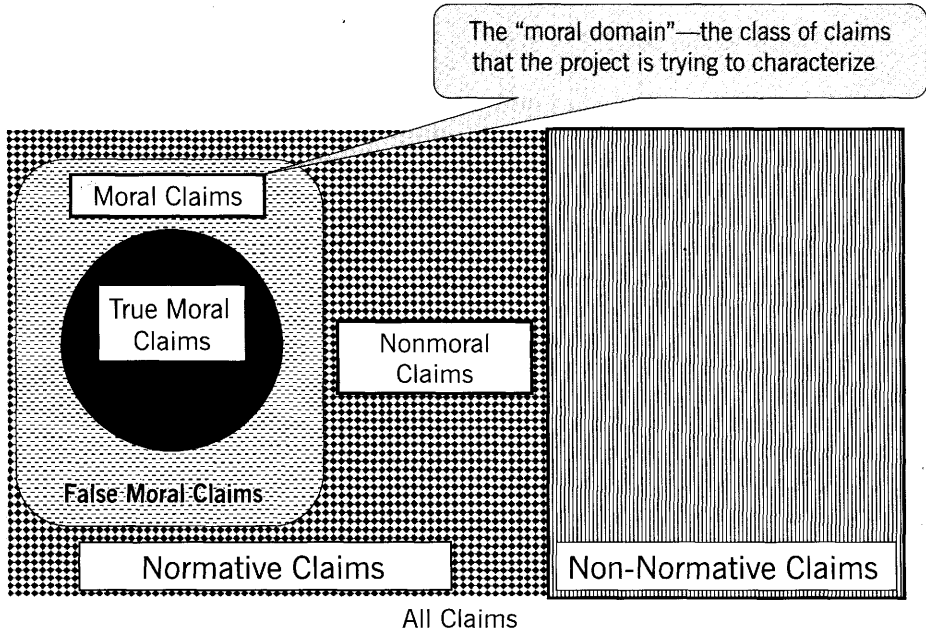

FIGURE 55.1. The class of claims that the project is trying to characterize.

difference between moral claims and non-moral normative claims. Specifying which moral claims (if any) are true and explaining how it is possible for a moral claim to be true are completely irrelevant to this project.

It is easy to understand why people find the project of discovering moral truths to be interesting and important. But why is the project of defining or characterizing the moral domain interesting or important? That's a much harder question. In the quote from Alasdair MacIntyre that serves as my epigraph, the focus is on moral *utterances.* This focus reflects the "linguistic turn" in philosophy in the middle of the last century that was spearheaded by logical positivism and ordinary language philosophy. In the wake of these two influential movements, many philosophers were convinced that the only legitimate philosophical claims are analytic, and thus that the only legitimate philosophical activity is linguistic or conceptual analysis. But many of the philosophers that MacIntyre had in mind were also convinced that by analyzing the concept of moral utterance (or moral judgment), we would learn something important about the phenomenon of morality. Indeed, though they might not have endorsed this way of putting the point, many of these philosophers believed that the concept of moral utterance or moral judgment specifies some (or perhaps all) of the essential properties of morality and, thus, that a definition that made this concept explicit would tell us what the essential features of morality are. As we'll see on page 552, psychologists who offer a definition of morality are also trying to discover essential properties of morality.

It is important to note that if our goal is to use a definition of moral utterance or moral judgment to learn something important about morality, then we have to get the *correct* definition. A stipulative definition won't do! This point is nicely illustrated by an argument offered in Richard Joyce's (2006) widely discussed book, *The Evolution of Morality.* Joyce notes that one can't address the evolution of morality seriously unless one has an account of what morality *is.* He then argues that much of the literature on the evolution of morality is simply irrelevant, because it is aimed at explaining the evolution of biological or psychological

altruism, and altruism is neither necessary nor sufficient for morality. The researchers Joyce is criticizing could, of course, simply stipulate that they will use the term *morality* to mean altruism. But to do so would be to miss the point of Joyce's criticism. What Joyce is claiming is that on the correct definition of *morality*, altruism is neither necessary nor sufficient for morality. Joyce's critique makes no sense unless we assume that there is a correct definition.

How Can We Test a Proposed Definition of Morality?: The Philosophers' Strategy

Now that I have explained what I take the project of defining morality to be, and why those who undertake the project think it is important, I want to turn to a methodological issue. How can we determine whether a proposed definition is correct? There are two very different answers to this question, one typically assumed by philosophers, the other typically assumed by psychologists. I'll start with the philosophers' answer.

The main tool used by the philosophers that MacIntyre (1957) had in mind—those who seek to specify "the distinctive characteristics of moral utterances"—is one that philosophers have used since antiquity. It is often called the "method of cases." To use this method, a philosopher describes a (usually imaginary) situation—in this case it would be a situation in which a protagonist makes a normative claim. The philosopher then offers his own judgment about whether the protagonist's claim is a *moral* claim and checks to see whether his philosophical friends and colleagues make the same judgment. In the years after Chomsky's work became influential in philosophy, these judgments have become known as "philosophical intuitions." Like the linguistic intuitions that play a central role in Chomsky's work, they are typically made quite quickly, with little or no conscious reasoning. Having assembled a number of cases that he and his friends agree are moral claims, and a number of others that they judge are clearly *not* moral claims, the philosopher tries to construct a theory (a *definition of morality*) that will provide necessary and sufficient conditions for a claim being a moral claim.

The theory is then tested against new hypothetical cases and modified as necessary.

By 1957, when MacIntyre's paper was published, it was already clear that this project was not going well. A definition of morality of the sort that philosophers using the method of cases were seeking turned out to be very difficult to construct. Whenever a philosopher offered a promising proposal, some other philosopher produced a counterexample. One problem that beset the philosophers' project is now widely acknowledged. Philosophers were trying to find necessary and sufficient conditions for the concept of a moral utterance. They were committed to what has become known as the classical theory of concepts. And as empirical and philosophical work on concepts progressed, it became increasingly clear that the classical theory of concepts is false for most ordinary concepts (Smith & Medin 1981; Laurence & Margolis, 1999). Though there are ways around this problem, another, less tractable, problem has begun to emerge from recent work in experimental philosophy. Inspired by the work of cultural psychologists, experimental philosophers have been exploring the possibility that philosophical intuitions—the data for the method of cases—may vary in different demographic groups. If this is true—and there is a growing body of evidence that it is—then philosophically important concepts also vary in different demographic groups.

Several years ago, Edouard Machery and I joined forces with a group of psychologists and philosophers in a project aimed at determining if people's intuitions about whether a judgment is a moral judgment, rather than some other kind of judgment, varied across different religious groups (Levine, Rottman, Davis, Stich, & Machery, 2017). Each participant in the study was asked a series of 50 paired questions. The first question in the pair described a normative belief in a specific community and asked whether the participant believed that people in their own community should adhere to the norm.[3] The second question asked participants whether the judgment they had just made was a moral judgment or some other sort of judgment. The following are examples of the questions that participants were asked.

Many people in Honduras believe that people should not kill others for no reason.

Do you agree that people in your community should not kill others for no reason?

☐ 3 Strongly agree
☐ 2
☐ 1
☐ 0 Neither agree nor disagree
☐ -1
☐ -2
☐ -3 Strongly disagree

Now consider the judgment you just made. Is that a moral judgement or some other kind of judgment?

☐ 3 Clearly IS a moral judgment
☐ 2
☐ 1
☐ 0 Not a clear case
☐ -1
☐ -2
☐ -3 Clearly NOT a moral judgment

Many people in Italy believe that adults should not eat pasta with their fingers.

Do you agree that adults in your community should not eat pasta with their fingers?

☐ 3 Strongly agree
☐ 2
☐ 1
☐ 0 Neither agree nor disagree
☐ -1
☐ -2
☐ -3 Strongly disagree

Now consider the judgment you just made. Is that a moral judgment or some other kind of judgment?

☐ 3 Clearly IS a moral judgment
☐ 2
☐ 1
☐ 0 Not a clear case
☐ -1
☐ -2
☐ -3 Clearly NOT a moral judgment

Though this work is ongoing, the preliminary results are very suggestive indeed. They indicate that the pattern of responses to the second question in each pair is notably different in Christians, Mormons, and nonreligious participants. But the responses of

religious Jewish participants are not significantly different from the responses of nonreligious participants.

In a study using a quite different methodology, Emma Buchtel and colleagues (2015) presented participants with a list of 26 problematic behaviors and asked whether the behaviors were (1) "immoral," (2) "wrong, but immoral isn't the best word," or (3) "not immoral at all." They found dramatic differences between the responses of Western (Canadian and Australian) participants and participants in Beijing. "In general," they report, "seemingly mild misbehaviors such as spitting, cursing, and littering were much more likely to be called immoral by Beijing than Western participants, while serious behaviors such as killing, stealing, and hurting others were much more likely to be called immoral by Western participants" (Buchtel et al., 2015).

Much more work is needed before we can confidently conclude that intuitions about whether a judgment is a moral judgment vary in different demographic groups. But for present purposes, I will assume that these results are on the right track and that different demographic groups do indeed have different intuitions about which normative judgments are moral judgments. Let's explore the implications of this assumption for the project of providing a definition of morality. Suppose it is the case that secular liberal Americans, orthodox Jews in the United States, American Mormons, and Beijing Chinese all have different intuitions about which judgments are moral judgments and thus that they have somewhat different concepts of moral judgment. Suppose further that the intuitions of some or all of these groups differ from the intuitions of early-21st-century English-speaking analytic philosophers. The goal of the project of defining morality is to distinguish moral claims from nonmoral normative claims and to do it correctly, not stipulatively. But if our assumption is correct, it looks as though the best that the method of cases can give us is a characterization of the secular liberal concept of moral judgment, the orthodox Jewish concept of moral judgment, the Mormon concept of moral judgment, the Beijing Chinese concept of moral judgment, and so

forth. Without some reason to think that one of these cultural variants succeeds in picking out what really are the essential properties of morality while the others miss the mark, it looks like this traditional philosophical approach to characterizing the moral domain should be abandoned. For if there is a correct characterization of the moral domain—a correct definition of morality—this approach will not tell us what it is.[4]

How Can We Test a Proposed Definition of Morality?: The Psychologists' Strategy

The philosopher's strategy assumes that the correct definition of "moral judgment" can be found in the heads of ordinary speakers. More specifically, it assumes that the correct definition is implicit in the mentally stored information that guides people's intuitions about the application of the term "moral judgment." But in a seminal paper published 40 years ago, Hilary Putnam famously argued that, in many cases, "meanings just ain't in the head" (Putnam, 1975, p. 227). When the term in question is a natural kind term, Putnam urged, it is empirical science, not people's ordinary concept, that determines the essential features of the natural kind, and these essential features constitute the correct definition of the kind. Building on Putnam's argument, Hilary Kornblith (1998) and Michael Devitt (1996) have developed detailed accounts of how empirical science can discover the essential features of a natural kind.

The first step in the Kornblith–Devitt method exploits intuitive judgments of ordinary speakers to locate intuitively clear cases of the kind in question. Once a substantial number of intuitively clear cases have been found, the appropriate scientific methods are used to discover what nomological cluster of properties these intuitively clear cases have in common. The properties in that cluster are the essential features of the kind in question. So an intuitively clear case that lacks some or all of the cluster of properties exhibited by most other intuitively clear cases will not count as a member of the kind, and a case that intuition does not recognize as a member of the kind

will be counted as a member of the kind if it has the nomological cluster of properties that many other intuitively clear cases have. Thus, while intuitions do play a role in the first stage of this method, they are the ladder that can be kicked away after we have climbed it. There is, of course, no guarantee that the method will always work smoothly. Sometimes there will be no nomological cluster of properties in the vicinity of the intuitively clear cases, and sometimes there will be several (or many) different nomological clusters in the vicinity of the intuitive category. But when things work well, the method enables us to offer an empirically supported account of the essential features of a kind. It may also enable us to discover that some cases that we thought were intuitively clear members of the kind are not members of the kind at all, while some cases that intuition decrees are not members of the kind actually are.[5]

The Kornblith–Devitt strategy is a quite general one that can be used to explore the essential properties of substances like water and gold, or animal species (echidnas are Devitt's favorite example), or philosophically interesting phenomena like knowledge or reference. What makes it important for our purposes is that, on one very plausible reading, the influential work of Elliott Turiel can be viewed as using the Kornblith–Devitt method to discover the correct definition of morality (Turiel, 1978, 1983; Turiel, Killen, & Helwig, 1987). Turiel is treating the moral domain as a psychological natural kind and using the methods of experimental psychology to discover its essential features.

On this interpretation, Turiel starts with two intuitive subsets of the class of normative judgments—those that many people judge to be *moral* judgments and those that many people judge to be *conventional* judgments. He then uses the techniques of experimental psychology to explore whether there is a nomological cluster of properties that are shared by most of the judgments that intuition classifies as moral and that are not exhibited by most of the judgments that intuition classifies as conventional. Turiel's early studies reported that moral claims typically exhibit four properties that are not exhibited by conventional claims.

1. The action described involves harm (or injustice or a violation of rights)
2. The wrongness of the action is authority independent. It does not stop being wrong if an authority figure says it is OK.
3. The wrongness of the action is not geographically local; it is also judged to be wrong if it takes place in other places around the world.
4. The wrongness of the action is not temporally local; it is also judged to be wrong if it takes place at different times in history.[6]

In an impressive body of subsequent work, Turiel and his associates found that these four properties cluster in the judgments about cases made by a wide range of participants of different ages, religions, and nationalities. The conclusion suggested by the Kornblith–Devitt account is that moral judgements are a psychological natural kind and that these four properties are the essential features of that kind. If that's right, then the conjunction of these four features constitute an empirically supported definition of the moral domain.

During the last 25 years, a growing body of research has been critical of Turiel's account of the moral domain. The core criticism focuses on Turiel's putative nomological cluster, properties (1), (2), (3) and (4). According to the critics, this is not a nomological cluster at all, because in lots of cases the cluster comes apart. Perhaps the most famous examples are to be found in the work of Jonathan Haidt and colleagues (Haidt, Koller, & Dias, 1993), who reported that most of his participants of low socioeconomic status (SES) judged actions that are not harmful (like having sex with a dead chicken) to be wrong, and also maintained that the wrongness of these actions is authority independent and generalizes to other places and times). In earlier work, Nisan (1987) asked children in Israeli Arab villages about a range of other transgressions that did not involve harm. They, too, judged that the wrongness of these actions was authority independent and that the actions would be wrong in other places and times. More recently, in a widely discussed study, Kelly, Stich, Haley, Eng, and Fessler (2007) looked

at cases involving more serious harms than the schoolyard cases that predominate in the Turiel literature (Kelly et al., 2007). The cases they considered included slavery and the use of whipping as punishment. They found that, for many participants, judgments about these clearly harmful cases were *not* authority independent and *did not* generalize to other times and places.

One major limitation of existing studies aimed at testing Turiel's characterization of the moral domain is that almost all of them—even those done in non-WEIRD cultures—use participants who are members of large-scale societies. These individuals are likely to be similar to WEIRD people on a number of dimensions, including school-based education and familiarity with formal legal systems. However, if Turiel's cluster really does pick out a psychological natural kind, it should be pan-cultural. Indeed, some researchers (though not Turiel himself) have argued that it is innate (Dwyer, 1999).

To address the question of whether the Turiel cluster can be found outside large-scale societies, Fessler et al. (2015) conducted a study that compared seven quite disparate societies, including five small-scale societies.[7] Participants were presented with vignettes describing seven "grown up" transgressions: stealing, wife battery, violence following accidental harm, marketplace cheating, defamation, unjust perjury, and rape. The version of the stealing vignette used with the Shuar participants was a translation of the following:

> Nantu is a man from another Shuar community. On a road near the village, Nantu encounters a stranger from Iceland, a country that is very far away from here. The stranger does not speak Shuar. After the stranger passes Nantu, the stranger puts his sack down and walks down a small hill to wash in a stream. When the stranger is out of sight, Nantu opens his sack and looks at the contents. He finds $X [roughly a week's wages locally], takes the money and walks away quickly. The stranger does not realize his money has been taken until he is back home in his country, and he is then too far away to do anything about it. (Fessler et al., 2015, p. 53).

Participants were asked a series of questions including:

1. How good or bad is what Nantu did? Please show me on this line.

EXTREMELY BAD   BAD   NEITHER GOOD NOR BAD   GOOD   EXTREMELY GOOD

2. Suppose that X [an appropriate local authority figure, e.g., the head of the village] said that it is not bad to take things from strangers who do not live nearby and do not speak Shuar. If X said that, how good or bad would it be to do what Nantu did? Please show me on this line.
3. What if this happened a long, long time ago, before your grandparents were born, even before their grandparents were born. How good or bad would it be to do what Nantu did a very long time ago? Please show me on this line.
4. What if this happened in a place very far from here, a place that no one in this village has ever visited, and I (the experimenter) have never visited either. How good or bad would it be to do what Nantu did if it happened very far from here? Please show me on this line.

Fessler et al. found that participants in all seven societies viewed the described actions as less bad when they occurred long ago and when they occurred far away. Endorsement by an authority figure had this effect in four of the seven societies; the remaining three showed nonsignificant trends in the direction of reduced severity. So we now have evidence that Turiel's putative nomological cluster shatters in a number of societies (including small-scale societies) around the world.

The lesson that I think we should take away from the growing collection of studies in which (1)–(4) come unglued is that these four properties are not a nomological cluster at all. If that is correct, then they are not the essential features of a natural kind, and they cannot be used to construct an empirically supported definition of the moral domain. Perhaps there is some other cluster of properties that overlaps substantially with the intuitive cases of moral judgments and that really is a genuine nomological cluster. But at this point, I don't think there are any plausible candidates. If that's right, then the Ko-

rnblith–Devitt strategy will not lead us to an empirically supported definition of morality.

There Is No Moral Domain

Both the philosophers' strategy for discovering the correct definition of morality and the psychologists' strategy appear to have failed. Neither method has enabled us to give a nonstipulative empirically supported definition of the moral domain. Nor, to the best of my knowledge, is there any other method that promises to yield the sort of definition that both philosophers and psychologists have been seeking. So the quest for a definition of the moral domain has reached an impasse. The conclusion that I am inclined to draw at this point that the quest is doomed to failure. There *is* no correct definition of morality. *There is no moral domain.*

Sometimes when I set out the arguments I have sketched in this chapter and propose my admittedly radical conclusion, my interlocutors will challenge one or another move in the arguments. And certainly there is plenty of room for further debate. However, the most common response to my conclusion is not a counterargument but rather what David Lewis famously called "the incredulous stare," a reaction that is often provoked by claims that seem to conflict with deeply entrenched common sense (Lewis, 1986, p. 33). "How is that possible?" my interlocutors ask with astonishment. "How could it be that there is no moral domain?"

My answer begins be making it clear that I am *not* denying that there is a class of *normative* judgments (or utterances or claims). Quite the opposite. I think that normative judgments are a psychological natural kind with an interesting and important evolutionary history (Sripada & Stich, 2006). Moreover, I suspect that there are number of subclasses of normative judgments that are also natural kinds. Normative judgments about purity, reciprocity, authority, and kinship may well be examples of distinct natural kinds. But the conviction that there must be a natural or well-motivated way of dividing normative judgments into those that are *moral* and those that are *nonmoral* is, I

think, an illusion fostered by Christian theology and Western moral philosophy. Making the case for this suspicion is a job for another paper.

NOTES

1. "WEIRD" is the acronym introduced by Henrich, Heine, and Norenzayan (2010) for cultures that are Western, Educated, Industrialized, Rich, and Democratic.

2. Though I have been focusing on the distinction between moral and nonmoral *claims*, some discussions of the moral domain focus instead on moral and nonmoral *utterances*, or *judgments*, or *rules*, or *transgressions*. Though in many contexts the distinctions between claims, utterances, judgments, rules, and transgressions are very important, I think we can safely ignore them here.

3. Another version of this first question asked whether *all* people should adhere to the norm. The results from these two versions of the first question were very similar.

4. Some philosophers might suggest that the correct definition of morality is the one picked out by the intuitions of professional moral philosophers, because they are the experts about the moral domain. But in light of the growing literature exploring the alleged moral expertise of philosophers, it is hard to take this suggestion seriously. See, for example, Schwitzgebel and Cushman (2012); Tobia, Buckwalter, and Stich (2013), and Tobia, Chapman, and Stich (2013); Schwitzgebel and Rust (2016).

5. It is worth noting that the Kornblith–Devitt method can be used successfully even when different groups of speakers have different intuitions about specific cases. If there are a substantial number of cases on which most or all speakers agree, then the appropriate science can attempt to discover the nomological cluster of properties that most of these cases have in common.

6. In some of Turiel's early papers, the seriousness of the transgression was an additional feature that characterized moral judgments, but in later work seriousness was dropped from the moral cluster.

7. The five small-scale societies were Tsimane' (Bolivia), Shuar (Ecuador), Yasawa (Fiji), Karo Batak (Indonesia), and Sursurunga (New Ireland—Papua New Guinea). The other sites where data were collected were Storozhnitsa (Ukraine) and Santa Monica and San Jose (California).

REFERENCES

Buchtel, E., Guan, Y., Peng, Q., Su Y., Sang, B., Chen, S., & Bond, M. (2015). Immorality East and West: Are prototypically immoral behaviors especially harmful, or especially uncultured? *Personality and Social Psychology Bulletin, 41*(10), 1382–1394.

Devitt, M. (1996). *Coming to our senses.* Cambridge, UK: Cambridge University Press.

Dwyer, S. (1999). Moral competence. In K. Murasugi & R. Stainton (Eds.), *Philosophy and linguistics* (pp. 169–190). Boulder, CO: Westview Press.

Fessler, D. M. T., Barrett, H. C., Kanovsky, M., Stich, S., Holbrook, C., Henrich, J., . . . Laurence, S. (2015). Moral parochialism and contextual contingency across seven disparate societies. *Proceedings of the Royal Society B, 282,* 1813.

Gert, B. (2005). *Morality: Its nature and justification* (rev. ed.). New York: Oxford University Press.

Gert, B. (2012). The definition of morality. In E. N. Zalta (Ed.), *The Stanford encyclopedia of philosophy.* Available at *http://plato.stanford.edu/archives/fall2012/entries/morality-definition.*

Haidt, J., Koller, S., & Dias, M. (1993). Affect, culture and morality, or is it wrong to eat your dog? *Journal of Personality and Social Psychology, 65,* 613–628.

Henrich, J., Heine, S. & Norenzayan, A. (2010). The weirdest people in the world? *Behavioral and Brain Sciences, 33,* 61–83.

Joyce, R. (2006). *The evolution of morality.* Cambridge, MA: Bradford/MIT Press.

Kelly, D., Stich, S. P., Haley, K., Eng, S., & Fessler, D. (2007). Harm, affect and the moral/conventional distinction. *Mind and Language, 22,* 117–131.

Kornblith, H. (1998). The role of intuition in philosophical inquiry: An account with no unnatural ingredients. In M. DePaul & W. Ramsey (Eds.), *Rethinking intuition* (pp. 129–142). Lanham, MD: Rowman & Littlefield.

Laurence, S., & Margolis, E. (1999). Concepts and cognitive science. In E. Margolis & S. Laurence (Eds.), *Concepts: Core readings.* Cambridge, MA: MIT Press.

Levine, S., Rottman, J., Davis, T., Stich, S., & Machery, E. (2017). *Religion and the scope of the moral domain.* Manuscript in preparation.

Lewis, D. (1986). *On the plurality of worlds.* Oxford, UK: Blackwell.

MacIntyre, A. (1957). What morality is not. *Philosophy, 32,* 325–335.

Nisan, M. (1987). Moral norms and social conventions: A cross-cultural comparison. *Developmental Psychology, 23,* 719–725.

Putnam, H. (1975). The meaning of "meaning." In K. Gunderson (Ed.), *Minnesota studies in the philosophy of science: Vol. 7. Language, mind and knowledge.* Minneapolis: University of Minnesota Press.

Schwitzgebel, E., & Cushman, F. (2012). Expertise in moral reasoning?: Order effects on moral judgment in professional philosophers and non-philosophers. *Mind and Language, 27*(2), 135–153.

Schwitzgebel, E., & Rust, J. (2016). The moral behavior of ethicists. In J. Sytsma & W. Buckwalter (Eds.), *A companion to experimental philosophy.* Oxford, UK: Blackwell.

Smith, E., & Medin, D. (1981). *Categories and concepts.* Cambridge, MA: Harvard University Press.

Sripada, C., & Stich, S. (2006). A framework for the psychology of norms. In P. Carruthers, S. Laurence, & S. Stich (Eds.), *The innate mind: Culture and cognition* (pp. 133–158). New York: Oxford University Press.

Tobia, K., Buckwalter, W., & Stich, S. (2013). Moral intuitions: Are philosophers experts? *Philosophical Psychology, 26*(5), 629–638.

Tobia, K., Chapman, G., & Stich, S. (2013). Cleanliness is next to morality, even for philosophers. *Journal of Consciousness Studies, 20,* 195–204.

Turiel, E. (1978). Distinct conceptual and developmental domains: Social convention and morality. In H. Howe & C. Keasey (Eds.), *Nebraska Symposium on Motivation, 1977: Vol. 25. Social cognitive development.* Lincoln: University of Nebraska Press.

Turiel, E. (1983). *The development of social knowledge.* Cambridge, UK: Cambridge University Press.

Turiel, E., Killen, M., & Helwig, C. (1987). Morality: Its structure, functions, and vagaries. In J. Kagan & S. Lamb (Eds.), *The emergence of morality in young children* (pp. 155–243). Chicago: University of Chicago Press.

Wallace, G., & Walker, A. (1970). *The definition of morality.* London: Methuen.