

## SECTION I

# Science

“Hume’s Law” urges us to distinguish statements of how things are and attempts to explain why they are that way from evaluations of the world so described and any policy proposals, recommendations, or decisions we might premise in our evaluations (Hume, T 3.1.1). So understood, Hume’s Law poses challenges for a putative science of moral epistemology, for it is extraordinarily difficult to advance claims about morality without therein taking a stance on what is and what is not immoral, where statements of immorality are commonly treated as evaluative if not prescriptive in character. Nevertheless, despite the historical challenges to the enterprise, academic scientists have returned to the study of morality en masse. Is this because we have found a way to respect Hume’s Law? Or have we simply grown comfortable flouting its demands? Is a science of morality really possible?

The scientist authors of this section’s chapters have values; and, like the rest of us, they have views about what we ought to be doing as a community. Surely these values have affected their choice of what to describe, how to describe it, and which explanations of the putative data to report. Moreover, though these authors have attempted to describe and explain things as they have found them, they have had to assess whether one putative explanation of what they’ve described is better than another. To be fair, the judgment that one explanation of the data is better than another is supposed to emerge from an epistemic assessment of the theories in play, not a moral evaluation of the goodness or badness of the reality these theories are meant to explain. But the theories in question have “our” moral views and practices as their object, and the authors in question are members of “us.” So it is entirely appropriate to wonder whether they’ve succeeded in isolating their epistemic evaluations from the moral values or principles that find expression in their nonacademic lives, as when they reprimand others or defend their actions from judgment, or endorse certain political candidates and criticize others, or advocate for changes in our laws or public policies.

To their credit, academics have developed methods for achieving some level of “objectivity” in their assessment of morality. First, philosophers have attempted to give accounts of what knowledge of right and wrong would have to be were we to have such knowledge without taking a stand on the reality of what they’ve described. These projects are usually

framed as analyses of our concepts of moral knowledge or accounts of the social practices in which these concepts are applied. Following J. L. Mackie, many label these projects “second-order” or “metaethical” theories. It should be noted that some of those who advance metaethical theories do not think of themselves as scientists, perhaps because they don’t feel the need to conduct or analyze experiments to lend credence to their claims. But many philosophers do consider metaethics a science. And even those uncomfortable with this label tend to advance their accounts as true or accurate representations of our moral thinking. Authors in both these camps must either reject Hume’s Law or show that their metaethical hypotheses have no implication for the “first-order” morality of reader or author. Since philosophers agree on very little, it is not surprising that they continue to debate whether an author’s metaethics can be isolated from her ethics in the manner proposed.

But a second approach to the scientific study of morality is suggested by the traditional analysis of knowledge itself. Though E. Gettier demonstrated to the satisfaction of most contemporary epistemologists that knowledge does not reduce to justified, true belief, most of us continue to posit belief as the core psychological component of knowledge. (It seems reasonable to suppose that you must be convinced of something to know it.) So a scientist of morality might eschew talk of “moral knowledge” in favor of “moral belief” or “moral judgment.” Or, in an attempt to denote her target subject in full generality, the theorist might write of the genesis, development, and operations of “moral cognition.” On this understanding of the terrain, psychology, ethology, anthropology, and sociology are the scientific components of moral epistemology.

Stephen Stich uses Chapter 1 of this volume to recount the recent history of this “Philosopher’s Project”: the attempt to distinguish our moral psychologies from other components of our minds without making substantive assumptions about what is right and wrong. Many philosophers analyzed paradigmatic moral cognitions as *representations of rules*, but to distinguish *moral* rules from rules of etiquette and the like, R. M. Hare, inspired by Kant, argued that you don’t think of a norm as moral in character unless you treat it as “universalizable” and “prescriptive.” Other theorists turned their attention to isolating distinctively moral modes of thinking. For instance, to distinguish genuinely moral reasoning from prudential calculation, W. Frankena proposed that some *consideration of others* and their interests is essential. N. Cooper and P. Taylor joined Frankena in hypothesizing that a person’s moral code consists of those prescriptions she treats as “overriding or supremely important.” And A. Gewirth added Kant’s idea of *categoricity*: to think of a rule as moral you must think that you are bound to follow it even when obedience would thwart your ends or frustrate your desires. Several theorists added some *susceptibility to guilt or ostracism* in the wake of a norm’s violation, and additional conditions were proposed.

Predictably, the philosophers failed to achieve consensus. But when psychologists, led by E. Turiel, eventually extracted a working definition of “moral cognition” from the philosophical literature, they abandoned the attempted neutrality of most analyses by requiring some relation to harm, welfare, justice, or rights. Though many psychologists found evidence that people conceptualize distinctively moral rules in the way Turiel supposed, and many theorists still distinguish moral rules from “mere” conventions in this way, Stich articulates the seeds of the project’s ruin. First, the formal analyses or definitions of “moral cognition” have fallen into question. For example, Stich reports evidence that some people

think of rules against corporal punishment as authority-dependent. Are prohibitions on the practice not then moral rules? Stich rejects this conclusion and proposes that “moral cognition” is not associated with a unitary concept. Instead, this phrase and others like it correspond to different concepts in the minds of different people, though our paradigms of moral rules, moral reasoning, or moral motivation may be similar. Second, Stich describes how J. Haidt and others demonstrated that morality is not limited to harm, welfare, rights, and justice in the minds of illiberal people. For instance, disgusting acts are often “moralized,” even when they neither cause harm nor constitute injustice.

Elizabeth O’Neill and Edouard Machery, the authors of Chapter 2, agree with Stich’s critique of the moral/conventional distinction. They join Haidt in broadening the category of “moral cognition” beyond rules related to harm, care, fairness, and reciprocity to include also “groupish” norms of patriotism, ideals of loyalty, authority, and respect, and rules preserving purity or sanctity. They go on to add to Haidt’s list our concern for privacy and honesty.<sup>1</sup> There is even evidence that some non-Western people “lump” all of their rules together; that there is no difference in their minds between norms associated with these “foundations” and norms of other kinds. They conclude that some people fail to draw a distinction of any kind between moral norms and nonmoral conventions.

Instead of trying to define distinctively moral cognition, O’Neill and Machery try to isolate a broader phenomenon: *normative* cognition, which they define as the capacity to learn social rules, the disposition to follow them, the tendency to punish rule-breakers, and some susceptibility to a range of characteristic emotions, including admiration, disgust, guilt, and shame. When it is defined in this way, normative cognition is indeed a human universal, and O’Neill and Machery report evidence that all people “externalize” at least some of their norms to some degree. They also describe a significant overlap in the contents of norms embraced by diverse cultures; recount how almost every culture judges an agent’s intentions relevant to the propriety of punishing her for a given violation; and report a study suggesting that no community treats ignorance of the norms in play as an excuse for violating its rules. But O’Neill and Machery find a great deal of variation in the content of norms beyond these areas of overlap, along with significant differences with regard to the importance of an agent’s intentions for judgments of her blameworthiness. They also report substantive variation with regard to how much of a community’s life is governed by norms of any kind. Some societies are more rule-heavy than others.

Though the universality of normative cognition among humans does not imply the innateness of a shared “normative sense,” it lends credence to a continuity hypothesis of some sort. Mightn’t *homo sapiens* have inherited our proclivity toward rule governance from the human-like apes from whom we evolved? In Chapter 3, Sarah Vincent, Rebecca Ring, and Kristin Andrews shed light on this question by describing the “ought thoughts” of other animals. They find these normative cognitions implicated in various “normative practices,” which are defined as “patterns of behavior shared by members of a community that demonstrate they value certain ways of doing things as opposed to others.”

According to Vincent, Ring, and Andrews, when the leader of a wolf pack prevents a female member from breeding with a strange male, she is enforcing a *norm of obedience*. When a dog tucks her tail and hides her face because she anticipates a scolding for stealing cake, her guilt evidences her susceptibility to these same norms. Indeed, when an older ape

critiques a youngster's initially unsuccessful attempts at termite fishing, norms of obedience are expressed and taken to heart in the process. In partial contrast, when capuchin monkeys protest getting a cucumber after observing a fellow receive a more highly valued grape for the same task, the capuchin has expressed her aversion to injustice, implicating a *norm of reciprocity*. This second group of norms is supposed to guide monkeys' exchanges of food and grooming services and their expressed dissatisfaction with unfair deals. Animal acts of self-sacrifice and consolation are instead guided by *norms of altruism* or caring, as when orcas attack ships to save their pod-mates, humpback whales save a seal from the pursuit of these same orcas, or a polar bear mourns the death of a mate. *Norms of social responsibility* are manifested in various distributions of goods and divisions of labor. One example might be the sentinels among a scurry of Belding field squirrels, who draw danger upon their own heads by whistling warning of a hawk's approach. In this vein, Vincent et al. report a chimp in the Kansas City Zoo who propped a log against the enclosure's wall to serve as a ladder and then "beckoned to another six chimps to join him" in his escape. Finally, the authors posit *norms of solidarity* that reinforce the common identity of the communities in which they live. These norms are invoked to help explain why a group of cetaceans might develop an in-group language of whistles and clicks or beach themselves collectively. The authors go on to examine the lives of chimpanzees and cetaceans in detail, reporting an array of normative practices, which are in turn supposed to provide evidence of a similarly complex manifold of "ought thoughts" in the minds of those animals enacting them. They conclude by rebutting various deflationary explanations of the practices they report and then arguing, against C. Korsgaard (but in keeping with M. Rowlands), that an animal can be guided by norms even if she doesn't have the ability to introspect, interrogate, and modify her initial reactions to the behavior of a conspecific so as to comply with rules she accepts "from" her reflective endorsement of them.

So as not to beg questions against the parties to this dispute, let us use "explicit norm guidance" to refer to the reflective capacity we have just described. If one chimp wants to mate with another and refrains from doing so *because* she represents this as something forbidden—or something she ought not do—we will say that she is explicitly guided by the norm in question. But what is it to think of some proposed action as wrong or forbidden? If we analyze this thought in the way proposed by O'Neill and Machery in Chapter 2, we must look for evidence that the chimp in question enforces the mating hierarchy in cases in which she is not personally implicated or that she now complies with it because of the guilt and remorse she experienced after prior indiscretions. Vincent, Ring, and Andrews do not argue that explicit rule guidance (so understood) is manifest among populations of nonhuman animals, but they also don't rule it out, and there is some intriguing evidence in favor of the hypothesis. De Waal, for example, argues that other animals exhibit "willpower," as when they forgo a present reward to secure more remote advantages (2016, 221–229), and willpower would seem to implicate explicit norm guidance of a sort. To be fair, when an animal suppresses an experienced appetite for one grape in the hopes of therein securing ten, she is explicitly guided by prudential rather than moral norms. But when these norms implicate others, suppression of appetites and aversions in their service might be thought to constitute genuinely "moral" norm guidance. (This is a possibility to which Vincent et al. remain open.) In the end, we might join Darwin in awarding full marks to dogs who overcome fear to save their owners.

The conception of normative cognition that emerges from an evolutionary perspective is more detailed, less idealized, and more realistic than those assumed by traditional moral epistemologists. But these theories remain relatively abstract and conjectural. Since the posited psychological processes of norm guidance are often supposed to be introspectively inaccessible to us, only neuroscience can tell us whether aspects of our moral lives that might have evolved via natural mechanisms of selection operating upon populations of our hunter-gatherer ancestors really did so evolve and persist to this day.

Unfortunately, to describe the characteristic inputs, functions, and outputs of a neurological process in psychological or computational terms, we need to utilize concepts drawn from outside neuroscience. Joanna Demaree-Cotton and Guy Kahane explain the relevance of this realization to normative moral epistemology in Chapter 4 when they argue that findings in neuroscience cannot be used to evaluate the reliability or adaptivity of a set of moral intuitions unless we can infer which psychological process a neural network is implementing. The authors go on to assess the role that neuroscience has played within the cognitive science of morality (CSM) more generally and the prospects that the CSM will have “normative significance” by affecting the first-order moralities of those of us who have been exposed to it.

Demaree-Cotton and Kahane describe how difficult it is to “map” neurological processes onto “higher-level” psychological processes to confirm or infirm hypotheses about the proximate causes of our relatively automatic normative intuitions and judgments. First, a psychological process may be differently realized in the nervous systems of different people or groups of people. “For example, emotional processing that is normally supported by paralimbic brain areas in nonclinical populations might be supported by the lateral frontal cortex in psychopaths.” Second, a discrete brain area or neural network can support many different processes, which both complicates the attempt to assign that area a unified computational function and undermines efforts to assess whether the cognition it enables is reliable or adaptive. It may turn out that a single network participates in both reliable and unreliable (or adaptive and maladaptive) cognitive processes. Finally, we may have good reason to recognize cognitive processes that are not neurologically discrete.

As Demaree-Cotton and Kahane see it, CSM is organized around three main paradigms. Two of these are the “two-systems” or dual-process models of Haidt and Greene discussed in detail in many of this volume’s chapters. The third is J. Mikhail’s “universal moral grammar” approach—critiqued by Stich in Chapter 1—which posits an innate faculty for representing intent, harm, and the intuitive wrongness of intentionally inflicting harm. The authors are relatively dubious of the use to which neuroscience has been put in defense of these models, arguing that a range of different neural networks are implicated in moral judgment. According to their review of the relevant literature, we currently lack neurological evidence of those specialized or “dedicated” mechanisms of normative judgment posited by evolutionary psychologists. Instead, the evidence suggests that moral rules are learned and imbued with emotion in the way Nichols describes in Chapter 6.

Demaree-Cotton and Kahane also use neurological evidence to further undercut the reason-emotion dichotomy placed into question by May and Kumar in Chapter 7. The neurological evidence suggests that emotions are key components of reflection, decision, and choice. And this undermines Greene’s attempt to identify utilitarian calculation with the neural correlates of “reason” in order to dismiss the processes responsible for deontic

intuition as unreliable “passions” or emotions. The neurology responsible for aversion to killing one to save five isn’t correlated with knee-jerk reactions of disgust or fear. Instead, these “emotional” processes, mediated by areas such as the right temporoparietal junction (rTPJ), amygdala, and ventromedial prefrontal cortex (vmPFC), allow us to assign intention, distinguish relevant from irrelevant information, and weigh different pieces of relevant information against others to arrive at “all things considered” judgments. What emerges is not a “preponent” emotional response, as Greene has maintained, but a deployment of Aristotelian phronesis or practical wisdom. Evidence for this reconceptualization is provided by observation of clinical populations. Damage to the aforementioned brain regions is indeed correlated with higher rates of utilitarian judgment, but it is also correlated with psychopathology, blindness to the import of intention, diminished empathy, and an increased tendency to punish perceived slights. “In nonclinical populations, so-called ‘utilitarian’ judgments that Greene associated with the dorsolateral prefrontal cortex (dlPFC) are not associated with genuinely utilitarian, impartial concern for others but rather with rational egoism, endorsement of clear ethical transgressions, and lower levels of altruism and identification with humanity.” For this reason, “current evidence suggests that both emotions and reasoning contribute to moral judgment and that moral judgment may operate at its best when reasoning and emotion interact.”

To adopt an evolutionary perspective on morality we must focus on the accumulation of changes in populations of animals over “deep” time. By studying the norms enacted by our ancestor species, trying to correlate these behaviors with similar activities in humans, and looking for the shared neurological structures that enable these interactions to proceed as they do, scientists are trying to uncover the origins and underlying reality of the human moralities enacted across the globe today. But to fully understand our moralities we must incorporate a more proximate developmental perspective. Each person develops a set of moral views and dispositions over her lifetime. Developmental psychologists examine these processes. What have they discovered? And how do these discoveries mesh with the other sciences of morality that we have examined?

In Chapter 5, Julia W. Van de Vondervoort and J. Kiley Hamlin address these questions by recounting relatively recent paradigms in developmental moral psychology. Their history overlaps with Stich’s to some extent as we read of how psychologists from Piaget to Kohlberg imputed the history of political thought from Hobbes to Kant into the activities of the children they studied. Moral reasoning was supposed to emerge from self-interested calculation when the needs and interests of other children were made salient during playground disputes. The “highest” form of moral development was to be found in Rawlsian calls to limit rules to those no reasonable person would reject. Though Van de Vondervoort and Hamlin are less critical than our other authors of the Kohlbergian tradition in developmental moral psychology, they report a number of the criticisms that have been brought against it. Children do not always focus on outcomes to the exclusion of intentions in their evaluations of behavior, and when a child’s parents distinguish moral rules from nonmoral conventions by treating violations of the former as more serious than violations of the latter, their children cognize this distinction at a much earlier age than Kohlberg allowed.

Van de Vondervoort and Hamlin also question Kohlberg’s assumption that Kantian reasoning marks the pinnacle of moral development by reviewing a large body of evidence

linking emotion to moral cognition. They begin with the analyses of Hume and Adam Smith, who both argued that emotional reactions of approval and disapproval acquire moral content when we judge that we would continue to experience them were we to adopt a “general view” or imagine ourselves impartial spectators to the events to which we are reacting. Hypothesizing that judgment-sculpted condemnation and approval of these kinds evolved to facilitate cooperation, Van de Vondervoort and Hamlin report evidence that our attraction to helpers and aversion to hinderers originates in preverbal infancy. Two-year-olds distinguish intentional harm from accidental damage, and they protest intentional acts of harm and injustice no matter who commits them, but they limit their condemnation of unconventional acts to those “in-group” members who are party to the convention. Infants as young as 3 months old seem to track the social valence of an act, preferring helpful graphics and puppets to characters shown hindering the pursuits of others. And infants as young as 18 months manifest intuitions of fairness, exhibiting a preference for equal distributions of goods except when the labor or success of some would merit their receiving a larger share, (a finding which resonates with the evolutionary hypotheses defended by Cosmides et al. in Chapter 9). Chapter 5 concludes with an assessment of the evidence that infants distinguish morality from prudence. In one study, children in their first year rejected more snacks from a hinderer, choosing instead fewer snacks offered by a helper. Do these infants prefer helpers over hinderers because they hope to benefit from future interactions with the partners they’ve chosen? Are infants capable of the more disinterested evaluations that Hume and Smith equated with distinctively “moral” judgment? According to our authors, the jury is still out on these matters.

Van de Vondervoort and Hamlin conclude that both emotions and reasoning are implicated in moral development. But exactly how do these elements interact in the development of the reader’s more or less “mature” morality? In Chapter 6, Shaun Nichols tries to put emotions in their proper place by embracing a relatively cognitivist account of the genesis of our intuitive moral judgments. On J. Greene’s dual-processing account, our resistance to killing one to save five is primarily constituted by a “preponent” aversion to the act. But why, asks Nichols, do we feel this way toward killing as a means? Drawing on work in machine learning, Nichols distinguishes the kind of “model-free” learning that inculcates habits and instinctive responses from the “model-based” learning that provides animals with the kind of information they need to navigate their environments in more flexible ways. Habits tend to persist for a time even when they don’t serve our ends, but rational animals overcome their habits when the benefits of doing so outweigh the costs. When an animal makes this calculation but nevertheless “defers” to the habit blocking her ends, Nichols classifies her behavior as weakness of will. He cites, as an example, someone who desperately wants to scuba dive but gives up trying because she finds it difficult to surmount her instinctive aversion to breathing underwater.

Do our deontic intuitions originate in the kind of model-free learning responsible for our instinctive aversions? Nichols reports work by F. Cushman that is supposed to support a positive answer. For instance, subjects are reluctant to smash an obviously fake hand when they have no qualms about striking a nut. Of course, subjects are even more averse to striking real hands, but the point remains: people are averse to an action that tends to be harmful (e.g., hitting what looks to be a hand) even in cases in which it is manifestly not harmful.

But Nichols provides reasons for doubting consequentialist attempts to debunk deontic judgment as an “overlearned” response of this same kind. Though instinctive or habituated aversion is a regular component of deontic intuition, we are often averse to an action we do not judge wrong. (Think of someone spitting into a cup and then drinking it. This is disgusting but not obviously immoral.) Nichols’s positive proposal is that we don’t judge an act wrong unless we represent it as the intentional performance of an action known to be prohibited by rule, where model-free learning is insufficiently robust to account for our knowledge of social rules and agents’ intentions.

To argue that an understanding of social rules is necessary for moral judgment, Nichols critiques P. Railton’s cognitive account of the “broad affective system” implicated in those of our moral judgments that are not grounded in conscious reasoning. Railton insists that emotions are not “dumb,” as they are often attuned to risks and rewards, obstacles and affordances. For example, mightn’t our relatively automatic aversion to incest attune us to the dangers of this practice? According to Railton, this assessment is not undermined by the observations of Haidt and colleagues that subjects remain averse to a described act of incest even when researchers stipulate that it hasn’t caused harm. After all, we are similarly averse to games of “Russian roulette” that do not end in suicide. But Nichols insists that affect is “less flexible and sensitive to evidence” than “general cognition” and he casts doubt on Railton’s diagnosis. Risky behavior isn’t typically conceptualized as immoral, and few of us have the kind of experience with incest that would engender an emotional memory of its deleterious effects, so it’s unlikely that our belief in incest’s immorality can be chalked up to “affective attunement.” Instead, variation in incest norms across cultures suggests that most of us have been taught the particular incest prohibitions operative in our communities.

Nichols concludes his chapter with a description of the kind of statistical learning implicated in a child’s mastery of the social rules operative in her milieu. Though we are prone to various inductive and probabilistic fallacies, recent studies show that children utilize valid heuristics like a “size principle” to shape their expectations. Nichols suggests, on this basis, that youngsters can extract subtle principles (such as the greater immorality of harming in comparison with allowing harm to occur) from their exposure to a range of moral judgments without explicit instruction in these principles. Perhaps parental and religious prohibitions on incest are facilitated by this kind of implicit pattern recognition, which functions alongside a more or less innate aversion to sexual encounters with siblings to yield the kind of aversion Haidt has recorded. In either event, a typical person’s belief in the immorality of incest is not a wholly “system 1” product.

In Chapter 7, Joshua May and Victor Kumar look at how reasoning and emotions interact in mature moral agents who have learned a morality in the ways Nichols describes. They begin by endorsing a “two-systems” model of cognition in general and moral cognition in particular. According to May and Kumar, when Humean philosophers and psychologists deemphasize the role of reasoning in the genesis, modification, and entrenchment of our moral judgments, they are neglecting unconscious or “system 1” reasoning, which is supposed to be quicker, more automatic, and less flexible than the “system 2” reasoning of which we are introspectively aware. When you think like a utilitarian, calculating the likely impact of a proposed course of action on those you know will be affected, you are utilizing your slow, effortful system 2. In contrast, system 1 processes account for the automatic aversion you experience to the prospect of killing one person to save five others. While



they admit that your reaction to a “Trolley case” of the relevant sort will have an emotional component, May and Kumar suggest this feeling might be an effect of your belief that the ends don’t justify the means rather than its cause. As an example of this phenomenon, they discuss ideological vegetarians who only come to experience disgust at the sight of meat after accepting arguments against killing animals for food. The authors also consider the possibility that affective processing runs “in parallel” with the unconscious reasoning responsible for our deontic intuitions. In offering their analysis, May and Kumar take aim at other “two-systems” theorists, like Haidt, who identify system 1 with “emotion” rather than “reason.” But the authors also assign a substantive role to conscious inference, rejecting Haidt’s suggestion that system 2 reasoning is limited to the lawyerly defense of a moral judgment that has been challenged. System 2 reasoning is supposed to allow us to achieve greater levels of consistency between our moral intuitions, attain reflective equilibrium between our intuitive judgments of particular actions and the moral principles we embrace, and help us suppress or even eliminate automatic reactions to one another we reject as racist, sexist, or unduly prejudicial.

What then of emotion? Psychopaths have a diminished capacity for sympathy and guilt. Doesn’t this distort their moral thinking? May and Kumar are not convinced, hypothesizing that deficits in empathic concern and allied emotions may adversely affect the moral development of psychopathic children, even if these emotions do not play a significant, proximate role in our adult capacity for moral thought. As evidence for this, they focus on those patients analyzed by Demaree–Cotton and Kahane in Chapter 4, who are thought to retain their capacity for moral judgment despite sustaining damage to the ventromedial cortex in adulthood, which deadens their emotional sensitivities. The authors also cast doubt on the significance and replicability of studies that are supposed to show that disgust and anger magnify moral condemnation. But May and Kumar allow that emotions influence reasoning in cases of wishful thinking, self-deception, and confirmation bias and that damage to the brain areas most directly implicated in the experience of emotion are correlated with impairments in deliberation and decision. They therefore conclude “that the way to attain and maintain moral knowledge will require improving both reasoning and emotion.”

AuQ2

To account for our intuitive judgments, May and Kumar posit automatic, effortless bouts of reasoning that we cannot introspectively identify or describe. In Chapter 8, Piotr Patrzyk explores morally relevant heuristics in greater depth, recounting the pioneering work of G. Gigerenzer and colleagues. He begins with a critique of the kind of excessive idealization in moral psychology that results from modeling people who have not been exposed to Kant and Mill as “tacit” deontologists or consequentialists. Kohlberg, in particular, is criticized for not distinguishing the reasoning utilized to defend or justify a judgment from the cognitive processes implicated in its genesis. To get at the real causes of our moral verdicts, Patrzyk insists that we begin with an assessment of the “computational feasibility” of a proposed mechanism or decision rule. For example, we rarely have the information we would need to calculate expected utilities, so it is reasonable to suppose that we rarely do so. Instead, limited bodies of information trigger “domain-specific” mechanisms that in turn account for the different sorts of normative judgment we render, where we can assume, in advance, that the calculations or inferences instantiated by these mechanisms are tractable, robust, frugal, and quick.

Patrzyk extends his critique to prominent advocates of the “two systems” approach to moral judgment. Haidt is criticized for saying nothing about how system 1 takes us from the description of a scenario to an intuition of the rightness or wrongness of the actions portrayed, an allegation lent force by the cognitivist accounts of system 1 advanced by May, Kumar, and Nichols in previous chapters. Theorists have also tended to assume that system 2 processing “corrects” system 1 intuitions, rendering reflective moral judgment more reliable, but studies show that time pressure sometimes augments virtuous choice, as in a public goods game. The problem, Patrzyk claims, is that “system 1” and “system 2” are vague, overly idealized labels. But instead of replacing these terms with more descriptively adequate theories that might account for the variable effects of time pressure, some two-systems theorists implausibly claim that as a general matter our intuitive responses are self-interested, then quickly change to incorporate the interests of others, and then revert to amorality when more time is devoted to choice. According to Patrzyk, this is “data fitting” at its worst, as experimental results are shoe-horned into the two-systems framework they in fact undermine. The dominant paradigm in behavioral economics is even more idealized than the two-systems view, and Patrzyk goes on to offer devastating criticisms of the economist’s penchant for claiming that we make our decisions “as if” we are trying to maximize expected utility. “Such research does little to answer questions about how humans perceive dilemmas, what information they look for and in what order, and how they combine information to make decisions.”

How then should we model the mechanisms or processes implicated in the genesis and revision of our moral intuitions? To answer this question, Patrzyk cites work by Delton, Krasnow, Cosmides, and Tooby on why people cooperate with strangers. Delton et al. “contextualize” the decision problem by assuming that the mechanisms responsible for a decision to cooperate initially evolved under conditions of selection. A disposition to cooperate is an adaptation, as hunter-gatherers augmented their reproductive fitness by cooperating with fellow tribe members for mutual benefit. But when we utilize these strategies in our present context, we cooperate in ways that often fail to advance individual fitness. Those who ignore human history when crafting their models of moral judgment and choice entirely overlook this possibility.

Patrzyk concludes by urging researchers to use what is known about the evolution of humanity to frame more realistic accounts of judgment and choice. We should assume, in particular, that processes of judgment and choice are domain specific, that the mechanisms executing these processes evolved under conditions of selection because they solved the problems humans faced in those conditions, and that utilizing these mechanisms was “rational” in these environments insofar as it secured solutions that were better or more adaptive in comparison with their tractable alternatives. But there is no way to determine whether a mechanism or process of judgment or choice is both tractable and plausibly realized in a human’s mind or brain without describing that mechanism in algorithmic detail. Those advancing serious hypotheses need to describe the “search rules” that guide acquisition of the inputs to decision, the “stopping rules” that determine when the search for information gives way to decision, and the “decision rules” that take a mental mechanism from premises to conclusion. As an example of best practices, Patrzyk describes a (2017) study conducted by Tan, Luan, and Katsikopoulos on the conditions under which we will forgive someone for a perceived indiscretion.

In Chapter 9, Leda Cosmides and John Tooby, leading researchers in this field of study, join Ricardo Guzmán to provide a masterful overview of what evolutionary moral psychology has achieved to date. They begin by arguing for the domain specificity of various mechanisms of normative judgment on evolutionary grounds. “It is hard to see how natural selection would favor a single, unitary system for generating and regulating our choices—moral or otherwise—when programs tailored for ‘tracking’ and promoting fitness in one domain (e.g., cooperative hunting, followed by sharing) require features that are not required to track and promote fitness in other domains (e.g., courtship, with competition for exclusive access to mates) . . . it is reasonable to predict as many domain-specific cognitive adaptations as there are domains in which the definitions of (evolutionarily) ‘successful’ behavioral outcomes are incommensurate.”

After describing the foundations of the evolutionary approach, Cosmides, Guzmán, and Tooby describe its application to norms of incest and familial obligation. Inbreeding diminishes the reproductive fitness of families over time. Because of this, “natural selection will favor mutations that introduce motivational design features that cost-effectively reduce the probability of incest.” Some primates have solved this problem by mixing populations, as animals of one sex (typically males) leave the troop to breed. “But for species like ours, in which close genetic relatives who are reproductively mature are commonly exposed to each other, an effective way of reducing incest is to make cues of genetic relatedness reduce sexual attraction.” And this dynamic is not limited to sexual intercourse. Because our foraging ancestors typically lived with close kin throughout their lives, opportunities abounded for helping and hurting those related to them, where the fitness benefits of aiding or hindering kin often coincided with their degree of genetic relatedness. Cosmides et al. sketch the kind of “kin selection” operative in these contexts in detail and argue that foragers needed some means for discerning the genetic relatedness of individuals in their tribes to settle on adaptive policies for social interaction. A “kin detection mechanism” evolved, consisting of some “monitoring circuitry” designed to register evidence of genetic relatedness and a “kinships estimator” that transformed these cues into an index of family ties. On the basis of studies conducted by Lieberman and others, the authors suggest that maternal perinatal association (MPA) was used to discern genetic relatedness of mother and child but that younger children used the cumulative duration of coresidence to detect the relatedness of their older siblings.

As we saw in Nichols’ Chapter 6, several authors have used cultural variability in incest norms as evidence that they’re learned or enculturated. Cosmides et al. try to deepen this analysis by providing a principled basis for predicting the variation in question. Lieberman et al. report that the degree of disgustingness and moral wrongness that older siblings assign to incest with younger opposite-sex siblings is directly proportional to the “MPA cue” provided by maternal care for the sibling in question but that the judgments of disgustingness and wrongness that younger siblings assign to incest with older siblings is better correlated with duration of coresidence. They also report that MPA and coresidence do not predict judgments of disgustingness or wrongness of same-sex incest. Moreover, “The same cues that regulate moral intuitions about incest—MPA and coresidence duration—regulate how often people sacrifice to help their siblings (as measured by favors done in the last month) and their willingness to incur large costs (such as donating a kidney), whether they are true

biological siblings, stepsiblings, or unrelated children being raised together on a kibbutz.” The concept of kin selection is then used to make general predictions about the way in which judgments of familial obligation will vary in accordance with genetic relatedness and other factors.

Chapter 9 concludes with an extensive “cook’s tour” of the psychological mechanisms that have been proposed by evolutionary psychologists to account for the various ways in which humans cooperate with one another. When distantly related people cooperate, one party often enough augments the reproductive fitness of another at some cost to herself. Because the parties are not closely related, the psychology responsible for this behavior will not have its origins in kin selection. Famously, Darwin invoked non-kin group selection to account for patriotic sacrifice and the triumph of “civilized” moralities over those operative among “savages.” Indeed, the power of group selection led Darwin to an optimistic assessment of the prospects for Christian ideals of universal brotherhood and their secular counterpart: the generalized benevolence preached by Utilitarians. In contrast, our authors argue that we do not need to appeal to selection among extrafamilial groups to explain why some people sacrifice for the benefit of relative strangers. Cooperation evolved between unrelated individuals when we developed a way of detecting cheaters and a way of distinguishing dedicated cheaters from those who have accidentally failed to keep their ends of mutually beneficial bargains. “Generosity in one-shot interactions evolves easily when natural selection shapes decision systems for regulating two-person reciprocity (exchange) under conditions of uncertainty.” If the costs of an initial act of altruism were relatively low, and the costs of missing out on reciprocation were fairly high, and the opportunities for reciprocation in hunter-gatherer societies were sufficiently great, cooperating by default with those who did not strike us as cheaters would have been adaptive and so favored by selection amongst individuals within a unitary population.

Our authors are similarly dismissive of the use of group selection to explain why people engage in costly acts of rule enforcement. Studies suggest that people mainly blame and punish those with whom they plan to interact in the future, and this suggests to Cosmides et al. that our investment of resources into punishment are similarly conditional in origin. The authors use “partner choice” of this kind to account for a wide variety of moral intuitions, including judgments of virtue and the respective roles we assign to effort and luck when evaluating the fairness of a distribution of goods or the character of a compatriot. In contrast, they argue that “partner control” is necessary to mitigate free riding in large societies. They conclude that each form of social interaction they have addressed is enabled by domain-specific modes of social conceptualization and dedicated modes of inference. Again, there are exactly as many “modules” of moral judgment as there were problems of cooperation to be surmounted by our foraging ancestors.

These conclusions help Cosmides et al. explain why our moral intuitions are better preserved by particularist and pluralistic normative theories than those frameworks derived from a “first principle” of morality. Of course, as Hume would insist, the evolutionary psychologists’ account of why we have the moral intuitions we do is neither a vindication of those intuitions nor a reason to abandon them. To address this frankly philosophical matter we must turn to Section II of this volume, where questions about the epistemological significance of what is now known about our moral psychologies are addressed in some detail.

**Note**

1. Note that in the third chapter of this volume, Vincent et al. cite Iyer et al. (2012), who add concern with liberty and oppression as a distinct category of normative cognition.

**Sources Cited**

- de Waal, F. B. M. (2016). *Are We Smart Enough to Know How Smart Animals Are?* New York: W. W. Norton.
- Iyer, R., S. Koleva, J. Graham, P. Ditto and J. Haidt, "Understanding libertarian morality: the psychological dispositions of self-identified libertarians," *PloS ONE* 7/8, e42366/.

Taylor & Francis  
Not for distribution



## 1

# THE QUEST FOR THE BOUNDARIES OF MORALITY\*

*Stephen Stich*

Alasdair MacIntyre begins his paper, “What Morality Is Not,” with a claim that may strike many philosophers as very surprising indeed.

The central task to which contemporary moral philosophers have addressed themselves is that of listing the distinctive characteristics of moral utterances.

(1957, 26)

MacIntyre is indulging in a bit of literary license here. The philosophers he has in mind were not just concerned with moral *utterances*, they were also concerned to give accounts of moral judgments, moral principles, moral norms and moral issues, and of what is required for a set of “action guiding”<sup>1</sup> principles to be a moral code—a morality. With this caveat noted, MacIntyre was surely right. The philosophical literature in the late 1950s was chock-a-block with discussion of what is required for an utterance (or judgment, or principle, etc.) to count as *moral*. Much of this literature was inspired by R. M. Hare’s enormously influential book, *The Language of Morals* (1952), and his article, “Universalizability” (1954–1955). Moreover, the outpouring of philosophical literature in this area continued long after MacIntyre’s essay, with important articles appearing throughout the 1960s, ’70s and ’80s.

The existence of this bountiful literature—which has largely disappeared from the philosophical curriculum over the last quarter century—raises a number of questions, including:

1. What were these philosophers trying to do?
2. Why did they want to do it? Why was it thought to be important?
3. How did they propose to discover the distinctive characteristics of moral judgments, principles and the rest? What sorts of evidence or argument did they rely on?
4. What characteristics were proposed; which were agreed on?
5. Why did contributions to this literature gradually diminish?
6. How is more recent work relevant to the project that these philosophers were pursuing?



I'll try to answer the first five of these questions in part 1, and the sixth in part 2. In part 3, I'll explain how this philosophical literature—a bit of it—was woven into the foundation of a psychological project that also sought to characterize the “distinctive characteristics” of moral judgments, rules and transgressions and that has had an important influence on contemporary empirical moral psychology. In part 4, a preliminary conclusion, I'll review what we've done. Subsequent parts address more contemporary theories as I ask what lessons can be learned from the six decades of philosophical and psychological research we'll be reviewing.

### 1. The Philosophers' Project (≈1952–≈1990): What Were These Philosophers Trying to Do?

To understand what these philosophers were trying to do, we must begin with a crucial distinction. Often, when we ask whether a person's judgment is moral, what we want to know is whether her moral judgment is *true*—or something in that vicinity: correct, or valid, or justified, or wise. What we are asking, to use Frankena's (1967) useful terminology, is whether the judgment is moral as opposed to *immoral*. It is hardly surprising that philosophers often want to know whether a judgment or a principle is moral (as opposed to immoral). Limning the contours of the moral (in this sense), has been a goal of philosophy since antiquity.<sup>2</sup> But it is very important to keep in mind that this was *not* the goal of the writers engaged in what I'm calling “The Philosophers' Project.” Rather, borrowing again from Frankena, what they were trying to do was to distinguish moral judgments, principles, etc. from *nonmoral* judgments or principles. So, for example, they wanted to know how to determine whether an action guiding rule that is widely accepted in a given culture is a moral rule or some other sort of rule—a religious rule, for example, or an aesthetic rule, or a prudential rule. Whether the rule is true, or valid, or justified, etc., was simply not their concern. Similarly, confronted with the unfamiliar, largely egoistic action guiding rules described in John Ladd's (1957) detailed study of the Navajo, they wanted to know whether this system of rules was a morality. If it was not, then, arguably, the Navajo did not have a moral code at all, and thus having a moral code is not a human universal. Closer to home, these philosophers wanted to specify how to distinguish a moral rule from a rule of etiquette. Are the tacit rules specifying appropriate behavior for people waiting on line to board a bus or to buy a coffee at Starbucks moral rules or just rules of etiquette?<sup>3</sup> How about rules specifying appropriate clothing to wear at important events, like funerals? They also wanted some principled way of determining which legal rules are also moral rules.

### 2. Why Did They Want to Do It? Why Was It Thought to Be Important?

The philosophers we are concerned with wanted to give an account of the conditions required for a judgment or a rule to be moral as opposed to nonmoral. Why? One reason, on which there was wide agreement, was that the account would enable us to give principled answers to the sorts of questions mentioned in the previous paragraph. It would, for example, tell us whether the Navajo, as described by Ladd, had a moral code.<sup>4</sup> It would also tell us whether rules about how to behave while waiting on line are moral rules, whether a specified legal rule is also a moral rule, etc. Another, more controversial reason was that the account would be a specification





of the *essence of morality*. While a number of authors endorsed this view,<sup>5</sup> others adamantly rejected it. According to Paul Taylor, “The importance of classifying moral principles . . . does not lie in the discovery of the essence of morality. (There is no such essence)” (1978, 52).

With the explosion of research in empirical moral psychology over the last two decades and philosophers’ growing interest in the area, many new questions have been raised that seem to require the sort of account that philosophers engaged in the Philosophers’ Project were seeking. One clear example can be found in Richard Joyce’s influential book, *The Evolution of Morality* (2006). Joyce wants to provide an account of the evolution of the “moral sense,” which he characterizes as “a faculty for making moral judgments” (44). But we can’t undertake an inquiry into the evolution of the moral sense, Joyce maintains, without an account of what moral judgments are.

Any attempt to understand how our ability to make moral judgments evolved will not get far if we lack a secure understanding of what a moral judgment is. (To neglect this would be like writing a book called *The Origin of Virtue* without any substantial discussion of what virtue is).

(44)

He goes on to offer his own chapter-length account of “the nature of morality,” which includes a detailed attempt to answer the question, “What is a moral judgment?”<sup>6</sup>

Another example that has garnered a great deal of attention grows out of some provocative and problematic claims by Jonathan Haidt. About a decade ago, Haidt, who has been one of the most influential moral psychologists in recent years, accused his fellow moral psychologists of politically motivated bias. Here is a quote that nicely summarizes Haidt’s critique.

[S]tudents of morality are often biased by their own moral commitments. . . . One problem is that the psychological study of morality, like psychology itself, has been dominated by politically liberal researchers (which includes us). The lack of moral and political diversity among researchers has led to an inappropriate narrowing of the moral domain to issues of harm/care and fairness/reciprocity/ justice. . . . Morality in most cultures (and for social conservatives in Western cultures), is in fact much broader, including issues of in-group/loyalty, authority/respect, and purity/sanctity. . . .

This article is about how morality might be partially innate. . . . We begin by arguing for a broader conception of morality and suggesting that most of the discussion of innateness to date has not been about morality per se; it has been about whether the psychology of *harm* and *fairness* is innate.

(Haidt & Joseph, 2007, 367)

To make their case for a broader conception of morality, Haidt and Joseph offer a brief overview of norms that prevail in other cultures. These norms include “rules about clothing, gender roles, food, and forms of address” and a host of other matters as well (371). They emphasize that people in these cultures care deeply about whether or not others follow these rules. But this is a puzzling way to defend their accusation. For surely Haidt



and Joseph don't think that the "politically liberal researchers" responsible for the "inappropriate narrowing" of the moral domain are *unaware* that rules governing these matters are widespread in other cultures. They don't think that these liberal researchers don't read the newspaper or that they are anthropological ignoramuses. The issue in dispute is not whether rules like these exist or whether people care deeply about them. What is in dispute is whether these rules are *moral* rules. To resolve that dispute, we need an account of what it is for a rule to be a moral rule.

In recent years, the philosophical literature has been awash in claims about the semantics of moral judgments (Boyd, 1988; Horgan & Timmons, 1992; Schroeder, 2008), the function of moral judgments (Roskies, 2003; Prinz, 2015), the evolutionary history of moral judgments (Joyce, 2006; Kitcher, 2011) and the psychological mechanisms underlying moral judgments (Nichols, 2004a; Prinz, 2007). In order to evaluate these claims, we need to know which normative judgments they apply to—which ones are *moral* judgments. And that is exactly what the Philosophers' Project is trying to provide.

### 3. How Did They Propose to Discover the Distinctive Characteristics of Moral Judgments?

Most of the philosophers who participated in the debate over the definition of morality during the last half of the twentieth century agreed that an analysis of ordinary linguistic usage had an important role to play in discovering and defending an appropriate definition. If a proposed definition classified as *moral* a judgment that we would not ordinarily describe as a moral judgment—or if it classified as *nonmoral* a judgment that we would ordinarily describe as moral—that was a consideration that counted against the definition. For Hare and some of the other leading figures in the debate, these sorts of linguistic considerations were the only source of evidence relevant to evaluating a definition, since the goal of the exercise was to capture the concept of moral judgment underlying ordinary usage. However, other central figures in the debate urged that this sort of descriptive conceptual analysis is one of two quite different goals that a philosopher might have when attempting to defend a definition of morality. The other goal is *conceptual revision*—characterizing a new concept of morality that will be better suited to playing a role in philosophical theory construction. William Frankena drew the distinction between these two projects very clearly and argued that conceptual revision is both legitimate and important.

[O]ur question and our answer to it may take two forms. For when we ask what morality is or what is to be regarded as built into the concept of morality, we may be asking what our ordinary concept of it is or entails, what we actually mean by "moral" and "morality" in their relevant uses, or what the prevailing rules are for the use of these terms. . . . However, when one asks what morality is or how it is to be conceived, one may be interested, not so much in our actual concept or linguistic rules, as in proposing a way of conceiving it or a set of rules for talking about it, not so much in what our concept and uses are, as in what they should be. If the questions are taken in the first way, the discussion will be a descriptive-elucidatory one, and the arguments pro and con will have a corresponding character; if they are taken in the second sense, the inquiry will be normative, and the arguments will have a different character, though, of



course, one may still take the fact that we actually think and talk in a certain way as an argument for continuing to do so.

Now, most recent philosophers who have dealt with our topic have been shy about making proposals of a normative sort. . . . Though some of them do at least favor one way of speaking against another, they tend to try to rest wholly on the basis of actual use and its rules. Indeed, they have tended to think that philosophers as such should not venture to propose revisions of our moral concepts, since to do so is to make a normative or value judgment, . . . and the business of philosophy is or should be (a normative judgment!) “analysis” or “logic.” . . . But if one may or must be normative at all, then in principle there is no reason why one may not be revisionary, especially if one finds difficulties and puzzles in our ordinary manners of thought and expression. In what follows, at any rate, I shall take it to be appropriate for a philosopher to ask whether something should be built into our concept of morality, even if it is not. . . . I shall take our problem to be primarily a normative rather than a descriptive–elucidatory one.

(Frankena, 1967, 149–150)

In an earlier paper, Frankena offers a memorable summary of this approach:

Defining terms like “moral judgment” may be part of an attempt to understand, re-think, and possibly even to revise the whole institution which we call morality, just as defining “scientific judgment” may be part of an attempt to do this for science.  
(1958, 45)

As Frankena notes, he is not alone in viewing the project of defining “moral” and “moral-ity” as primarily revisionary and normative. Von Wright (1963, 4–5) had adopted a similar view, and in later years Cooper (1970, 93), Rawls (1971, §23) and Paul Taylor (1978) did so as well.

#### 4. What Characteristics Were Proposed; Which Were Agreed on?

Since some of the philosophers engaged in the debate over the definition of morality adopted a “descriptive–elucidatory” approach while others viewed the project as revisionary and normative, it is hardly surprising that no consensus was reached on how “moral rule,” “moral judgment” and the rest should be defined. There is a long list of features that were argued to be necessary conditions. Perhaps the most widely discussed of these was Hare’s proposal that moral rules must be “universalizable.” As Hare unpacked the notion, it required that there be no names or definite descriptions in moral rules, only predicates. While the predicates could have a very restricted extension—“people who have four left-handed grandparents” would be fine—the rule applies to *everyone* to whom the predicate applies, no matter where they might be or when they might live. Another widely discussed proposal, also due to Hare, was that moral judgments are “prescriptive.” What this means is that the

action-guiding force [of moral rules] derives from the fact that they entail imperatives: my acceptance of the principle “One ought to do X” commits me to accepting



the imperative “Let me do X”; and my acceptance of the imperative commits me in turn to doing X in the appropriate circumstances.

*(Wallace & Walker, 1970, 9)*

A third proposal was that if an action guiding principle is a moral principle for a person, then she must regard it as “overriding or supremely important” (Frankena, 1967, 155). Moral norms “outweigh, as grounds of reasons-for-action, all other kinds of norms. In cases of conflict between moral and nonmoral principles, the former are necessarily overriding” (Taylor, 1978, 44; for a similar proposal, see Cooper, 1970, 95). A related idea is that moral judgments are “categorical.” According to Gewirth (1978, 24) “Judgments of moral obligation are categorical in that what persons ought to do sets requirements for them that they cannot rightly evade by consulting their own self-interested desires or variable opinions, ideals, or institutional practices.” Another frequently discussed necessary condition was that moral rules are behavior guiding rules whose violation is met with social sanctions, “the reproach of one’s neighbors” (Cooper, 1966, 73) or something more serious, like ostracism (Sprigge, 1964, 129 ff). This was sometimes paired with the idea that moral transgressions are followed by the transgressor sanctioning himself with feelings of guilt or shame or disliking himself (Wallace & Walker, 1970, 14; Sprigge, 1964, 130). Yet another proposed necessary condition was that if two people share the same factual beliefs then their moral judgments will be the same. So if people who share their factual beliefs continue to disagree, then at least one of them is not really expressing a moral judgment (Frankena, 1963, 5–6). All of these proposals were “formal” in the sense that they did not impose any constraints on the contents of moral rules or moral judgments. And this is far from a complete list of the formal conditions that were proposed; there were many more.<sup>7</sup>

There was no shortage of critics for these formal conditions. Wittgensteinians, who maintained that “moral” was a family resemblance term, denied that there are any strictly necessary conditions for the application of the term. MacIntyre (1957), inspired by Sartre, argued that many moral judgments were neither universalizable nor (in Hare’s sense) prescriptive. Sprigge (1964) offered a quite different argument against universalizability. And so it went. I think it is fair to say that nothing on this list of proposed formal conditions achieved anything even close to consensus during the three decades during which the Philosophers’ Project was most active.

Even more controversial was the question of whether more substantive social requirements should be built into the definition of morality. For example, Frankena urged that a necessary conditions for a set of rules being a morality should be that

it includes judgments, rules, principles, ideals, etc., which [(i)] concern the relations of one individual . . . to others [and (ii)] involve[s] or call[s] for a consideration of the effects of his actions on others (not necessarily all others), not from the point of view of his own interests or aesthetic enjoyments, but from their own point of view.

*(Frankena, 1967, 156)*

This condition allows in a wide variety of deontological and utilitarian moralities, but “it rules out as non-moral . . . such [action guiding systems] as pure egoism or prudentialism,



pure aestheticism, and pure religion” (157). It does not rule out “Nazi ethics,” which requires an individual to consider the effects of his actions on fellow Germans, but on some readings of Nietzsche, on which the proposed action guiding rules are purely egoistic or aesthetic, the condition entails that Nietzsche is not proposing a morality at all. Baier (1958, 199ff) proposed a similar but stronger condition on which moral rules “must be for the good of everyone alike.” Earlier, Toulmin (1950) had argued that a concern for the harmony of society is part of the meaning of “moral.” On these substantive principles, too, it is clear that no agreement was reached.

### 5. Why Did Contributions to this Literature Gradually Diminish?

According to General Douglas MacArthur, “Old soldiers never die, they just fade away.” Much the same could be said for many philosophical debates. During the last decade of the twentieth century, discussion of the definition of morality gradually faded from the philosophical literature.<sup>8</sup> The reason for this was certainly not that the problem of defining morality had been solved or that agreement had been reached. Nor was it the case that the importance of the issue had declined. Quite the opposite, as we saw in §3. Rather, I suspect, it was because most of the main options had been pretty thoroughly explored and promising new ideas and arguments were hard to come by. Moral philosophers turned their attention to newer issues. Perhaps the waning of the positivist-inspired prohibition against philosophers making “value judgements” also played a role. Whatever the reason, debates over the definition of morality no longer loomed large in leading journals. However, as philosophical discussion of the definition of morality wound down, the topic was moving to center stage in empirical moral psychology. That will be our topic in subsequent sections. But before getting to that, I want to briefly discuss a more recent challenge to the Philosophers’ Project.

### 6. Some Recent Work Relevant to the Philosophers’ Project

Those engaged in the Philosophers’ Project were trying to provide an analysis of concepts like moral judgment and moral rule, and it is clear that for most of these philosophers, the analysis they sought would provide necessary and sufficient conditions.<sup>9</sup> Moreover, those who took the project to be “descriptive-elucidatory” rather than normative wanted their account to capture the concept we actually use. That project did not meet with much success. As Jerry Fodor has famously noted, such projects rarely do (Fodor, 1981, 283). However, it might be thought that the failure of the Philosophers’ Project could be traced to the quest for an analysis providing necessary and sufficient conditions. The view that most concepts can be analyzed in this way has become known as the classical theory of concepts, and both empirical and philosophical work on concepts over the last four decades has made a convincing case that the classical theory of concepts is false for most ordinary concepts (Smith & Medin, 1981; Laurence & Margolis, 1999). There are, however, a variety of other ways of analyzing concepts utilizing prototypes, exemplars, commonsense theories or other approaches (Machery, 2009). So perhaps the descriptive-elucidatory project could be successfully revived by dropping the demand for necessary and sufficient conditions and adopting one of these alternative approaches to conceptual analysis.



However, recent work in moral psychology and experimental philosophy poses a challenge to this hopeful thought. It raises another, less tractable, problem for the descriptive-elucidatory project. Inspired by the work of cultural psychologists, experimental philosophers have been exploring the possibility that philosophical intuitions—spontaneous judgments about whether a familiar term applies to a real or hypothetical case—may vary in different demographic groups.<sup>10</sup> It is widely assumed that concepts play a central role in generating philosophical intuitions (Goldman, 2007). So if intuitions vary across demographic groups—and there is a growing body of evidence that they do—then philosophically important concepts may also vary in different demographic groups.

In a recent study, Levine et al. (under review) explored whether there were demographic differences in people's concept of a moral judgment. They asked American participants of different religious faiths—Mormon, Muslim, Hindu, Jewish and secular—to judge whether a long list of normative judgments were moral judgments or some other kind of judgment, and they found striking differences between these five groups. On the basis of this work, the authors suggest that there are important differences in how the adherents of different religions conceive of morality. Using very different methods, Buchtel et al. (2015) have shown that Chinese and Westerners classify different transgressions as moral, and Wright et al. (2013) have shown that there is considerable variation when American college students are asked whether an issue is a moral issue.<sup>11</sup>

If, as this work suggests, different people and different groups have different concepts of morality, then the goal of the descriptive-elucidatory project is underspecified in an important way. That goal, as we've seen, is to capture "our" concept of morality, the concept of morality that "we" actually use. But who are "we"—secular people, Jews, Mormons, Muslims or Hindus? Chinese or Westerners? And however this question is answered, why is *our* concept of morality more important than the concept employed by other groups? Why is it that *our* concept provides the answer to the philosophical questions posed in parts 1 and 2 of this chapter? I have no idea how to answer these questions. Without convincing answers, the descriptive-elucidatory project, when no longer committed to the classical theory of concepts, may be a fascinating exercise in cognitive anthropology, but it is hard to see why it is of any philosophical interest.

## 7. The Psychologists' Project (≈1970–the present): Turiel's Account of Moral Judgment

The psychologists' project that will be center stage in this section grows out of the work of Elliot Turiel and his colleagues. Turiel was a student of Lawrence Kohlberg whose influential work on moral reasoning and moral development was widely discussed and enormously influential in the 1970s, '80s and '90s. Following Piaget, Kohlberg held that moral reasoning emerged in stages. For young children, according to Kohlberg, morality is largely a matter of obedience and punishment. Children judge that certain behaviors are wrong because they know those behaviors are likely to be punished, and their understanding of wrongness is, near enough, exhausted by the idea of punishment: wrong behavior just is behavior that is typically punished.<sup>12</sup> Turiel, by contrast, was convinced that moral cognition is distinct from other sorts of cognition and that it emerges quite early in development. In order to make the case for this claim, he had to show that children could make characteristically



moral judgments. And to do that Turiel needed a test that would indicate when an experimental participant—child or adult—was making a moral judgment.

It was at this point that the Philosopher's Project played a crucial role in the development of the Psychologists' Project, as Turiel turned to the philosophical literature for a characterization of moral judgments. Several of the necessary conditions that philosophers had proposed were endorsed by Turiel and incorporated into his own account of moral judgments. One of these was universalizability. "Moral prescriptions," he tells us, "are *universally applicable* in that they apply to everyone in similar circumstances" (Turiel, 1983, 36; italics in the original). So if a young participant in an experiment judges that it is wrong for a child in her own school to push someone off a swing, and if that judgment is a moral judgment, we would expect the participant to say that it is also wrong for a child in another school to push someone off a swing. A second feature discussed in the philosophical literature that was adopted by Turiel was the categoricalness of moral judgments. He quotes with approval the passage from Gewirth that I quoted in §4:

Judgments of moral obligation are categorical in that what persons ought to do sets requirements for them that they cannot rightly evade by consulting their own self-interested desires or variable opinions, ideals, or institutional practices.

(Gewirth, 1978, 24, quoted in Turiel, 1983, 35)

Since institutional practices cannot alter moral obligations, we should expect that if an experimental participant has judged that it is wrong to push someone off a swing and that judgment is a moral judgment, then the participant would judge that it would be wrong in another school where there was no rule against pushing people off a swing, and it would be wrong even if the principal in her own school said that there was no rule against it. In the jargon that has developed in the literature growing out of Turiel's work, these questions are said to probe for "authority independence." The test that Turiel proposed to determine whether a judgment is a moral judgment includes one or more questions assessing whether the participant takes her judgment to be universalizable and one or more questions assessing whether she takes her judgment to be authority independent.

Both universalizability and categoricalness are "formal"—they do not impose any constraints on the content of moral rules or moral judgments. But Turiel also held that there are substantive features that all moral judgments share. They all, he maintained, deal with issues linked to harm, justice or rights. Thus if an experimental participant has made a genuinely moral judgment and is asked to explain why the behavior in question is wrong, she will typically appeal to the harm that has been done or to injustice or the violation of someone's rights. In building substantive features into his characterization of moral judgments, Turiel is siding with Toulmin, Frankena, Baier and others who argued against a purely formal characterization of morality, though there is no indication that Turiel was aware of the debate between the formalists and their philosophical critics. Moreover, Turiel's choice of substantive features—those linked to harm, justice and rights—was quite different from those proposed by the philosophical anti-formalists and was motivated by his account of how children acquire moral rules.

With these three putative features of moral judgments in hand, Turiel proceeded to construct an empirical test to determine whether an experimental participant's judgment



about a transgression is a *moral* judgment. The test typically begins with a brief vignette describing a hypothetical transgression. Since Turiel was interested in determining whether young children made moral judgments, the transgressions almost always involve events that would be familiar to kids. The participant is then asked a series of questions aimed at determining whether she thinks the action described is wrong, whether she thinks wrongness of the action in the vignette is “authority independent,” and whether the participant would universalize the judgment, making the same judgment if the transgression occurred at another place or time. These questions can be asked in a variety of ways depending on the age of the participant and the goals of the study. The participant is also asked to explain why the transgression is wrong, and responses are assessed to determine whether the participant invokes harm, justice or rights or whether she invokes other considerations (including custom, tradition, appeal to authority, disrupting social coordination or the likelihood of punishment) that, Turiel maintains, are the sorts of justifications that are to be expected for “conventional” transgressions (Turiel, 1983, 67). This experimental paradigm, in which a transgression is described and participants are asked questions to determine (i) whether they think it is wrong, (ii) how they would justify that judgment, (iii) whether their judgment is authority independent and (iv) whether they universalize the judgment, is frequently referred to as the *moral/conventional task*.

Another question often asked along with the four listed here is aimed at determining how serious the participant thinks the transgression is. Of course some moral transgressions are more serious than others, and some conventional transgressions are more serious than others. But since a number of philosophers have proposed that moral considerations are “overriding,” one might think that moral transgressions should always be considered more serious than conventional transgressions. Turiel and his followers reject this idea (Tisak & Turiel, 1988, 356) and report a number of studies in which participants judge that egregious conventional transgressions, like a boy wearing a dress to school, are more serious than minor moral transgressions like stealing an eraser (Turiel, 1983, 71). Thus, as Smetana notes, “the severity of the transgression is not considered to be a formal criterion for distinguishing moral and conventional rules and transgressions” (1993, 117).

Before proceeding, let me introduce a bit of jargon (mine, not Turiel’s) that will prove useful. The pattern of responses in the moral/conventional task that Turiel takes to be characteristic of a moral judgment are universalizability (U), authority independence (I) and justification by appeal to harm, justice or rights (H). I will call this the *UIH response pattern*. Turiel takes the opposite pattern—not universalizable, not authority independent and not justified by appeal to harm, justice or rights—to be characteristic of conventional normative judgments. I’ll call that the *~U~I~H response pattern*.

By using the moral/conventional task with youngsters, Turiel and his collaborators were able to show that they typically gave the UIH response pattern to vignettes describing what they thought adults would consider moral transgressions, and the *~U~I~H* response pattern to vignettes describing what they thought adults would describe as conventional transgressions. Turiel concluded that children can indeed make moral judgments at an age when Kohlberg’s theory predicted that they were only capable of conceptualizing morality in terms of punishment. More importantly, he concluded that young children have a basic grasp of the distinction between moral and conventional rules and transgressions.





## 8. A Critique of Turiel's Account of Moral Judgment, and a Response

Against the backdrop of the philosophical literature discussed in §1, one might well think that there is something seriously wrong with all this. Philosophers spent decades debating how “moral judgment,” “moral rule” and the rest should be defined without reaching any widely accepted conclusion. Turiel offered no additional evidence about the ordinary usage of these terms; he contributed nothing to the “descriptive-elucidatory” project of analyzing our ordinary concept of moral judgment. Nor did he offer any normative argument aimed at showing how our ordinary concept should be revised. Rather, it seems, he simply *stipulated* that moral judgments are universalizable, authority independent and justified by appeal to harm, justice or rights, and that the UIH response pattern can be used to identify moral judgments. But if the term “moral judgment” is supposed to have its ordinary meaning, then one *can't* just stipulate how moral judgments are to be identified. If, on the other hand, Turiel proposes to use “moral judgment” as a *technical term*, he is free to make whatever stipulations he wishes about how moral judgments (in the technical sense) are to be identified. However, if “moral judgment” is a technical term, then one cannot infer that moral judgments (in this technical sense) have anything to do with moral judgments as the term is usually used. So showing that children make moral judgments (in the technical sense of judgments exhibiting the UIH pattern) tells us exactly nothing about whether they make moral judgments in the ordinary sense. And, of course, the same is true of adults. Without some further argument, one cannot infer from the fact that an adult's judgment exhibits the UIH pattern to the conclusion that the adult has made a moral judgment.<sup>13</sup>

All of this, I think, is exactly right. But there is another way of construing Turiel's project—and much of the literature that it generated—that avoids these problems.<sup>14</sup> In a seminal paper published 40 years ago, Hilary Putnam (1975) famously argued that in many cases, “meanings just ain't in the head.” When the term in question is a natural kind term, like “water” or “fish” or “gold,” Putnam urged, it is the job of empirical science to determine the essential features of the natural kind, and these essential features constitute the correct definition of the kind. Other philosophers, notably Devitt (1996) and Kornblith (1998), have provided insightful accounts of how this process works. Very roughly, their story goes like this. To begin, the scientist focuses on intuitively prototypical examples of the kind in question. She then looks for properties that are shared by most of these prototypical examples. If she finds a cluster of properties that are present in most prototypical examples and absent in most things that, intuitively, are not members of the kind, she hypothesizes that that cluster of properties are the essential features of the kind.

It is a reasonable hypothesis that the ordinary term “moral judgment” is a natural kind term, picking out a psychological natural kind. If so, it is the job of science—psychology in this case—to determine the essential features of the natural kind. One way to do this would be for psychologists to discover a cluster of nomologically linked properties that are shared by many (but perhaps not all) cases of what they would intuitively take to be prototypical moral judgments and that are missing in many (but perhaps not all) cases of what they would intuitively take *not* be a moral judgment.<sup>15</sup>

With this by way of background, let's return to Turiel. In his book-length exposition of his research program, Turiel tells us that the “strategy of the research” he reviews in several chapters was to present subjects with “prototypical examples” of moral and conventional



transgressions as a means of investigating whether UIH judgments<sup>16</sup> are evoked by moral transgressions and  $\sim U \sim I \sim H$  judgments are evoked by conventional transgressions (Turiel, 1983, 55). His claim that UIH judgments are moral judgments can be interpreted as a hypothesis about the essential features of moral judgments. If the hypothesis is true, we would expect that the three components of UIH judgments are nomologically linked—they typically occur together. We would also expect that many UIH judgments are intuitively classified as prototypical moral judgments, and many  $\sim U \sim I \sim H$  judgments are intuitively classified as prototypical conventional judgments.

To make a persuasive case for that hypothesis, we would need lots of experiments, using a wide range of prototypical transgressions and many different participant populations. Over the years Turiel and his colleagues have conducted moral/conventional task experiments on many different groups of experimental participants. Findings supporting the hypothesis that the UIH response pattern is a nomological cluster, and thus that the UIH pattern captures the essence of moral judgments, have been found in participants ranging in age from toddlers to adults (Nucci & Turiel, 1978; Smetana, 1981; Nucci & Nucci, 1982), in participants of a number of different nationalities and religions (Nucci et al., 1983; Hollos et al., 1986; Yau & Smetana, 2003; for reviews, see Smetana, 1993; Tisak, 1995; Nucci, 2001) and in children with a variety of developmental disorders, including autism (Blair, 1996; Blair et al., 2001; Nucci & Herman, 1982; Smetana et al., 1984; Smetana et al., 1999). In response to this impressive body of evidence, many psychologists and a growing number of philosophers have accepted the moral/conventional task as a reliable way of identifying moral judgments.<sup>17</sup>

## 9. The Case against the Hypothesis that Moral Judgments Are a Natural Kind Evoking the UIH Response

While there are many studies that can be interpreted as supporting the hypothesis that moral judgments are a natural kind evoking the UIH response, the evidence for the claim that the components of the UIH package form a nomological cluster is far from uniform. Early studies indicating that UIH components do not always occur together focused on transgressions that do not involve harm (or justice or rights). Nissan (1987) used the moral/conventional task in a study that included children in traditional Arab villages in Israel. Among the transgressions that Nissan used were mixed-sex bathing and addressing a teacher by his first name—behaviors in which no one is harmed. He found that these children considered those transgressions to be universalizable (U) and authority independent (I). So, contrary to the hypothesis that the UIH package is a nomological cluster, in this study, U and I are not linked to H. In another study, Nucci and Turiel (1993) found that orthodox Jewish children in the USA judged a number of religious rules to be authority independent (I) even though they did not involve harm (or justice or rights). So in this study, I and H are not linked, contrary to the nomological cluster hypothesis. And in what is surely the most famous and most memorable study aimed at showing that the UIH cluster comes apart, Jonathan Haidt and colleagues used transgressions like washing the toilet bowl with the national flag and masturbating with a dead chicken (Haidt et al., 1993). Though Haidt's participants agreed that none of these behaviors were harmful, his low socioeconomic status participants in Brazil and in the USA nonetheless said that the behaviors were wrong and indicated that their judgment was universalizable (U) and authority



independent (I)—again U and I without H. In another important study, Nichols (2002) used examples of disgusting but harmless etiquette transgressions. He found that American children judged them to be universalizable (U) and authority independent (I)—still another example of U and I without H. Moreover, in the same study, Nichols found that American college students judged these etiquette transgressions to be authority independent though *not* universalizable. So with these participants, I has become detached from both U and H. Taken together, these studies pose a serious challenge to the claim that the elements of the UIH package form a nomological cluster.

All of the studies mentioned in the previous paragraph used transgressions that did not involve harm but nonetheless evoked other elements of the UIH package. In a 2007 study, Kelly et al. set out to explore participants' reactions to transgressions that do involve harm. There had, of course, been many studies by Turiel and his followers in which a harmful transgression was linked to U and I. But Kelly and colleagues noted that in almost all of these studies the harmful transgressions were restricted to the sorts of behaviors that young children might encounter. This was true even of a study in which the participants included incarcerated psychopathic murderers (Blair, 1995). So Kelly and colleagues decided to focus on transgressions that are not encountered in the schoolyard, including slavery, serious corporal punishment (whipping a sailor who was drunk on duty) and physically abusing military trainees. They found that many participants judged that these sorts of transgressions were *not* authority independent. According to these participants, it is OK to physically abuse military trainees if it is not prohibited by the authorities, but it is not OK if it is prohibited. Kelly and colleagues also found that the judgments of many participants do not generalize over time and space. Whipping a drunken sailor is not acceptable now but was acceptable 300 years ago. Slavery is not acceptable now but was acceptable in ancient Greece and Rome. So in this study, too, the UIH package comes unstuck. We find H without U or I.

The Kelly et al. study was motivated by the observation that previous moral/conventional task studies had not used a wide range of harmful transgressions; they were almost all of the "schoolyard" variety. Another, more recent, study which also used "grown-up" transgressions was undertaken because previous studies, though they included a number of different demographic groups, had all focused on participants in large-scale, relatively modern societies (Fessler et al., 2015). Fessler and colleagues decided to explore what would happen to the UIH package if grown-up transgressions were used in small-scale societies. Using transgressions like stealing, wife battery, marketplace cheating, defamation, perjury and rape, they collected data in five small-scale societies and two large-scale modern societies. They 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, again challenging the claim that there is a nomological link between H and U. Endorsement by an authority figure had this effect in four of the seven societies, with the remaining three showing nonsignificant trends in the direction of reduced severity—another challenge to the nomological link between H and I. So we now have evidence that Turiel's putative nomological cluster comes apart with grown-up transgressions in a number of societies, including small-scale societies.

The lesson that I am inclined to draw from the studies discussed in the last three paragraphs is that the UIH pattern is not a nomological cluster and thus that the elements of that cluster are not the essential features of a natural kind. If that's right, then they can't be used to construct an empirically supported definition of morality. One way in which this



conclusion might be challenged is to critique the methods or analyses of the studies cited. This has been done by a number of authors, and lively debates have ensued. My own view is that the critics have not been very successful. But I am hardly an impartial observer, so I'd encourage you to make your own assessment.<sup>18</sup>

### 10. Another Natural Kind Account of Moral Judgment

Another reaction to the studies reviewed in the previous section would be to offer an empirically informed alternative to the UIH cluster—a different account of the essential features of moral judgments. That's the strategy adopted by Kumar (2015). The first step in Kumar's proposed revision is to urge that the third element in Turiel's cluster, the requirement that moral judgments be justified by appeal to harm (or justice or rights) should be abandoned. His argument for this move seems to turn on intuition, or on how things "seem":

[F]olk theories about how moral claims are justified do not seem to be part of the concept of morality. . . . [I]t would seem that many people gain a facility with moral concepts before they have any theory about what grounds them. Justificatory grounds, whatever role they may play in marking important boundaries in moral philosophy, are not internal to the ordinary concept of morality.

*(Kumar, 2015, §3)*

I confess that I do not have intuitions on such rarefied matters as what is internal to the ordinary concept of morality. But there is no need to dispute these claims since, as Kumar makes clear, he is offering an alternative to Turiel's hypothesis about the essential features of moral judgment, and he is free to include, or exclude, whatever features he wishes. The crucial question is whether the set of features he proposes actually do form a nomological cluster. Dropping the requirement that moral judgments must be justified by harm is certainly a strategically wise move for Kumar, for it enables him to ignore some of the best-known and most persuasive critiques of Turiel. The fact that Haidt's low SES participants judge that transgressions not involving harm are authority independent and universalizable is not a problem for Kumar, since his theory—which he calls "MCT"—does not predict that U and I will be nomologically linked to harm.

The second step in Kumar's revision is to add a feature that does not occur in Turiel's account. Over the last decade, there has been growing interest in the question of whether ordinary folk are moral objectivists or moral relativists. To explore the issue, a number of investigators have presented participants with moral claims like "Consciously discriminating against someone on the basis of race is morally wrong," along with factual claims like "Homo sapiens evolved from more primitive primate species" and conventional claims like "Wearing pajamas and bathrobe to a seminar meeting is wrong behavior."

After determining that a participant agrees with the statement, the participant is told about someone who disagrees, and asked to choose among the following options:

1. The other person is surely mistaken.
2. It is possible that neither you nor the other person is mistaken.
3. It could be that you are mistaken, and the other person is correct.



If the participant selects (1) or (3), it is taken to be evidence that the participant is an objectivist about the claim. Selecting (2) is taken to be evidence that the participant is a relativist.<sup>19</sup> In the earliest studies (Nichols, 2004b; Goodwin & Darley, 2008), participants' responses in what the investigators took to be prototypical moral cases were usually similar to their responses in the scientific cases—they thought that one of the disputants must be wrong. Based on these findings, Kumar hypothesizes that objectivity is a feature of the nomological cluster that defines the concept MORAL.

A third step in Kumar's revision is to upgrade seriousness to an essential feature of moral judgments. To justify the move, he says that "research suggests that morality is unlike convention in that morality is serious" (Kumar, 2015, §2) and cites several studies in the Turiel tradition in which a seriousness question was included in a moral/conventional task experiment. His review of the moral/conventional task literature also leads him to endorse the claims that universality (being "general") and authority independence are features of moral judgments. The upshot of all this is summarized in the following passage.

We are now in a position to say what defines MORAL. A moral wrong, for instance, is a wrong that is

- (1) serious
- (2) general
- (3) authority-independent
- (4) objective

[T]he four features that define MORAL are stable and mutually reinforcing. Moral judgment, like other natural kinds, is a homeostatic property cluster. . . . The human cognitive system is organized in such a way that the four features have a nomological tendency to cluster together.

(Kumar, 2015, §4)

For three quite different reasons, I find Kumar's hypothesis unconvincing. First, the literature on folk moral objectivism is much more contested than Kumar suggests. Since the Nichols (2004b) and Goodwin and Darley (2008) papers were published, there have been a number of studies suggesting that the folk are not moral objectivists (Sarkissian et al., 2011) or that they are objectivists on some moral issues and not on others (Goodwin & Darley, 2012) and that participants' responses in experiments like these are influenced by a wide range of factors including the age of the participant, how the person who disagrees is described, how controversial the issue is and whether the moral claim in question is about a bad action or a good one (Wright et al., 2013; Beebe, 2014, 2015; Beebe & Sackris, 2016).

The second reason is that the format of the studies cited in the previous paragraph offers no evidence that objectivity forms a nomological cluster with other items on Kumar's list. In these studies participants are presented with a sentence like "Consciously discriminating against someone on the basis of race is morally wrong" and asked questions designed to determine whether they are objectivists about those statements. Participants are *not* asked anything about seriousness, generalizability or authority independence. So Kumar is simply *speculating* that if they had been asked participants would judge that the transgression



described is serious, generalizable and authority independent.<sup>20</sup> Of course, Kumar's speculation *might* be true. But at this point there is no evidence at all that it is. To make a serious case that Kumar's four features form a nomological cluster we would need studies that test participants on all four features, and at this writing there are no such studies.

Finally, Kumar has misinterpreted the role that seriousness plays in the Turiel tradition.<sup>21</sup> As noted in §7, Turiel and his followers do not take seriousness to be "a formal criterion for distinguishing moral and conventional rules and transgressions." This was clearly a wise move on their part. For while there are a number of studies in which participants judge that the schoolyard moral transgressions used are more serious than the conventional, there are also studies in which conventional transgressions are judged to be more serious than moral transgressions. Moreover, when one reflects on the vast range of possible non-schoolyard transgressions, the claim that moral transgressions are more serious than conventional transgressions is singularly implausible. Though no one has done the experiment, I would be willing to bet that most people would judge that showing up naked at your grandmother's funeral is more serious than stealing an eraser! Kumar tells us that his "moral/conventional pattern is not supposed to be exceptionless," and that his MCT only claims that "the features usually cluster together" (Kumar, 2015, §5). But it is hard to come up with a sensible interpretation of what "usually" could mean here. If the claim is that most *actual* moral transgressions are more serious than most *actual* conventional transgressions, then we have no evidence that would support the claim, and we never will, since most transgressions of both sorts have never been recorded. If the claim is that most *possible* moral transgressions are more serious than most *possible* conventional transgressions, then Kumar will have to explain how we are to go about comparing these two infinite sets of transgressions. Perhaps there is some more plausible interpretation of Kumar's claim. But I have no idea what it is.

The bottom line, I think, is that Kumar's MCT is no more successful than Turiel's theory in providing a defensible account of a nomological cluster of properties that can be used in an empirically supported definition of moral judgment.

## 11. Summing Up and Looking Beyond: A Future Without "Morality"

Most of the philosophers who contributed to the Philosopher's Project were convinced that there is a correct or well-motivated way of dividing normative judgments into those that are *moral* and those that are *nonmoral*. But, as we saw in §1–§5, those philosophers who took their project to be "descriptive-elucidatory"—aimed at providing an analysis of the concept of moral judgment that we actually use—did not meet with much success. Though some of the necessary conditions that were proposed were widely accepted, no set of necessary and sufficient conditions convinced more than a handful of contributors to the literature. Those who took their project to be normative were, if anything, less successful. Most of the normative analyses were, at best, very sketchy. And more often than not they were not endorsed by anyone but the author.

In §6 we noted that the failure of both the descriptive-elucidatory and the normative projects might be blamed on a commitment to the classical theory of concepts and that things might go better if that commitment was abandoned in favor of some other account of concepts. But there are other challenges facing those pursuing descriptive-elucidatory project. There is some evidence that people in different religious or cultural groups, and



even perhaps people who share their religion and culture, have notably *different* concepts of moral judgment. It is, I believe, too early to draw any confident conclusions from the evidence available; much more work is needed. But if it is true that there are religious, cultural and individual differences in people's concept of moral judgment, then the descriptive-elucidatory project is both poorly specified and poorly motivated. The goal of that project is to analyze the concept of moral judgment that *we* actually use. But if there are significant interpersonal and intergroup differences, we need to be told who "we" refers to. We also need to be told why *our* concept—however "our" is unpacked—is of any special philosophical importance. Why, for example, should *our* concept be the one to use in deciding whether the Navajo have a moral code? Philosophers are very clever people. So perhaps this challenge can be met. But at this point, I know of no serious attempts.

In §7–§10 we explored the idea that "moral judgment" might be a natural kind term with a definition that can be discovered by psychologists. Turiel's project fits comfortably into this picture. But a growing body of evidence suggests that Turiel's UIH cluster shatters in a variety of ways and thus that it is not a nomological cluster at all. Here too, much more work is needed. For as John Doris has eloquently reminded us, in any given experiment in psychology, there is a *lot* that can go wrong. So it is wise to wait until there are many experiments all pointing in the same direction (Doris, 2015, 44–49). Kumar's alternative natural kind account of moral judgment is, I think, less promising than Turiel's. It requires that objectivity judgments form a nomological cluster with seriousness, universality and authority independence judgments, and at this writing there is no evidence at all for that claim. But though I'm critical of Kumar's theory, I think his strategy is a good one. If we are to find a well-motivated way of defining "moral judgment" and related terms, our best hope is to locate a nomological cluster of properties exhibited by many intuitively prototypical moral judgments but not by most intuitively prototypical nonmoral normative judgments. Finding such a cluster would be an important discovery for both moral philosophy and moral psychology.

There is, of course, no guarantee that the quest for a nomological cluster account of "moral judgment" will succeed. For it may turn out that there simply is no natural kind to be found in this vicinity or that there are numerous natural kinds, none of which can sustain a compelling argument that it specifies the essential features of moral judgments. What would be the consequences if *that* is how things unfold? To make things easier, let's also assume that neither the descriptive-elucidatory nor the normative project is successful, and that these three projects are the only options available for those who seek a well-motivated way of defining "moral judgment."

Perhaps the most obvious implication of the failure of these projects is that debates that turn on whether specific normative judgments are really moral judgments will turn out to be irresolvable because they are based on a mistaken assumption. Consider, for example, Jonathan Haidt's accusation that the preponderance of politically liberal researchers has led to "an inappropriate narrowing of the moral domain." As we saw in §2, Haidt's accusation turns on his insistence that norms governing such matters as clothing, gender roles, food and forms of address are *moral* norms, and whether judgments about such matters are *moral* judgments. Haidt insists they are. Turiel insists they aren't. If our assumptions are correct, then there is simply no fact to the matter. Much the same is true for those who would debate whether the Navajo, as described by Ladd, have a moral code at all.

Let's turn, now, to those many philosophers who debate the semantics of moral judgments, the function of moral judgments, the evolutionary history of moral judgments and the psychological mechanisms underlying moral judgments. How would their projects be impacted if our assumptions are correct? Here the consequences are less dire. To be sure, if the parties to these debates focus on different examples, and if one side insists that the examples used by the other side are not really *moral* judgments at all, then the debate is irresolvable, since once again there is no fact to the matter. But this is not how most debates on these topics unfold. Rather, in most cases at least, the philosophers involved agree that the examples of moral judgments advanced by their opponents really are moral judgments. So what they are debating is the semantics, or the function, or the evolutionary history or the psychological mechanisms of judgments *like those*. And progress can be made without specifying the boundaries of that class. However, if it turns out, and I'm betting it will, that there are actually a number of different natural kinds included in that vaguely specified class, then future philosophers and psychologists may simply drop the term "moral judgment" and focus instead on judgments of these separate natural kinds. If that's the way things unfold, both philosophers and psychologists may be destined for a future without "morality."<sup>22</sup>

### Notes

- \* This paper is dedicated to the memory of William Frankena, my esteemed colleague at the University of Michigan during the first decade of my career.
- 1. The terms "action guide" and "action guiding" are borrowed from Frankena (1967).
- 2. Oddly, this project was not the primary focus of moral philosophers in the analytic tradition during the 1950s and 1960s. More on this later.
- 3. The example is borrowed from Stohr (2012).
- 4. Many of these philosophers would have agreed with Frankena, who maintained that having an account is the *only* way to settle this question. "One cannot say that the Navaho have a morality until after one has formed some conception of morality and found that the Navajo have such an institution" (Frankena, 1963, 17).
- 5. See, for example, Wallace and Walker (1970, 1) and MacIntyre (1957, 26).
- 6. Joyce's account is one of the few philosophically sophisticated analyses to appear since the turn of the century. For some critical thoughts about that analysis, see Stich (2008). Southwood (2011) offers another philosophically sophisticated analysis. See too, this volume, Chapter 2 ("The Normative Sense: What is Universal? What Varies?"); Chapter 3 ("Normative Practices of Animals"); Chapter 12 ("The Denial of Moral Knowledge"); and Chapter 14 ("Nihilism and the Epistemic Profile of Moral Judgment") for analyses of moral judgment.
- 7. Frankena (1963) offers a much more extensive list, along with many references.
- 8. Though they did not completely disappear from the literature. See Gert's *Stanford Encyclopedia of Philosophy* article, "The Definition of Morality," which was first published in 2002 and has undergone four "substantive content changes," in 2005, 2008, 2011 and 2016.
- 9. Frankena, for example, tells us that the question he is asking is "What are we to take as the necessary and sufficient conditions for something's being or being called moral or a morality" (Frankena, 1967, 146–147).
- 10. See Chapter 18 of this volume for more on the variability of intuitions.
- 11. For further discussion of research on variability in conceptions of morality, see Chapter 2 of this volume.
- 12. For an informed and insightful account of Kohlberg's work, see Lapsley (1996), chapters 3 & 4; see too Chapters 5 and 16 of this volume.



13. It is a striking fact that a number of philosophers engaged in the Philosophers' Project insisted that the definition of "moral judgment" must include a "material condition" that reflects "a concern for others or a consideration of social cohesiveness and the common good" (Frankena, 1963, 9). For Turiel, being justified by appeal to social cohesiveness is part of the definition of a *conventional* judgment.
14. I know of no evidence that Turiel or any of his followers would construe their project in this way. It is offered here as a friendly amendment that avoids the challenge posed in the previous paragraph.
15. Why "many (*but perhaps not all*)"? Because commonsense intuition can't be counted on to be a flawless detector of natural kinds. Intuition told people that fool's gold was gold and that whales were fish. But when the relevant sciences discovered the essential features of gold and fish, it turned out that intuition was wrong about fool's gold and whales. For more on the way psychologists and other scientists might discover the essential features of a natural kind, see Stich (2018), §3.
16. I'll use this as shorthand for judgments that exhibit the UIH response pattern.
17. Philosophers include Dwyer (2006); Dwyer et al. (2010); Joyce (2006); Levy (2005); Nichols (2004a); Prinz (2007). Psychologists are too numerous to mention.
18. For a critique of Nissan (1987), see Turiel et al. (1988). For a critique of Kelly et al. (2007), see Sousa et al. (2009); for a response, see Stich et al. (2009). Kumar (2015) offers a rather different critique of Kelly et al. (2007). For a critique of Fessler et al. (2015), see Piazza and Sousa (2016); for a response, see Fessler et al. (2016).
19. This is a somewhat simplified version of the method employed in Goodwin and Darley (2008). Other investigators have used similar methods.
20. Isn't that speculation supported by the findings in the Turiel tradition? No, it's not. Turiel and his followers describe a behavior, but they never ask participants whether they think that behavior is "morally wrong."
21. In earlier papers, including Kelly et al. (2007), I have made the same mistake!
22. The authors of Chapter 2 of this volume adopt a strategy of this kind by shifting focus from the category of moral judgments to an analysis of normative judgments more generally.

## References

- Baier, K. (1958). *The Moral Point of View*. Ithaca, NY: Cornell University Press. Excerpts reprinted in G. Wallace and D. Walker (eds.), *The Definition of Morality*. London: Methuen, 1970, 188–210. Page references are to the Wallace and Walker reprint.
- Beebe, J. (2014). "How Different Kinds of Disagreement Impact Folk Metaethical Judgments," in H. Sarkissian and J. Wright (eds.), *Advances in Experimental Moral Psychology*. London: Bloomsbury.
- . (2015). "The Empirical Study of Folk Metaethics," *Etyka*, 15, 11–28.
- Beebe, J. and Sackris, D. (2016). "Moral Objectivism Across the Lifespan," *Philosophical Psychology*, 29, 6, 912–929.
- Blair, R. (1995). "A Cognitive Developmental Approach to Morality: Investigating the Psychopath," *Cognition*, 57, 1–29.
- Blair, R. (1996). "Brief Report: Morality in the Autistic Child," *Journal of Autism and Developmental Disorders*, 26, 571–579.
- Blair, R., Monson, J. and Frederickson, N. (2001). "Moral Reasoning and Conduct Problems in Children with Emotional and Behavioural Difficulties," *Personality and Individual Differences*, 31, 799–811.
- Boyd, R. (1988). "How to Be a Moral Realist," in G. Sayre-McCord (ed.), *Essays on Moral Realism*. Ithaca, NY: Cornell University Press.
- Buchtel, E., Guan, Y., Peng, Q., Su, Y., Sang, B. Chen, S. and Bond, M. (2015). "Immorality East and West: Are Prototypically Immoral Behaviors Especially Harmful, or Especially Uncultured?" To appear in *Personality and Social Psychology Bulletin*, 41, 10, 1382–1394.



- Cooper, N. (1966). "Two Concepts of Morality," *Philosophy*, 19–33. Reprinted in G. Wallace and D. Walker (eds.), *The Definition of Morality*. London: Methuen, 1970, 72–90. Page references are to the Wallace and Walker reprint.
- . (1970). "Morality and Importance," in G. Wallace and D. Walker (eds.), *The Definition of Morality*. London: Methuen, 91–97.
- Devitt, M. (1996). *Coming to Our Senses*. Cambridge: Cambridge University Press.
- Doris, J. (2015). *Talking to Our Selves: Reflection, Ignorance, and Agency*. Oxford: Oxford University Press.
- Dwyer, S. (2006). "How Good Is the Linguistic Analogy?" in P. Carruthers, S. Laurence and S. Stich (eds.), *The Innate Mind: Culture and Cognition*. New York: Oxford University Press.
- Dwyer, S., Huebner, B. and Hauser, M. (2010). "The Linguistic Analogy: Motivations, Results, and Speculations," *Topics in Cognitive Science*, 2, 486–510.
- Fessler, D., Barrett, H., Kanovsky, M., Stich, S., Holbrook, C., Henrich, J., Bolyanatz, A., Gervais, M., Gurven, M., Kushnick, G., Pisor, A., von Rueden, C. and Laurence, S. (2015). "Moral Parochialism and Contextual Contingency Across Seven Disparate Societies," *Proceedings of the Royal Society B*, 282 (1813) (August 22). doi:10.1098/rspb.2015.0907.
- . (2016). "Moral Parochialism Misunderstood: A Reply to Piazza and Sousa," *Proceedings of the Royal Society B*, 283. <http://dx.doi.org/10.1098/rspb.2015.2628>
- Fodor, J. (1981). "The Present Status of the Innateness Controversy," in J. Fodor (ed.), *Representations*. Cambridge, MA: Bradford Books.
- Frankena, W. (1958). "MacIntyre on Defining Morality," *Philosophy*, 158–162. Reprinted in G. Wallace and D. Walker (eds.), *The Definition of Morality*. London: Methuen, 1970, 40–46. Page references are to the Wallace and Walker reprint.
- . (1963). "Recent Conceptions of Morality," in H. Castañeda and G. Nakhnikian (eds.), *Morality and the Language of Conduct*. Detroit: Wayne State University Press, 1–24.
- . (1967). "The Concept of Morality," in University of Colorado Studies, *Series in Philosophy*, No. 3, 1–22. Reprinted in G. Wallace and D. Walker (eds.), *The Definition of Morality*. London: Methuen, 1970, 146–173. Page references are to the Wallace and Walker reprint.
- Gert, B. (2002). "The Definition of Morality," in Edward N. Zalta (ed.), *The Stanford Encyclopedia of Philosophy*. <http://plato.stanford.edu/archives/sum2002/entries/morality-definition/>
- Gewirth, A. (1978). *Reason and Morality*. Chicago: University of Chicago Press.
- Goldman, A. (2007). "Philosophical Intuitions: Their Target, Their Source, and Their Epistemic Status," *Grazer Philosophische Studien*, 74.
- Goodwin, G. and Darley, J. (2008). "The Psychology of Meta-Ethics: Exploring Objectivism," *Cognition*, 106, 1339–1366.
- Goodwin, G. and Darley, J. (2012). "Why Are Some Moral Beliefs Perceived to Be More Objective Than Others?" *Journal of Experimental Social Psychology*, 48, 1, 250–256.
- Haidt, J. and Joseph, C. (2007). "The Moral Mind: How 5 Sets of Innate Moral Intuitions Guide the Development of Many Culture-Specific Virtues, and Perhaps Even Modules," in P. Carruthers, S. Laurence and S. Stich (eds.), *The Innate Mind, Volume 3*. New York: Oxford University Press, 367–391.
- Haidt, J., Koller, S. and Dias, M. (1993). "Affect, Culture and Morality, or Is It Wrong to Eat Your Dog?" *Journal of Personality and Social Psychology*, 65, 613–628.
- Hare, R. (1952). *The Language of Morals*. Oxford: Clarendon Press.
- . (1954/1955). "Universalizability," *Proceedings of the Aristotelian Society*, 55, 295–312.
- Hollos, M., Leis, P. and Turiel, E. (1986). "Social Reasoning in IJO Children and Adolescents in Nigerian Communities," *Journal of Cross-Cultural Psychology*, 17, 352–376.
- Horgan, T. and Timmons, M. (1992). "Troubles for New Wave Moral Semantics: The "Open Question Argument" Revived," *Philosophical Papers*, 21, 153–175.
- Joyce, R. (2006). *The Evolution of Morality*. Cambridge, MA: MIT Press.
- Kelly, D., Stich, S., Haley, K., Eng, S. and Fessler, D. (2007). "Harm, Affect and the Moral/Conventional Distinction," *Mind and Language*, 22, 117–131.
- Kitcher, P. (2011). *The Ethical Project*. Cambridge, MA: MIT Press.



- Kornblith, H. (1998). "The Role of Intuition in Philosophical Inquiry: An Account with No Unnatural Ingredients," in M. DePaul and W. Ramsey (eds.), *Rethinking Intuition*. Lanham, MD: Rowman & Littlefield.
- Kumar, V. (2015). "Moral Judgment as a Natural Kind," *Philosophical Studies*, published Online February 5. doi:10.1007/s11098-015-0448-7.
- Ladd, J. (1957). *The Structure of a Moral Code: A Philosophical Analysis of Ethical Discourse Applied to the Ethics of the Navaho*. Cambridge, MA: Harvard University Press.
- Lapsley, D. (1996). *Moral Psychology*. Boulder, CO: Westview Press.
- Laurence, S. and Margolis, E. (1999). "Concepts and Cognitive Science," in E. Margolis and S. Laurence (eds.), *Concepts: Core Readings*. Cambridge, MA: MIT Press.
- Levine, S., Machery, E., Rottman, J., Davis, T. and Stich, S. (under review) "Religion's Impact on the Moral Sense"
- Levy, N. (2005). "Imaginative Resistance and the Moral/Conventional Distinction," *Philosophical Psychology*, 18, 231–241.
- Machery, E. (2009). *Doing Without Concepts*. Oxford: Oxford University Press.
- MacIntyre, A. (1957). "What Morality Is Not," *Philosophy*, 32, 325–335. Reprinted in G. Wallace and D. Walker (eds.), *The Definition of Morality*. London: Methuen, 1970, 26–39. Page references are to the Wallace and Walker reprint.
- Nichols, S. (2002). "Norms with Feeling: Towards a Psychological Account of Moral Judgment," *Cognition*, 84, 2, 221–236.
- Nichols, S. (2004a). *Sentimental Rules: On the Natural Foundations of Moral Judgment*. Oxford: Oxford University Press.
- . (2004b). "After Objectivity: An Empirical Study of Moral Judgment," *Philosophical Psychology*, 17, 3–26.
- Nissan, M. (1987). "Moral Norms and Social Conventions: A Cross-Cultural Comparison," *Developmental Psychology*, 23, 719–725.
- Nucci, L. (2001). *Education in the Moral Domain*. Cambridge: Cambridge University Press.
- Nucci, L. and Herman, S. (1982). "Behavioral Disordered Children's Conceptions of Moral, Conventional, and Personal Issues," *Journal of Abnormal Child Psychology*, 10, 411–425.
- Nucci, L. and Nucci, M. (1982). "Children's Social Interactions in the Context of Moral and Conventional Transgressions," *Child Development*, 53, 403–412.
- Nucci, L. and Turiel, E. (1978). "Social Interactions and the Development of Social Concepts in Preschool Children," *Child Development*, 49, 400–407.
- Nucci, L. and Turiel, E. (1993). "God's Word, Religious Rules, and Their Relation to Christian and Jewish Children's Concepts of Morality," *Child Development*, 64, 5, 1475–1491.
- Nucci, L., Turiel, E. and Encarnacion-Gawrych, G. (1983). "Children's Social Interactions and Social Concepts in the Virgin Islands," *Journal of Cross-Cultural Psychology*, 14, 469–487.
- Piazza, J. and Sousa, P. (2016). "When Injustice Is at Stake, Moral Judgments Are Not Parochial," *Proceedings of the Royal Society B*, 283, 20152037.
- Prinz, J. (2007). *The Emotional Construction of Morals*. Oxford: Oxford University Press.
- . (2015). "An Empirical Case for Motivational Internalism," in G. Bjornsson, C. Strandberg, R. Ollinder, J. Eriksson and F. Bjorklund (eds.), *Motivational Internalism*. Oxford: Oxford University Press.
- Putnam, H. (1975). "The Meaning of 'Meaning'," in K. Gunderson (ed.), *Language, Mind and Knowledge. Minnesota Studies in the Philosophy of Science*, vol. 7. Minneapolis: University of Minnesota Press.
- Rawls, J. (1971). *A Theory of Justice*. Cambridge, MA: Harvard University Press.
- Roskies, A. (2003). "Are Ethical Judgments Intrinsically Motivational? Lessons from "Acquired Sociopathy,'" *Philosophical Psychology*, 16, 51–66.
- Sarkissian, H., Park, J., Tien, D., Wright, J. C., Knobe, J. (2011). "Folk Moral Relativism," *Mind and Language*, 26, 4, 482–505.
- Schroeder, M. (2008). *Being for: Evaluating the Semantic Program of Expressivism*. Oxford: Oxford University Press.

- Smetana, J. (1981). "Preschool Children's Conceptions of Moral and Social Rules," *Child Development*, 52, 1333–1336.
- . (1993). "Understanding of Social Rules," in M. Bennett (ed.), *The Development of Social Cognition: The Child as Psychologist*. New York: Guilford Press.
- Smetana, J., Kelly, M. and Twentymen, C. (1984). "Abused, Neglected, and Nonmaltreated Children's Conceptions of Moral and Social-Conventional Transgressions," *Child Development*, 55, 277–287.
- Smetana, J., Toth, S., Cicchetti, D., Bruce, J., Kane, P. and Daddis, C. (1999). "Maltreated and Nonmaltreated Preschoolers' Conceptions of Hypothetical and Actual Moral Transgressions," *Developmental Psychology*, 35, 269–281.
- Smith, E. and Medin, D. (1981). *Categories and Concepts*. Cambridge, MA: Harvard University Press.
- Sousa, P., Holbrook, C. and Piazza, J. (2009). "The Morality of Harm," *Cognition*, 113, 80–92.
- Southwood, N. (2011). "The Moral/Conventional Distinction," *Mind*, 120, 761–802.
- Sprigge, T. (1964). "Definition of a Moral Judgment," *Philosophy*, 301–322. Reprinted in G. Wallace and D. Walker (eds.), *The Definition of Morality*. London: Methuen, 1970, 118–145. Page references are to the Wallace and Walker reprint.
- Stich, S. (2008). "Some Questions About the Evolution of Morality," *Philosophy and Phenomenological Research*, 77 (1), 228–236.
- Stich, S. (2018). "The Moral Domain," in K. Gray and J. Graham (eds.), *The Atlas of Moral Psychology*. New York: Guilford Press, 547–555.
- Stich, S., Fessler, D. and Kelly, D. (2009). "On the Morality of Harm: A Response to Sousa, Holbrook and Piazza," *Cognition*, 113, 93–97.
- Stohr, K. (2012). *On Manners*. New York: Routledge.
- Taylor, P. (1978). "On Taking the Moral Point of View," *Midwest Studies in Philosophy*, III, 35–61.
- Tisak, M. (1995). "Domains of Social Reasoning and Beyond," in R. Vasta (ed.), *Annals of Child Development*, vol. 11. London: Jessica Kingsley.
- Tisak, M. and Turiel, E. (1988). "Variations in Seriousness of Transgressions and Children's Moral and Conventional Concepts," *Developmental Psychology*, 24, 352–357.
- Toulmin, S. (1950). *An Examination of the Place of Reason in Ethics*. Cambridge: Cambridge University Press.
- Turiel, E. (1983). *The Development of Social Knowledge: Morality and Convention*. Cambridge: Cambridge University Press.
- Turiel, E., Nucci, L. and Smetana, J. (1988). "A Cross-Cultural Comparison About What? A Critique of Nisan's (1987). Study of Morality and Convention," *Developmental Psychology*, 24, 140–143.
- von Wright, G. (1963). *The Varieties of Goodness*. New York: The Humanities Press.
- Wallace, G. and Walker, D. (1970). "Introduction," in G. Wallace and D. Walker (eds.), *The Definition of Morality*. London: Methuen.
- Wright, J., Grandjean, P. and McWhite, C. (2013). "The Meta-ethical Grounding of Our Moral Beliefs: Evidence for Meta-Ethical Pluralism," *Philosophical Psychology*, 26, 336–361.
- Yau, J. and Smetana, J. (2003). "Conceptions of Moral, Social-Conventional, and Personal Events Among Chinese Preschoolers in Hong Kong," *Child Development*, 74 (1), 647–658.

### Further Readings

For an excellent overview of early work on the Philosophers' Project, see W. Frankena, "Recent Conceptions of Morality," in H. Castañeda and G. Nakhnikian (eds.), *Morality and the Language of Conduct* (Detroit: Wayne State University Press, 1963), 1–24. G. Wallace and D. Walker, *The Definition of Morality* (London: Methuen, 1970) is a collection of important papers debating the Philosophers' Project. For a definitive account of Turiel's version of the Psychologists' Project, see E. Turiel, *The Development of Social Knowledge: Morality and Convention* (Cambridge: Cambridge University Press, 1983). Two important empirical critiques of the Psychologists' Project are J. Haidt, S. Koller, and M. Dias, "Affect, Culture and Morality, or Is It Wrong to Eat Your Dog?" *Journal of Personality and Social Psychology*, 65, 613–628, 1993, and D. Kelly, S. Stich, K. Haley, S. Eng, and D. Fessler, "Harm, Affect and the Moral/Conventional Distinction," *Mind and Language*,

22, 117–131, 2017. V. Kumar, “Moral Judgment as a Natural Kind,” *Philosophical Studies*, published online February 5, 2015 (doi: 10.1007/s11098-015-0448-7) is a recent attempt to avoid the problems that beset the Psychologists’ Program. E. O’Neill, E. “Kinds of Norms,” *Philosophical Compass*, 12, e12416, 2017 (doi: 10.1111/phc3.12416) is a valuable discussion of the many different kinds of norms found in cultures around the world.

### Related Chapters

Chapter 2 The Normative Sense: What is Universal? What Varies; Chapter 5 Moral Development in Humans; Chapter 6 Moral Learning; Chapter 7 Moral Reasoning and Emotion; Chapter 8 Moral Intuitions and Heuristics; Chapter 9 The Evolution of Moral Cognition; Chapter 14 Nihilism and the Epistemic Profile of Moral Judgment; Chapter 15 Relativism and Pluralism in Moral Epistemology; Chapter 16 Rationalism and Intuitionism: Assessing Three Views about the Psychology of Moral Judgment; Chapter 20 Methods, Goals and Data in Moral Theorizing; Chapter 21 Moral Theory and its Role in Everyday Moral Thought and Action; Chapter 22 Moral Knowledge as Know-How; Chapter 25 Moral Expertise; Chapter 30 Religion and Moral Knowledge.

Taylor & Francis  
Not for distribution