

The Secret Joke of Kant's Soul

Joshua Greene

“Two things fill the mind with ever new and increasing wonder and awe, the oftener and more steadily we reflect on them: the starry heavens above me and the moral law within me.”

--Immanuel Kant (Kant, 1993)

“That such an unnatural use (and so misuse) of one's sexual attributes is a violation of one's duty to himself and is certainly in the highest degree opposed to morality strikes everyone upon his thinking of it... However, it is not so easy to produce a rational demonstration of the inadmissability of that unnatural use, and even the mere unpurposive use, of one's sexual attributes as being a violation of one's duty to himself (and indeed in the highest degree where the unnatural use is concerned). The ground of proof surely lies in the fact that a man gives up his personality (throws it away) when he uses himself merely as a means for the gratification of an animal drive.”

--Immanuel Kant, “Concerning Wanton Self-Abuse” (Kant, 1994) pg. 86.

“*Kant's Joke*—Kant wanted to prove, in a way that would dumbfound the common man, that the common man was right: that was the secret joke of this soul. He wrote against the scholars in support of popular prejudice, but for scholars and not for the people.”

--Friedrich Nietzsche (Nietzsche, 1974)

There is a substantial and growing body of evidence suggesting that much of what we do, we do unconsciously, and for reasons that are inaccessible to us (Wilson, 2002). In one experiment, for example, people were asked to choose one of several panty-hose displayed in a row. When asked to explain their

preferences, people gave sensible enough answers, referring to the relevant features of the items chosen—superior knit, sheerness, elasticity, etc. But their choices had nothing to do with such features, as the items on display were in fact identical. People simply had a preference for items on the right side of the display (Nisbett and Wilson, 1977). What this experiment illustrates—and there are many, many such illustrations—is that people (1) make choices for reasons unknown to them and (2) make up reasonable sounding justifications for their choices, all the while unaware of their actual motivations and subsequent rationalizations.

Jonathan Haidt applies these psychological lessons to the study of moral judgment in his influential paper, “The Emotional Dog and Its Rational Tail: A Social Intuitionist Approach to Moral Judgment” (Haidt, 2001). He argues that, for the most part, moral reasoning is a *post-hoc* affair: We decide what’s right or wrong on the basis of emotionally-driven intuitions, and then, if necessary, we make up reasons to explain and justify our judgments. Haidt concedes that some people, some of the time, may actually reason their way to moral conclusions, but he insists that this is not the norm. More importantly for the purposes of this essay, Haidt does not distinguish among the various approaches to ethics familiar to moral philosophers: consequentialism, deontology, virtue ethics, etc. Rather, his radical thesis is intended, if only implicitly, to apply equally to the adherents of all moral philosophies, though not necessarily well to moral philosophers as a group (Kuhn, 1991).

Jonathan Baron (Baron, 1994), in contrast, draws a psychological distinction between consequentialist and non-consequentialist judgments, arguing that the latter are especially likely to be made on the basis of heuristics, simple rules of thumb for decision-making. Baron, however, does not regard emotion as essential to these heuristic judgments.

In this essay, I draw on Haidt’s and Baron’s respective insights in the service of a bit of philosophical psychoanalysis. I will argue that deontological judgments tend to be driven by emotional responses, and that deontological philosophy, rather than being grounded in moral *reasoning*, is to a large extent

an exercise in moral *rationalization*. This is in contrast to consequentialism, which, I will argue, arises from rather different psychological processes, ones that are more “cognitive,” and more likely to involve genuine moral reasoning. These claims are strictly empirical, and I will defend them on the basis of the available evidence. Needless to say, my argument will be speculative and will not be conclusive. Beyond this, I will argue that if these empirical claims are true, they may have normative implications, casting doubt on deontology as a school of normative moral thought.

I. Preliminaries

Defining Deontology and Consequentialism

Deontology is defined by its emphasis on moral rules, most often articulated in terms of *rights* and *duties*. Consequentialism, in contrast, is the view that the moral value of an action is, in one way or another, a function of its consequences alone. Consequentialists maintain that moral decision-makers should always aim to produce the best overall consequences for all concerned, if not directly than indirectly. Both consequentialists and deontologists think that consequences are important, but consequentialists believe that consequences are the *only* things that ultimately matter, while deontologists believe that morality both requires and allows us to do things that do not produce the best possible consequences. For example, a deontologist might say that killing one person in order to save several others is wrong, even if doing so would maximize good consequences (Kagan, 1997).

The above is a standard explanation of what deontology and consequentialism are and how they differ. In light of this explanation, it might seem that my thesis is false *by definition*. Deontology is rule-based morality, usually focused on rights and duties. A deontological judgment, then, is a judgment made out of respect for certain types of moral rules. From this it

follows that a moral judgment that is made on the basis of an emotional response simply cannot be a deontological judgment, although it may appear to be from the outside. Kant himself was adamant about this, at least with respect to his own brand of deontology. He notoriously claimed that an action performed merely out of sympathy and not out of an appreciation of one's duty lacks moral worth (Kant, 1959, Ch. 1)(Korsgaard, 1996a, Ch. 2).

The assumption behind this objection—and as far as I know it has never been questioned previously—is that consequentialism and deontology are, first and foremost, moral philosophies. It is assumed that philosophers know exactly what deontology and consequentialism are because these terms/concepts were defined by philosophers. Despite this, I believe it is possible that philosophers do not necessarily know what consequentialism and deontology really are.

How could this be? The answer, I propose, is that the terms “deontology” and “consequentialism” refer to *psychological natural kinds*. I believe that consequentialist and deontological philosophy are not so much philosophical inventions as they are philosophical manifestations of two dissociable psychological patterns, two different ways of moral thinking, that have been part of the human repertoire for thousands of years. According to this view, the moral philosophies of Kant, Mill, and others are just the explicit tips of large, mostly implicit, psychological icebergs. And if that's right, then philosophers may not really know what they're dealing with when they trade in consequentialist and deontological moral theories, and we may have to do some science to find out.

An analogy, drawing on a familiar philosophical theme: Suppose that in a certain tropical land they refer to water by this symbol: ♠. And in their *Holy Dictionary* it clearly states that ♠ is a clear and drinkable liquid. (That is, the dictionary defines ♠ in terms of its “primary intension” (Chalmers, 1996).) One day, an enterprising youngster journeys to the top of a nearby mountain and is the first of her people to encounter ice. Through a bit of experimentation, she discovers that ice is a form of water and excitedly tells the tribal Elders of her discovery. The next day she drags one of the Elders to the mountaintop, hands him some ice, and says, “Behold! ♠!” At which point the exasperated Elder

explains that ♠ is a liquid, that the hard stuff in his hand is clearly not a liquid, and that he doesn't appreciate having his time wasted.

In a narrow sense the Elder is correct. The *Holy Dictionary* is the authority on what the local symbols mean, and it states very clearly that ♠ refers to a clear, drinkable, liquid. But the Elder is missing the big picture. What he is forgetting, or perhaps never understood, is that many things in the world have underlying structures—"essences," if you prefer—that are responsible for making things appear and behave as they do, for giving them their functional properties. And because things have underlying structures, it is possible to refer to something, even make up a definition for it, without really understanding what it is (Kripke, 1980; Putnam, 1975). Of course, a linguistic community can insist that their definition is correct. No one's to stop them from using their symbols as they please. But, in doing this, they run the risk of missing the big picture, of denying themselves a deeper understanding of what's going on around them, or even within them.

Because I am interested in exploring the possibility that deontology and consequentialism are psychological natural kinds, I will put aside their conventional philosophical definitions and focus instead on their relevant functional roles. As noted above, consequentialists and deontologists have some characteristic practical disagreements. For example, consequentialists typically say that killing one person in order to save several others may be the right thing to do, depending on the situation. Deontologists, in contrast, typically say that it's wrong to kill one person for the benefit of others, that the "ends don't justify the means." Because consequentialists and deontologists have these sorts of practical disagreements, we can use these disagreements to define consequentialist and deontological judgments functionally. For the purposes of this discussion, we'll say that consequentialist judgments are judgments in favor of characteristically consequentialist conclusions (e.g. "Better to save more lives") and that deontological judgments are judgments in favor of a characteristically deontological conclusions (e.g. "It's wrong despite the benefits"). My use of "characteristically" is obviously loose here, but I trust that

those familiar with contemporary ethical debates will know what I mean. Note that kind of judgment made is largely independent of who is making it. A card-carrying deontologist can make a “characteristically consequentialist” judgment, as when Judith Jarvis Thomson says that it’s okay to turn a runaway trolley that threatens to kill five people onto a side track so that it will kill only one person instead (Thomson, 1986). This is a “characteristically consequentialist” judgment because it is easily justified in terms of the most basic consequentialist principles, while deontologists need to do a lot of fancy philosophizing in order to defend this position. Likewise, consider the judgment that it’s wrong to save five people who need organ transplants by swiping the organs out of an unwilling donor (Thompson, 1986). This judgment is “characteristically deontological,” not because many card-carrying consequentialists don’t agree, but because they have to do a lot of extra explaining in order to justify their agreement.

By defining “consequentialism” and “deontology” in terms of their characteristic judgments, we give our empirical hypothesis a chance. If it turns out that characteristically deontological judgments are driven by emotion (an empirical possibility), then that raises the possibility that deontological *philosophy* is also driven by emotion (a further empirical possibility). In other words, what we find when we explore the psychological causes of characteristically deontological judgments might suggest that what deontological moral philosophy really is, what it is *essentially*, is an attempt to produce rational justifications for emotionally-driven moral judgments, and not an attempt to reach moral conclusions on the basis of moral reasoning.

The point for now, however, is simply to flag the terminological issue. When I refer to something as a “deontological judgment” I am saying that it is a characteristically deontological judgment and am not insisting that the judgment in question necessarily meets the criteria that philosophers would impose for counting that judgment as deontological. In the end, however, I will argue that such judgments are best understood as genuinely deontological because they are produced by an underlying psychology that is the hidden essence of deontological philosophy.

Defining “Cognition” and Emotion

In what follows I will argue that deontological judgment tends to be driven by emotion, while consequentialist judgment tends to be driven by “cognitive” processes. What do we mean by “emotion” and “cognition,” and how do these things differ?

Sometimes “cognition” refers to information processing in general, as in “cognitive science,” but often “cognition” is used in a narrower sense that contrasts with “emotion,” despite the fact that emotions involve information processing. I know of no good off-the-shelf definition of “cognition” in this more restrictive sense, despite its widespread use. Elsewhere, my collaborators and I offered a tentative definition of our own (Greene et al., 2004), one that is based on the differences between the information-processing requirements of stereotyped vs. flexible behavior.

The rough idea is that “cognitive” representations are inherently neutral representations, ones that do not automatically trigger particular behavioral responses or dispositions, while “emotional” representations are ones that do have such automatic effects, that are behaviorally valenced. Highly flexible behavior requires “cognitive” representations that can be easily mixed around and recombined as situational demands vary, and without pulling the agent in sixteen different behavioral directions at once. For example, sometimes you need to avoid cars, and other times you need to approach them. It is useful, then, if you can represent CAR in a behaviorally neutral or “cognitive” way, one that doesn’t automatically presuppose a particular behavioral response. Stereotyped behavior, in contrast, doesn’t require this sort of flexibility and therefore doesn’t require “cognitive” representations, at least not to the same extent. To make things clear, I will use quotation marks to indicate the more restrictive sense of “cognitive” defined here, and I will drop the quotation marks when using this term to refer to information-processing in general.

While the whole brain is devoted to cognition, “cognitive” processes are especially important for reasoning, planning, manipulating information in working memory, controlling impulses, and “higher executive functions” more generally. Moreover, these functions tend to be associated with certain parts of the brain, primarily the dorsolateral surfaces of the prefrontal cortex and parietal lobes (Koechlin et al., 2003; Miller and Cohen, 2001; Ramnani and Owen, 2004). Emotion, in contrast, tends to be associated with other parts of the brain such as the amygdala and the medial surfaces of the frontal and parietal lobes (Adolphs, 2002; Maddock, 1999; Phan et al., 2002). And while the term “emotion” can refer to stable states such as moods, here we will primarily be concerned with emotions subserved by processes that, in addition to being valenced, are quick and automatic, though not necessarily conscious.

Here we are concerned with two different kinds of moral judgment (deontological and consequentialist) and two different kinds of psychological process (“cognitive” and emotional). Crossing these, we get four basic empirical possibilities. First, it could be that both kinds of moral judgment are generally “cognitive,” as Kohlberg’s theories suggest (Kohlberg, 1971).¹ At the other extreme, it could be that both kinds of moral judgment are primarily emotional, as Haidt’s view suggests (Haidt, 2001). Then there is the historical stereotype, according to which consequentialism is more emotional (emerging from the “sentimentalist” tradition of David Hume (Hume, 1978) and Adam Smith (Smith, 1976)) while deontology is more “cognitive” (encompassing the Kantian “rationalist” tradition (Kant, 1959)). Finally, there is the view for which I will argue, that deontology is more emotionally-driven while consequentialism is more “cognitive.” I hasten to add, however, that I don’t believe that either approach is strictly emotional or “cognitive” (or even that there is a sharp distinction between “cognition” and emotion). More specifically, I am sympathetic to Hume’s claim that all moral judgment (including consequentialist judgment) must have some emotional component (Hume, 1978). But I suspect that the kind

¹ Kohlberg was certainly partial to deontology and would likely say that it is more “cognitive” than consequentialism.

of emotion that is essential to consequentialism is fundamentally different from the kind that is essential to deontology, the former functioning more like a currency and the latter functioning more like an alarm. We will return to this issue later.

II. Scientific Evidence

Evidence from Neuroimaging

In recent decades, philosophers have devised a range of hypothetical moral dilemmas that capture the tension between the consequentialist and deontological viewpoints. A well-known handful of these dilemmas gives rise to what is known as the “Trolley Problem” (Foot, 1978; Thomson, 1986), which begins with the *trolley* dilemma:

A runaway trolley is headed for five people who will be killed if it proceeds on its present course. The only way to save these people is to hit a switch that will turn the trolley onto a side track where it will run over and kill one person instead of five. Is it okay to turn the trolley in order to save five people at the expense of one? The consensus among philosophers (Fischer and Ravizza, 1992), as well as people who have been tested experimentally (Petrinovich and O'Neill, 1996; Petrinovich et al., 1993), is that it is morally acceptable to save five lives at the expense of one in this case.

Next consider the *footbridge* dilemma (Thomson, 1986): As before, a runaway trolley threatens to kill five people, but this time you are standing next to a large stranger on a footbridge spanning the tracks, in between the oncoming trolley and the five people. The only way to save the five people is to push this stranger off the bridge and onto the tracks below. He will die as a result, but his body will stop the trolley from reaching the others. Is it okay to save the five people by pushing this stranger to his death? Here the consensus is that it's not okay to save five lives at the expense of one (Fischer and Ravizza, 1992; Greene

et al., 2004; Greene et al., 2001; Petrinovich and O'Neill, 1996; Petrinovich et al., 1993).

People exhibit a characteristically consequentialist response to the *trolley* case and a characteristically deontological response to the *footbridge* case. Why? Philosophers have generally offered a variety of *normative* explanations. That is, they have assumed that our responses to these cases are correct, or at least reasonable, and have sought principles that *justify* treating these two cases differently (Fischer and Ravizza, 1992). For example, one might suppose, following Kant (Kant, 1959) and Aquinas (Aquinas, 1988), that it's wrong to harm someone as a means to helping someone else. In the *footbridge* case the proposed action involves literally using the person on the footbridge as a trolley-stopper, whereas in the *trolley* case the victim is to be harmed merely as a side-effect. (Were the single person on the alternate track to magically disappear, we would be very pleased.) In response to this proposal, Thomson devised the *loop* case (Thomson, 1986). Here, the situation is similar to that of the *trolley* dilemma, but this time the single person is on a piece of track that branches off of the main track and then rejoins it at a point before the five people. In this case, if the person were not on the side track, the trolley would return to the main track and run over the five people. The consensus here is that it is morally acceptable to turn the trolley in this case, despite the fact that here, as in the *footbridge* case, a person will be used as a means.

There have been many such normative attempts to solve the trolley problem, but none of them has been terribly successful (Fischer and Ravizza, 1992). My collaborators and I have proposed a partial and purely descriptive solution to this problem and have collected some scientific evidence in favor of it. We hypothesized that the thought of pushing someone to his death in an "up close and personal" manner (as in the *footbridge* dilemma) is more emotionally salient than the thought of bringing about similar consequences in a more impersonal way (e.g. by hitting a switch, as in the *trolley* dilemma). We proposed that this difference in emotional response explains why people respond so differently to these two cases. That is, people tend toward consequentialism in

the case in which the emotional response is low and tend toward deontology in the case in which the emotional response is high.

The rationale for distinguishing between *personal* and *impersonal* forms of harm is largely evolutionary. “Up close and personal” violence has been around for a very long time, reaching back far into our primate lineage (Wrangham and Peterson, 1996). Given that personal violence is evolutionarily ancient, predating our recently-evolved human capacities for complex abstract reasoning, it should come as no surprise if we have innate responses to personal violence that are powerful, but rather primitive. That is, we might expect humans to have negative emotional responses to certain basic forms of interpersonal violence, where these responses evolved as a means of regulating the behavior of creatures who are capable of intentionally harming one another, but whose survival depends on cooperation and individual restraint (Sober and Wilson, 1998; Trivers, 1971). In contrast, when a harm is *impersonal* it should fail to trigger this alarm-like emotional response, allowing people to respond in a more “cognitive” way, perhaps employing a cost-benefit analysis. As Josef Stalin once said, “A single death is a tragedy; a million deaths is a statistic.” His remarks suggest that when harmful actions are sufficiently impersonal, they fail to push our emotional buttons, despite their seriousness, and as a result we think about them in a more detached, actuarial fashion.

This hypothesis makes some strong predictions regarding what we should see going on in people’s brains while they’re responding to dilemmas involving personal vs. impersonal harm (henceforth “personal” and “impersonal” moral dilemmas). The contemplation of personal moral dilemmas like the *footbridge* case should produce increased neural activity in brain regions associated with emotional response and social cognition, while the contemplation of impersonal moral dilemmas like the *trolley* case should produce relatively greater activity in brain regions associated with “higher cognition.”² This is exactly what was

² It turns out that determining what makes a moral dilemma “personal” and “like the footbridge case” vs. “impersonal” and “like the trolley case” is no simple matter, and in many ways reintroduces the complexities associated with traditional attempts to solve the trolley problem. For the purposes of this essay, however, I’m happy to leave the personal/impersonal distinction

observed (Greene et al., 2004; Greene et al., 2001). Contemplation of personal moral dilemmas produced relatively greater activity in three emotion-related areas, the posterior cingulate cortex, the medial prefrontal cortex, and the amygdala. This effect was also observed in the superior temporal sulcus, a region associated with various kinds of social cognition in humans and other primates (Allison et al., 2000; Saxe et al., 2004). At the same time, contemplation of impersonal moral dilemmas produced relatively greater neural activity in two classically "cognitive" brain areas, the dorsolateral prefrontal cortex and inferior parietal lobe.

This hypothesis also makes a prediction regarding people's reaction times. According to the view I've sketched, people tend to have emotional responses to personal moral violations, responses that incline them to judge against performing those actions. That means that someone who judges a personal moral violation to be *appropriate* (e.g., someone who says it's okay to push the man off the bridge in the *footbridge* case) will most likely have to override an emotional response in order to do it. This overriding process will take time, and thus we would expect that "yes" answers will take longer than "no" answers in response to personal moral dilemmas like the *footbridge* case. At the same time, we have no reason to predict a difference in reaction time between "yes" and "no" answers in response to impersonal moral dilemmas like the *trolley* case because there is, according to this model, no emotional response (or much less of one) to override in such cases. Here, too, the prediction has held. Trials in which the subject judged in favor of personal moral violations took significantly longer than trials in which the subject judged against them, but there was no comparable reaction time effect observed in response to impersonal moral violations (Greene et al., 2004; Greene et al., 2001).

Further results support this model as well. Next we subdivided the personal moral dilemmas into two categories on the basis of difficulty (i.e. based

as an intuitive one, in keeping with the evolutionary account given above. For the purposes of designing the brain-imaging experiment discussed below, however, my collaborators and I developed a more rigid set of criteria for distinguishing personal from impersonal moral violations (Greene *et al.*, 2001). I no longer believe that these criteria are adequate. Improving these is a goal of ongoing research.

on reaction time). Consider the following moral dilemma (the *crying baby* dilemma): It's wartime, and you and some of your fellow villagers are hiding from enemy soldiers in a basement. Your baby starts to cry, and you cover your baby's mouth to block the sound. If you remove your hand, your baby will cry loudly, the soldiers will hear, and they will find you and the others and kill everyone they find, including you and your baby. If you do not remove your hand, your baby will smother to death. Is it okay to smother your baby to death in order to save yourself and the other villagers?

This is a very difficult question. Different people give different answers, and nearly everyone takes a relatively long time. This is in contrast to other personal moral dilemmas, such as the *infanticide* dilemma, in which a teenage girl must decide whether to kill her unwanted newborn. In response to this case, people (at least the ones we tested) quickly and unanimously say that this action is wrong.

What's going on in these two cases? My colleagues and I hypothesized as follows. In both cases there is a prepotent, negative emotional response to the personal violation in question, killing one's own baby. In the *crying baby* case, however, a cost-benefit analysis strongly favors smothering the baby. After all, the baby is going to die no matter what, and so you have nothing to lose (in consequentialist terms) and much to gain by smothering it, awful as it is. In some people the emotional response dominates, and those people say "no." In other people, this "cognitive," cost-benefit analysis wins out, and these people say "yes."

What does this model predict that we'll see going on in people's brains when we compare cases like *crying baby* and *infanticide*? First, this model supposes that cases like *crying baby* involve an increased level of "response conflict," that is, conflict between competing representations for behavioral response. Thus, we should expect that difficult moral dilemmas like *crying baby* will produce increased activity in a brain region that is associated with response conflict, the anterior cingulate cortex (Botvinick et al., 2001). Second, according to our model, the crucial difference between cases like *crying baby* and cases

like *infanticide* is that the former evoke strong "cognitive" responses that can effectively compete with a prepotent, emotional response. Thus, we should expect to see increased activity in classically "cognitive" brain areas when we compare cases like *crying baby* to cases like *infanticide*, despite the fact that difficult dilemmas like *crying baby* are personal moral dilemmas, which were previously associated with emotional response (Greene et al., 2001).

These two predictions have held (Greene et al., 2004). Comparing high-reaction-time personal moral dilemmas like *crying baby* to low-reaction-time personal moral dilemmas like *infanticide* revealed increased activity in the anterior cingulate cortex (conflict) as well as the anterior dorsolateral prefrontal cortex and the inferior parietal lobes, both classically "cognitive" brain regions.

Cases like *crying baby* are especially interesting because they allow us to directly compare the neural activity associated with characteristically consequentialist and deontological responses. According to our model, when people say "yes" to such cases (the consequentialist answer) it is because the "cognitive" cost-benefit analysis has successfully dominated the prepotent emotional response that drives people to say "no" (the deontological answer). If that's right, then, we should expect to see increased activity in the previously identified "cognitive" brain regions (dorsolateral prefrontal cortex and inferior parietal cortex) for the trials in which people say "yes" in response to cases like *crying baby*. This is exactly what we found. In other words, people exhibit more "cognitive" activity when they give the consequentialist answer.³

To summarize, people's moral judgments appear to be products of at least two different kinds of psychological processes. First, both brain imaging and reaction time data suggest that there are prepotent negative emotional responses that drive people to disapprove of the personally harmful actions proposed in cases like the *footbridge* and *crying baby* dilemmas. These

³ It is worth noting that no brain regions, including those implicated in emotion, exhibited the opposite effect. First, it's not clear that one would expect to see such a result since the hypothesis is that everyone experiences the intuitive emotional response, while only some individuals override it. Second, it is difficult to draw conclusions from negative neuroimaging results because current neuroimaging techniques, which track changes in blood flow, are relatively crude instruments for detecting patterns in neural function.

responses are characteristic of deontology, but not of consequentialism. Second, further brain imaging results suggest that “cognitive” psychological processes can compete with the aforementioned emotional processes, driving people to approve of personally harmful moral violations, primarily when there is a strong consequentialist rationale for doing so, as in the *crying baby* case. The parts of the brain that exhibit increased activity when people make characteristically consequentialist judgments are those that are most closely associated with higher cognitive functions such as executive control (Koechlin et al., 2003; Miller and Cohen, 2001), complex planning (Koechlin et al., 1999), deductive and inductive reasoning (Goel and Dolan, 2004), taking the long view in economic decision-making (McClure et al., 2004), and so on. Moreover, these brain regions are among those most dramatically expanded in humans as compared to other primates (Allman et al., 2002).

Emotion and the Sense of Moral Obligation

In his classic article, “Famine, Affluence, and Morality,” Peter Singer (Singer, 1972) argues that we in the affluent world have an obligation to do much more than we do to improve the lives of needy people. He argues that if we can prevent something very bad from happening without incurring a comparable moral cost, then we ought to do it. For example, if one notices a small child drowning in a shallow pond, one is morally obliged to wade in and save that child, even if it means muddying up one’s clothes. As Singer points out, this seemingly innocuous principle has radical implications, implying that all of us who spend money on unnecessary luxuries should give up those luxuries in order to spend the money on saving and/or improving the lives of impoverished peoples. Why, Singer asks, do we have a strict obligation to save a nearby drowning child, but no comparable obligation to save faraway sick and starving children through charitable donations to organizations like Oxfam?

Many normative explanations come to mind, but none is terribly compelling. Are we allowed to ignore the plight of faraway children because they are citizens of foreign nations? If yes, then would it be acceptable to let the child drown, provided that the child was encountered while traveling abroad? Or in international waters? And what about the domestic poor? This argument doesn't relieve us of our obligations to them. Is it because of diffused responsibility—because many are in a position to help a starving child abroad, but only you are in a position to help this hypothetical drowning child? What if there were many people standing around the pond doing nothing? Would that make it okay for you to do nothing as well? Is it because international aid is ultimately ineffective, only serving to enrich corrupt politicians, or create more poor people? In that case, our obligation would simply shift to more sophisticated relief efforts incorporating political reform, economic development, family planning education, and so on. Are all relief efforts doomed to ineffectiveness? That's a bold empirical claim that no one can honestly make with great confidence.

Here, we find ourselves in a similar position to the one we faced with the trolley problem. We have a strong intuition that two moral dilemmas are importantly different, and yet we have a hard time explaining what that important difference is (Kagan, 1989; Unger, 1996). It turns out that the same psychological theory that makes sense of the trolley problem can make sense of Singer's problem. Note that the interaction in the case of the drowning child is "up close and personal," the sort of situation that might have been encountered by our human and primate ancestors. Likewise, note that the donation case is not "up close and personal," and is not the sort of situation that our ancestors could have encountered. At no point were our ancestors able to save the lives of anonymous strangers through modest material sacrifices. In light of this, the psychological theory presented above suggests that we are likely to find the obligation to save the drowning child more pressing simply because that "up close and personal" case pushes our emotional buttons in a way that the more impersonal donation case does not (Greene, 2003). As it happens, these two

cases were among those tested in the brain imaging study described above, with a variation on the drowning child case included in the *personal* condition and the donation case included in the *impersonal* condition (Greene et al., 2004; Greene et al., 2001).

Few people accept Singer's consequentialist conclusion. Rather, people tend to believe, in a characteristically deontological way, that they are within their moral rights to spend their money on luxuries for themselves, despite the fact that their money could be used to dramatically improve the lives of other people. This is exactly what one would expect if (1) the deontological sense of obligation is driven primarily by emotion, and (2) when it comes to obligations to aid, emotions are only sufficiently engaged when those to whom we might owe something are encountered (or conceived of) in a personal way.

Emotion and the Pull of Identifiable Victims

One aspect of someone's being "up close and personal" is that such a person is always, in some sense, an identifiable, determinate individual and not a mere statistical someone (Greene and Haidt, 2002; Greene et al., 2001). The drowning child, for example, is presented as a particular person, while the children you might help through donations to Oxfam are anonymous and, so far as you know, indeterminate.⁴ Many have observed a tendency to respond with greater urgency to identifiable victims, as compared to indeterminate, "statistical" victims (Schelling, 1968). This is known as the "identifiable victim effect."

You may recall, for example, the case of Jessica McClure, a.k.a. "Baby Jessica," who in 1987 was trapped in a well in Texas. Over \$700,000 was sent to her family in support of the rescue effort (Small and Loewenstein, 2003; Variety, 1989). As Small and Loewenstein point out, that amount of money, if spent on preventative healthcare, could have been used to save the lives of

⁴ Of course, some aid organizations deliberately pair individual donors with individual recipients in order to make the experience more personal.

many children. This observation raises a normative question, that is essentially the same as Singer's. Do we have a greater obligation to help people like Baby Jessica than we do to help large numbers of others who could be saved for less? If all else is equal, a consequentialist would say "no," while, most people apparently say "yes." Furthermore, most people, if pressed to explain their position, would probably do so in deontological terms. That is, they would probably say that we have a *duty* to aid someone like Baby Jessica, even if doing so involves great effort and expense, while we have no comparable duty to the countless others who might be helped using the same resources.

The same "up close and personal" theory of emotional engagement can explain this pattern of judgment. Others have proposed what amounts to the same hypothesis, and others still have gathered independent evidence in support of it. In Thomas Schelling's seminal article on this topic he observes that the death of a particular person invokes "anxiety and sentiment, guilt and awe, responsibility and religion, [but]... most of this awesomeness disappears when we deal with statistical death," (Schelling, 1968; Small and Loewenstein, 2003). Inspired by Schelling's observation, Small and Loewenstein conducted two experiments aimed at testing the hypothesis that "identifiable victims stimulate a more powerful emotional response than do statistical victims."

Their crucial move was to design their experiments in such a way that their results could count against all normative explanations of the identifiable victim effect, i.e. explanations that credit decision-makers with normatively respectable reasons for favoring identifiable victims. This is difficult because the process of identifying a victim inevitably provides information about that victim (name, age, gender, appearance, etc.) that could potentially serve as a rational basis for favoring that person. To avoid this, they sought to document a weaker form of the identifiable victim effect, which one might call the "determinate victim effect." They examined people's willingness to benefit determined vs. undetermined individuals under conditions in which all meaningful information about the victims is held constant.

Their first experiment worked as follows. Ten laboratory subjects were each given an “endowment” of ten dollars. Some subjects randomly drew cards that said “KEEP” and were allowed to retain their endowments, while other subjects drew cards that said “LOSE” and subsequently had their endowments taken away, thus rendering them “victims.” Each of the non-victim subjects was anonymously paired with one of the victims as a result of drawing that victim’s number. The non-victim subjects were allowed to give a portion of their endowments to their respective victims, and each could choose how much to give. However—the crucial manipulation—some non-victim subjects drew the victim’s number *before* deciding how much to give, while others drew the victim’s number *after* deciding, knowing in advance that they would do so later. In other words, some subjects had to answer the question, “How much do I want to give to person #4?” (determined victim), whereas other subjects had to answer the question, “How much do I want to give to the person whose number I will draw?” (undetermined victim). At no point did the non-victim subjects ever know who would receive their money. The results: The mean donation for the group who gave to determined victims was sixty percent higher than that of the group giving to undetermined victims. The median donation for the determined victim group was over twice as high.

It’s worth emphasizing the absurdity of this pattern of behavior. There is no rational basis for giving more money to “randomly determined person #4” than to “person #?, to be randomly determined,” and yet that’s what these people did.⁵ (Note that the experiment is designed so that none of the participants will ever know who chose what.) Why would people do this? Here, too, the answer

⁵ First, when I say that this behavior cannot be rationally defended, I do not mean that it is logically or metaphysically impossible for a rational person to behave this way. Someone could, for example, have a basic preference for helping determined victims and only determined victims. I am assuming, however, that none of the subjects in this experiment have such bizarre preferences, and that therefore their behavior is irrational. Second, I am not claiming that the general psychological tendency that produces this behavior has no “rationale” or that it is not adaptive. Rather, I am simply claiming that this particular behavior is irrational in this case. Few, if any, of the participants in this study would knowingly choose to respond to the experimental manipulation (determined vs. undetermined victim) by giving more to the determined victim. In other words, this experimental effect would have been greatly diminished, if not completely eliminated, had this experiment employed a within-subject design instead of a between-subject design.

implicates emotion. In a follow-up study replicating this effect, subjects reported on the levels of sympathy and pity they felt for the determined/undetermined victims with whom they were paired. As expected, their reported levels of sympathy and pity tracked their donation levels (personal communication).

One might wonder whether this pattern holds up outside the lab. To find out, Small and Loewenstein conducted a subsequent study in which people could donate money to Habitat for Humanity in order to provide a home for a needy family, where the family was either determined or to-be-determined. As predicted, the mean donation was twenty-five percent higher in the determined family condition, and the median donation in the determined family condition was double that of the undetermined family condition.

And then there's Baby Jessica. We can't say for sure that resources were directed to her instead of to causes that could use the money more effectively because of people's emotional responses (and not because of people's deontological reasoning about rights and duties), but what evidence there is suggests that that's the case. As Josef Stalin might have said, "A determinate individual's death is a tragedy; a million indeterminate deaths is a statistic."

Anger and Deontological Approaches to Punishment

While consequentialists and deontologists agree that punishment of wrongdoing is necessary and important, they disagree sharply over the proper justification for punishment. Consequentialists such as Jeremy Bentham (Bentham, 1982) argue that punishment is justified solely by its future beneficial effects, primarily through deterrence and (in the case of criminal law) the containment of dangerous individuals. While few would deny that the prevention of future harm provides a legitimate justification for punishment, many believe that such pragmatic considerations are not the *only* legitimate reasons to punish, or even the main ones. Deontologists such as Kant (Kant, 2002), for example, argue that the primary justification for punishment is *retribution*, to give wrong-doers what

they deserve based on what they've done, regardless of whether such retribution will prevent future wrong-doing.

One might wonder, then, about the psychology of the typical punisher. Do people punish, or endorse punishment, because of its beneficial effects, or do people punish because they are motivated to give people what they deserve, in proportion to their "internal wickedness," to use Kant's phrase (Carlsmith et al., 2002; Kant, 2002). Several studies speak to this question, and the results are consistent. People endorse both consequentialist and retributivist justifications for punishment in the abstract, but in practice, or when faced with more concrete hypothetical choices, people's motives appear to predominantly retributivist. Moreover, these retributivist inclinations appear to be emotionally-driven: People punish in proportion to the extent that transgressions make them angry.

First, let us consider whether punitive judgments are predominantly consequentialist or deontological/retributivist.⁶ Jonathan Baron and colleagues have conducted a series of experiments demonstrating that people's punitive judgments are, for the most part, retributivist rather than consequentialist. In one study Baron and Ritov (Baron and Ritov, 1993) presented people with hypothetical corporate liability cases in which corporations could be required to pay fines. In one set of cases a corporation that manufactures vaccines is being sued because a child died as a result of taking one of its flu vaccines. Subjects were given multiple versions of this case. In one version, it was stipulated that a fine would have a positive deterrent effect. That is, a fine would make the company produce a safer vaccine instead. In a different version, it was stipulated that a fine would have a "perverse" effect. Instead of causing the firm to make a safer vaccine available, a fine would cause the company to stop making this kind of vaccine altogether, a bad result given that the vaccine in question does more good than harm and that no other firm is capable of making

⁶ I am assuming that, within the domain of punishment, "deontological" and "retributivist" are effectively interchangeable, even though they are conceptually distinct. (For example, one could favor punishment as an end in itself, but in unpredictable ways that defy all normative rules.) So far as I know, all well-developed alternatives to consequentialist theories of punishment are, in one way or another, retributivist. Moreover, retributivism is explicitly endorsed by many noteworthy deontologists, including Kant (2002).

such a vaccine. Subjects indicated whether they thought a punitive fine was appropriate in either of these cases and whether the fine should differ between these two cases. A majority of subjects said that the fine should not differ at all. Baron and Ritov achieved similar results using a complementary manipulation concerning deterrent effects on the decisions of other firms. In a different set of studies Baron and colleagues found a similar indifference to consequentialist factors in response to questions about the management of hazardous waste (Baron et al., 1993).

The results of these studies are surprising in light of the fact that many people regard the deterrence of future harmful decisions as a major reason, if not the primary reason, for imposing such fines in the real world. The strength of these results is also worth emphasizing. The finding here is not simply that people's punitive judgments fail to accord with consequentialism, the view that consequences are ultimately the *only* things that should matter to decision-makers. Much more than that, it seems that a majority of people give *no weight whatsoever* to factors that are of clear consequentialist importance, at least in the contexts under consideration.

If people don't punish for consequentialist reasons, what motivates them? In a study by Kahneman and colleagues (Kahneman et al., 1998), subjects responded to a number of similar hypothetical scenarios (e.g. a case of anemia due to benzene exposure at work). Subjects rated for each scenario the extent to which the defendant's action is "outrageous." They also rated the extent to which the defendant in each case should be punished. The correlation between these scenarios' mean outrage ratings and their mean punishment ratings were near perfect, with a Pearson's correlation coefficient (r) of .98. (A value of 1 indicates a perfect correlation.) Kahneman and colleagues conclude that the extent to which people desire to see a corporation punished for its behavior is almost entirely a function of the extent to which people are emotionally outraged by that corporation's behavior.

Carlsmith and colleagues (Carlsmith et al., 2002) conducted a similar set of studies aimed explicitly at determining whether people punish for

consequentialist or deontological reasons. Here, as above, subjects were presented with scenarios involving morally and legally culpable behavior, in this case perpetrated by individuals rather than corporations. As before, subjects were asked to indicate how severe each person's punishment should be, first in abstract terms ("not at all severe" to "extremely severe") and then in more concrete terms ("not guilty"/no punishment to "life sentence"). The experimenters varied the scenarios in ways that warrant different levels of punishment, depending on the rationale for punishment. For example, a consequentialist theory of punishment considers the detection rate associated with a given kind of crime and the publicity associated with a given kind of conviction to be relevant factors in assigning punishments. According to consequentialists, if a crime is difficult to detect, then the punishment for that crime ought to be made more severe in order to counterbalance the temptation created by the low risk of getting caught. Likewise, if a conviction is likely to get a lot of publicity, then a law enforcement system interested in deterrence should take advantage of this circumstance by "making an example" of the convict with a particularly severe punishment, thus getting a maximum of deterrence "bang" for its punishment "buck."

The results were clear. For the experimental group as a whole, there was no significant change in punishment recommendations when the detection rates and levels of publicity were manipulated. In other words, people were generally indifferent to factors that, according to consequentialists, should matter, at least to some extent. This is in spite of the fact that Carlsmith *et al.*, as well as others (Weiner *et al.*, 1997), found that subjects readily expressed a general kind of support for deterrence-oriented penal systems and corporate policies.

In a follow-up study, subjects were explicitly instructed to adopt a consequentialist approach, with the consequentialist rationale explicitly laid out and with extra manipulation checks included to ensure that the subjects understood the relevant facts. Here, too, the results were striking. Subjects did modify their judgments when told to think like consequentialists, but not in a genuinely consequentialist way. Instead of becoming selectively sensitive to the

factors that increase the consequentialist benefits of punishment, subjects indiscriminately ratcheted up the level of punishment in all cases, giving perpetrators the punishment that they thought the perpetrators deserved based on their actions, plus a bit more for the sake of deterrence.

What motivated these subjects' punitive judgments? Here, too, an important part of the answer appears to be "outrage." Subjects indicated the extent to which they were "morally outraged" by the offenses in question, and the extent of moral outrage in response to a given offense was a pretty good predictor of the severity of punishment assigned to the perpetrator, although the effect here was weaker than that observed in Kahneman *et al.*'s study.⁷ Moreover, a structural equation model of these data suggests that the factors that had the greatest effect on people's judgments about punishment (severity of the crime, presence of mitigating circumstances) worked their effects through "moral outrage."

You will recall Small and Loewenstein's research on the "identifiable victim effect" discussed in the previous section. More recently they have documented a parallel effect in the domain of punishment. Subjects played an "investment game" in which individuals are given money that they may choose to put into a collective investment pool. The game allows individuals to choose the extent to which they will play cooperatively, benefiting the group at the chooser's expense. After the game, cooperators were given the opportunity to punish selfish players by causing them to lose money, but the punishing cooperators had to pay for the pleasure. As before, the crucial manipulation was between determined and

⁷ Some complications arise in interpreting the results of these two studies of "outrage" and punishment. It's not clear whether the "outrage" scale used by Kahneman *et al.* elicits a subjective report of the subject's emotional state or normative judgment concerning the defendant's behavior. A skeptic might say that the so-called "outrage" rating is really just a rating of the overall moral severity of the crime, which, not surprisingly, correlates with the extent to which people think it warrants punishment.

The Carlsmith *et al.* study addresses this worry (though not intentionally), and suggests that this worry may have some validity. The outrage measure used in the Carlsmith *et al.* study asks explicitly for a subjective report: "How morally outraged were you by this offense?" And, perhaps as a result of this change in tactic, the connection between "outrage" and punitive judgment is weakened from near perfect to pretty strong. Note, also that in choosing a strong word like "outrage" in a study of fairly mild, hypothetical crimes, the experimenters may have set the bar too high for subjective reports, thus weakening their results.

undetermined individuals, in this case the selfish players. Some subjects were asked, “How much would you like to punish uncooperative subject #4?” while other subjects were asked, “How much would you like to punish the uncooperative subject whose number you will draw?” Consistent with previous results, the average punishment was almost twice as high for the determined group, and once again subjects’ reports of their emotional responses (in this case a composite measure of anger and blame) tracked their behavior (Small and Loewenstein, in press).

Recent neuroimaging studies also suggest that the desire to punish is emotionally driven. Alan Sanfey, Jim Rilling, and colleagues (Sanfey et al., 2003) conducted a brain-imaging study of the ultimatum game in order to study the neural bases of people's sense of fairness. The ultimatum game works as follows. There is a sum of money, say \$10, and the first player (the proposer) makes a proposal as to how to divide it up between herself and the other player. The second player, the responder, can either accept the offer, in which case the money is divided as proposed, or reject the offer, in which case no one gets anything. Proposers usually make offers that are fair (i.e., fifty-fifty split) or close to fair, and responders tend to reject offers that are more than a little unfair. In other words, responders will typically pay for the privilege of punishing unfair proposers, even when the game is only played once. Why do people do this?

The answer, once again, implicates emotion. The experimenters found that unfair offers, as compared to fair offers, produced increased activity in the anterior insula, a brain region associated with anger, disgust, and autonomic arousal. Moreover, individuals' average levels of insula activity correlated positively with the percentage of offers they rejected and was weaker for trials in which the subject believed that the unfair offer was made by a computer program rather than a real person. Of course, it's conceivable that people were punishing in an attempt to deter unfair proposers from being unfair to others in the future, but that seems unlikely given the consistent finding that people are insensitive to manipulations that modulate the deterrent effects of punishment. Instead, it seems much more likely that people inflicted punishment for its own sake. And

once again, it seems that this retributivist tendency is emotionally-driven. A more recent neuroimaging study of punishment in response to violations of trust yields a similar conclusion (de Quervain et al., 2004). In this study, the extent of punishment was correlated with the level of activity in the caudate nucleus, a brain region associated with emotion and related more specifically to motivation and reward.

When people are asked in a general and abstract way about why it makes sense to punish, consequentialist arguments are prominent (Carlsmith et al., 2002; Weiner et al., 1997). But when people are presented with more concrete cases involving specific individuals carrying out specific offenses, people's judgments are largely, and in many cases completely, insensitive to factors affecting the consequences of punishment. This is so even when the consequentialist rationale for responding to these factors is highlighted and when people are explicitly instructed to think like consequentialists. It seems, then, that consequentialist thinking plays a negligible role in commonsense punitive judgment and that commonsense punitive judgment is almost entirely retributivist/deontological, so long as the matter is sufficiently concrete. Moreover, the available evidence, both from self-reports and neuroimaging data, suggests that people's deontological/retributivist punitive judgments are predominantly emotional, driven by feelings of anger or "outrage."

Emotion and the Moral Condemnation of Harmless Actions

According to consequentialists, actions are wrong because of their harmful consequences. In contrast, deontologists, along with many commonsense moralists, will condemn actions that do not cause harm in any ordinary sense. For example, a deontologist would likely say that it's wrong to break promises, regardless of whether doing so would have harmful consequences. Jonathan Haidt (Haidt et al., 1993) has conducted a series of studies of moral judgments made in response to harmless actions. Two themes relevant to the present

discussion emerge from this work. First, the moral condemnation of harmless action appears to be driven by emotion. Second, experience that encourages a more “cognitive” approach to moral decision-making tends to make people less willing to condemn harmless actions.

Haidt and two Brazilian colleagues conducted a cross-cultural study of moral judgment using a large set of subjects varying in socio-economic status (SES), nationality (Brazilian vs. American), and age (children vs. adults). The subjects were presented with a number of scenarios involving morally questionable, harmless actions:

1. A son promises his dying mother that he will visit her grave every week after she has died, but then doesn't because he is busy
2. A woman uses an old American/Brazilian flag to clean the bathroom
3. A family eats its dog after it has been killed accidentally by a car
4. A brother and sister kiss on the lips
5. A man masturbates using a dead chicken before cooking and eating it.

Subjects answered questions about each case: Is this action wrong? If so, why? Does this action hurt anyone? If you saw someone do this, would it bother you? Should someone who does this be stopped or punished? If doing this is the custom in some foreign country, is that custom wrong?

When people say that such actions are wrong, why do they say so? One hypothesis is that these actions are perceived as harmful, whether or not they really are (Turiel et al., 1987): Kissing siblings could cause themselves psychological damage. Masturbating with a chicken could spread disease, etc. If this hypothesis is correct, then we would expect people's answers to the question “Does this action hurt anyone?” to correlate with people's degree of moral condemnation, as indexed by affirmative answers to the questions: “Is this wrong?” “Should this person be stopped or punished?” “Is it wrong if it's the local custom?” Alternatively, if emotions drive moral condemnation in these cases, then we would expect people's answer the question “If you saw this,

would it bother you?” to better predict people’s answers to the above moral questions. As expected, Haidt and colleagues found that an affirmative answer to the “Would it bother you?” question was a better predictor of moral condemnation than an affirmative answer to the harm question.⁸

Equally interesting were the between-group differences. First, the high-SES subjects in Philadelphia and Brazil were far less condemning than their low-SES counterparts, so much so that the high-SES groups in Philadelphia and Brazil resembled each other more than they resembled their low-SES neighbors. Second, people from less “Westernized”⁹ cities tended to be more condemning. Third, children in both places tended to be more condemning than adults. In other words, education (SES), Westernization, and growing up were associated with more consequentialist judgments, in response to the scenarios used here. These three findings make sense in light of the model of moral judgment we’ve been developing, according to which intuitive emotional responses drive prepotent moral intuitions while “cognitive” control processes sometimes reign them in. Education is, to a large extent, the development of one’s “cognitive” capacities, learning to think in ways that are abstract, effortful, and often either non-intuitive or counter-intuitive. The Westernization factor is closely related. While Westerners may not be any more “cognitively” developed than members of other cultures, the Western tradition takes what is, from an anthropological perspective, a peculiarly “cognitive” approach to morality. Westerners are more likely than members of other cultures to argue for and justify their moral beliefs and values in abstract terms (Rozin, personal communication). Moreover, Western culture tends to be more pluralistic than other cultures, explicitly valuing

⁸ This result, however, only held for the subgroups that did the majority of the condemning. The subjects who were most reluctant to condemn harmless violations (chiefly high-SES, educated Westerners) found harm where others did not and cited that as a reason for condemnation, an effect that Haidt has documented elsewhere and which he has dubbed “moral dumbfounding” (Haidt, J., Bjorklund, F., and Murphy, S. (2000). Moral dumbfounding: When intuition finds no reason. (unpublished manuscript, university of virginia).

⁹ “Westernization” refers to “the degree to which each of three cities [Philadelphia and two Brazilian cities, Porto Alegre and Recife] has a cultural and symbolic life based on European traditions, including a democratic political structure and an industrialized economy,” (Haidt, J., Koller, S. H., and Dias, M. G. (1993). Affect, culture, and morality, or is it wrong to eat your dog? *J Pers Soc Psychol* 65, 613-628. Philadelphia is more westernized than Porto Alegre, which is more westernized than Recife.

multiple perspectives and an intellectual awareness that alternative perspectives exist. Finally, the capacity for “cognitive control” continues to develop through adolescence (Anderson et al., 2001; Paus et al., 1999). Children, like adults, are very good at feeling emotions such as anger, sympathy, disgust, etc., but unlike adults they are not very good at controlling their behavior when experiencing such feelings (Steinberg and Scott, 2003). Thus, as before, there seems to be a link between “cognition” and consequentialist judgment.

In this study, the connection between a reluctance to condemn and consequentialism is fairly straightforward. Consequentialists don’t condemn harmless actions.¹⁰ The connection between the tendency to condemn harmless actions and deontology is, however, less straightforward and more questionable. It’s not obvious, for example, that deontologists are any more likely than consequentialists to condemn flag-desecration or eating the family dog. Similar doubts apply to the case of kissing siblings and the man who masturbates with a dead chicken, though it’s worth noting that Kant argued that incest, masturbation, bestiality, and pretty much every other form of sexual experimentation are against the moral law (Kant, 1930; Kant, 1994). The broken promise case, however, is “downtown deontology.” Of course, not all deontologists would condemn harmlessly breaking a promise to one’s deceased mother, but anyone who would condemn such behavior (without appealing in some way to consequences) is exhibiting characteristically deontological behavior.¹¹ In light of this, it is worth examining this case more closely, and it turns out that this case fits the pattern for the inter-group differences quite well. Among high SES adults, the percentage of subjects in each city who said that this action should be stopped or punished ranged from 3% to 7%, while the percentage of low SES adults who said the same ranged from 20% (Philadelphia) to 57% (Recife, Brazil). Likewise, among high SES adults, the percentage who said that this

¹⁰ A consequentialist might favor a prohibition against a class of actions, some of which are not harmful, if the prohibition produces the best available consequences. Likewise, a consequentialist might *pretend* to condemn (or publicly condemn, while privately refraining from condemning) an action if this behavior were deemed beneficial.

¹¹ Subjects were asked to justify their answers, and typical justifications for condemning this action did not appeal to consequences, but rather simply stated that it’s wrong to break a promise.

behavior would be wrong even if it were the local custom ranged from 20% to 28%, while the corresponding percentages for low SES subjects ranged from 40% to 87%. The tendency to condemn this behavior also decreased with Westernization, and within every group children were more willing to condemn it than adults. (If you want someone to visit your grave when you're dead, you can't beat poor children from Recife, Brazil. 97% endorse punishing/stopping people who renege on grave-visiting promises, and 100% condemn cultures in which doing so is the custom.) Thus, the argument made above connecting "cognition" and consequentialism applies specifically to the case in which moral condemnation is most characteristically deontological. Haidt *et al.* did not provide data regarding the "Would it bother you?" question for this case specifically, but the fact that this case was not an exception to the general "cognitive" pattern (less condemnation in the presence of "cognition"-boosting factors) suggests that it is unlikely to be an exception to the general emotion-related pattern (condemnation correlated with negative emotions).

More powerful and direct evidence for the role of emotion in condemning harmless moral violations comes from two more recent studies. In the first of these, Thalia Wheatley and Jonathan Haidt (Wheatley and Haidt, 2005) gave hypnotizable individuals a post-hypnotic suggestion to feel a pang of disgust upon reading the word "often" (and to forget that they received this suggestion). The other subjects (also hypnotizable individuals) were given the same treatment, except that they were sensitized to the word "take." The subjects were then presented with scenarios, some of which involved no harm. In one scenario, for example, second cousins have a sexual relationship in which they "take/often go on weekend trips to romantic hotels in the mountains." The subjects who received the matching post-hypnotic suggestion (i.e. read the word to which they were hypnotically sensitized) judged this couple's actions to be more morally wrong than the other subjects.

In a second experiment, Wheatley and Haidt used a scenario in which the person described did nothing wrong at all. It was the case of a student council representative who "often picks" (or "tries to take up") topics of broad

interest for discussion. Many subjects who received matching post-hypnotic suggestions indicated that his behavior was somewhat wrong, and two subjects gave it high wrongness ratings. Subjects said things like: “It just seems like he’s up to something,” “It just seems so weird and disgusting,” and, “I don’t know [why it’s wrong], it just is.” Again, we see emotions driving people to non-consequentialist conclusions.

In a more recent (unpublished) study Simone Schnall, Jonathan Haidt, and Gerald Clore (Schnall et al., 2004) manipulated feelings of disgust, not with hypnosis, but by seating subjects at a disgusting desk while they filled out their questionnaires. (The desk was stained and sticky, located near an overflowing trashcan containing used pizza boxes and dirty-looking tissues, etc.) These subjects responded to a number of moral judgment scenarios including variations on the dog-eating and masturbation scenarios mentioned above. Here, as before, the disgust manipulation made people more likely to condemn these actions, though only for subjects who were rated as highly sensitive to their own bodily states.

Two Patterns of Moral Judgment

The experiments conducted by Greene *et al.*, Small and Loewenstein, Baron *et al.*, Kahneman *et al.*, Carlsmith *et al.*, Sanfey *et al.*, de Quervain *et al.*, and Haidt *et al.* together provide multiple pieces of independent evidence that deontological patterns of moral judgment are driven by emotional responses while consequentialist judgments are driven by “cognitive” processes. Any one of the results and interpretations described above may be questioned, but the convergent evidence assembled here makes a decent case for the association between deontology and emotion, especially since there is, to my knowledge, no empirical evidence to the contrary. Of course, deontologists may regard themselves and their minds as exceptions to the statistically significant and multiply convergent psychological patterns identified in these studies, but in my

opinion the burden is on them to demonstrate that they are psychologically exceptional in a way that preserves their self-conceptions.

Why should deontology and emotion go together? I believe the answer comes in two parts. First, moral emotion provides a natural solution to certain problems created by social life. Second, deontological philosophy provides a natural “cognitive” interpretation of moral emotion. Let us consider each of these claims in turn.

First, why moral emotions? In recent decades many plausible and complementary explanations have been put forth, and a general consensus seems to be emerging: The emotions most relevant to morality exist because they motivate behaviors that help individuals spread their genes *within a social context*. The theory of kin-selection explains why individuals have a tendency to care about the welfare of those individuals to whom they are closely related (Hamilton, 1964): Because close relatives share a high proportion of their genes, one can spread one’s own genes by helping close relatives spread theirs. The theory of reciprocal altruism explains the existence of a wider form of altruism: Genetically unrelated individuals can benefit from being nice to each other, so long as they are capable of keeping track of who is willing to repay their kindness (Trivers, 1971). More recent evolutionary theories of altruism attempt to explain the evolution of “strong reciprocity,” a broader tendency to reward cooperative behavior and punish uncooperative behavior, even in contexts in which the necessary conditions for kin-selection (detectable genetic relationships) and reciprocal altruism (detectable cooperative dispositions) are not met (Bowles and Gintis, 2004; Fehr and Rockenbach, 2004; Gintis, 2000). These theories explain the widespread human tendency to engage in cooperative behaviors (e.g. helping others and speaking honestly) and to avoid uncooperative behaviors (e.g. hurting others and lying), even when relatives and close associates are not involved. Moreover, these theories explain “altruistic punishment,” people’s willingness to punish anti-social behavior even when they cannot expect to benefit from doing so (Boyd et al., 2003; Fehr and Gächter, 2002; Fehr and Rockenbach, 2004). Other evolutionary theories make sense of other aspects of

morality. For example, the incest taboo can be explained as a mechanism for avoiding birth defects, which are more likely to result from matings between close relatives (Lieberman et al., 2003). Finally, the emerging field of cultural evolution promises to explain how moral norms (and cultural practices more broadly) develop and spread (Richerson and Boyd, 2005).

Such evolutionary accounts of moral phenomena have received a great deal of attention in recent years (Pinker, 2002; Sober and Wilson, 1998; Wright, 1994), and therefore I will not elaborate upon them here. I will simply assume that the general thrust of these theories is correct: that our most basic moral dispositions are evolutionary adaptations that arose in response to the demands and opportunities created by social life. The pertinent question here concerns the psychological implementation of these dispositions. Why should our adaptive moral behavior be driven by moral emotions as opposed to something else, such as moral reasoning? The answer, I believe, is that emotions are very reliable, quick, and efficient responses to re-occurring situations, whereas reasoning is unreliable, slow, and inefficient in such contexts. (Cf. Sober and Wilson (1998, Ch. 10) on altruistic emotions vs. hedonistic reasoning.)

Nature doesn't leave it to our powers of reasoning to figure out that ingesting fat and protein is conducive to our survival. Rather it makes us hungry, and gives us an intuitive sense that things like meat and fruit will satisfy our hunger. Nature doesn't leave it to us to figure out that fellow humans are more suitable mates than baboons. Instead, it endows us with a psychology that makes certain humans strike us as very appealing sexual partners, and makes baboons seem frightfully unappealing in this regard. And, finally, Nature doesn't leave it to us to figure out that saving a drowning child is a good thing to do. Instead, it endows us with a powerful "moral sense" that compels us to engage in this sort of behavior (under the right circumstances). In short, when Nature needs to get a behavioral job done, it does it with intuition and emotion wherever it can.

Thus, from an evolutionary point of view, it's no surprise that moral dispositions evolved, and it's no surprise that these dispositions are implemented

emotionally. Now, onto the second part of the explanation. Why should the existence of moral emotions give rise to the existence of deontological philosophy? To answer this question, we must appeal to the well-documented fact that humans are, in general, irrepressible explainers and justifiers of their own behavior. Psychologists have found repeatedly that when people don't know why they're doing what they're doing, they just make up a plausible sounding story (Haidt, 2001; Wilson, 2002).

Recall, for example, the panty-hose experiment described above. The subjects didn't know that they were drawn to items on the right side of the display, but when they were asked to explain themselves they made up perfectly rational, alternative explanations for their preferences (Nisbett and Wilson, 1977). In a similar experiment, Nisbett and Wilson (Nisbett and Wilson, 1977) induced subjects to prefer the laundry detergent Tide by priming them with word pairs like "ocean-moon" in a preceding memory test. When subjects explained their preferences, they said things like "Tide is the best-known detergent," or "My mother uses Tide," or "I like the Tide box." In an early experiment by Maier (Maier, 1931; Nisbett and Wilson, 1977), subjects had to figure out a way to tie together two cords hanging from the ceiling, a challenging task since the cords were too far apart to be reached simultaneously. The solution was to tie a heavy object to one of the cords so that it could swing like a pendulum. The subject could then hold onto one cord while waiting for the other one to swing into reach. Maier was able to help his subjects solve this problem by giving them a subtle clue. As he was walking around the room he would casually put one of the cords in motion. The subjects who were aided by this clue, however, were unaware of its influence. Instead, they readily attributed their insights to a different, more conspicuous cue (Maier's twirling a weight on a cord), despite the fact that this cue was demonstrated to be useless in other versions of the experiment. In a similar experiment Dutton and Aron (Dutton and Aron, 1974; Wilson, 2002) had male subjects cross a scary footbridge spanning a deep gorge, after which they were met by an attractive female experimenter. Control subjects rested on a bench before encountering the attractive experimenter. The subjects who had

just braved the scary bridge, with their sweaty palms and hearts a'pounding, were over twice as likely as the control subjects to call the experimenter later and ask her for a date. These individuals (many of them, at any rate) interpreted their increased physiological arousal as increased attraction to the woman they had met.

The tendency toward *post-hoc* rationalization is often revealed in studies of people with unusual mental conditions. Patients with Korsakoff's amnesia and related memory disorders are prone to "confabulation." That is, they attempt to paper over their memory deficits by constructing elaborate stories about their personal histories, typically delivered with great confidence and with no apparent awareness that they're making stuff up. For example, a confabulating patient seated near an air conditioner was asked if he knew where he was. He replied that he was in an air-conditioning plant. When it was pointed out that he was wearing pajamas, he said, "I keep them in my car and will soon change into my work clothes" (Stuss et al., 1978). Likewise, individuals acting under post-hypnotic suggestion will sometimes explain away their behaviors in elaborately rational terms. In one case, a hypnotized subject was instructed to place a lampshade on another person's head upon perceiving an arbitrary cue. He did as instructed, but when he was asked to explain why he did what he did, he made no reference to the post-hypnotic suggestion or the cue: "Well, I'll tell you. It sounds queer but it's just a little experiment in psychology. I've been reading on the psychology of humor and I thought I'd see how you folks reacted to a joke that was in very bad taste." (Estabrooks, 1943; Wilson, 2002). Perhaps the most striking example of this kind of *post-hoc* rationalization comes from studies of split-brain patients, people in whom there is no direct neuronal communication between the cerebral hemispheres. In one study, a patient's right hemisphere was shown a snow scene and instructed to select a matching picture. Using his left hand, the hand controlled by the right hemisphere, he selected a picture of a shovel. At the same time, the patient's left hemisphere, the hemisphere that is dominant for language, was shown a picture of a chicken claw. The patient was asked verbally why he chose the shovel with his left hand. He answered, "I saw

a claw and picked a chicken, and you have to clean out the chicken shed with a shovel.” (Gazzaniga and Le Doux, 1978; Wilson, 2002). Gazzaniga and LeDoux argue that these sorts of confabulations are not peculiar to split-brain patients, that this tendency was not created when these patients’ inter-cerebral communications lines were cut. Rather, they argue, we are all confabulators of a sort. We respond to the conscious deliverances of our unconscious perceptual, mnemonic, and emotional processes by fashioning them into a rationally sensible narrative, and without any awareness that we are doing so. This widespread tendency for rationalization is only revealed in carefully controlled experiments in which the psychological inputs and behavioral outputs can be carefully monitored, or in studies of abnormal individuals who are forced to construct a plausible narrative out of meager raw material.

We are now ready to put two and two together. What should we expect from creatures who (1) exhibit social/moral behavior that is driven largely by intuitive emotional responses, and (2) are prone to rationalization of their behaviors? The answer, I believe, is deontological moral philosophy. What happens when we contemplate pushing the large man off the footbridge? If I’m right, we have an intuitive emotional response that says “No!” This nay-saying voice can be overridden, of course, but as far as the voice itself is concerned, there is no room for negotiation. Whether or not we can ultimately justify pushing the man off the footbridge, it will always *feel* wrong. And what better way to express that feeling of non-negotiable absolute wrongness than via the most central of deontological concepts, the concept of a *right*: You can’t push him to his death because that would be a violation of his *rights*. Likewise, you can’t let that baby drown because you have a *duty* to save it.

Deontology, then, is a kind of moral confabulation. We have strong feelings that tell us in clear and uncertain terms that some things *simply cannot be done* and that other things *simply must be done*. But it’s not obvious how to make sense of these feelings, and so we, with the help of some especially creative philosophers, make up a rationally appealing story: There are these things called “rights” which people have, and when someone has a right you

can't do anything that would take it away. It doesn't matter if the guy on the footbridge is toward the end of his natural life, or if there are seven people on the tracks below instead of five. If the man has a right, then *the man has a right*. As John Rawls (Rawls, 1971) famously said, "Each person possesses an inviolability founded on justice that even the welfare of society as a whole cannot override" and, "In a just society the rights secured by justice are not subject to political bargaining or to the calculus of social interests" (pp. 3-4) These are applause lines because they make emotional sense. Deontology, I believe, is a natural "cognitive" expression of our deepest moral emotions.

This hypothesis raises a further question. Why just deontology? Why not suppose that all moral philosophy, even all moral reasoning, is a rationalization of moral emotion? (This is the strong form of the view defended by Jonathan Haidt (Haidt, 2001), whose argument is the model for the argument made here.¹²) The answer, I think, is that consequentialist moral judgment is not driven by emotion, or at least not driven by the sort of "alarm bell" emotion that drives deontological judgment. The evidence presented above supports this hypothesis, suggesting that consequentialist judgment is less emotional and more "cognitive," but it doesn't explain why this should be so. I argued above that there is a natural mapping between the content of deontological philosophy and the functional properties of alarm-like emotions. Likewise, I believe that there is a natural mapping between the content of consequentialist philosophy and the functional properties of "cognitive" processes. Indeed, I believe that consequentialism is inherently "cognitive," that it couldn't be implemented any other way.

Consequentialism is, by its very nature, systematic and aggregative. It aims to take nearly everything into account, and grants that nearly everything is negotiable. All consequentialist decision-making is a matter of balancing competing concerns, taking into account as much information as is practically feasible. Only in hypothetical examples in which "all else is equal" does consequentialism give clear answers. For real-life consequentialism, everything

¹² Haidt, however, believes that philosophers may be exceptional in that they actually do reason their way to moral conclusions (Kuhn, D., 1991).

is a complex guessing game, and all judgments are revisable in light of additional details. There is no moral clarity in consequentialist moral thought, with its approximations and simplifying assumptions. It is fundamentally actuarial.

Recall the definition of “cognitive” proposed above: “Cognitive” representations are inherently neutral representations, ones that, unlike emotional representations, do not automatically trigger particular behavioral responses or dispositions. Once again, the advantage of having such neutral representations is that they can be mixed and matched in a situation-specific way without pulling the agent in multiple behavioral directions at once, thus enabling highly flexible behavior. These are precisely the sorts of representations that a consequentialist needs to make a judgment based on aggregation, one that takes all of the relevant factors into account: “Is it okay to push the guy off the bridge if he’s about to cure cancer?” “Is it okay to go out for sushi when the extra money could be used to promote health education in Africa?” And so on. Deontologists can dismiss these sorts of complicated, situation-specific questions, but consequentialists cannot, which is why, I argue, that consequentialism is inescapably “cognitive.”

Some clarifications: First, I am not claiming that consequentialist judgment is emotionless. On the contrary, I am inclined to agree with Hume (Hume, 1978) that all moral judgment must have some affective component, and suspect that the consequentialist weighing of harms and benefits is an emotional process. But, if I’m right, two things distinguish this sort of process from those associated with deontology. First, this is, as I’ve said, a weighing process and not “alarm” processes. The sorts of emotions hypothesized to be involved here say, “Such-and-such matters this much. Factor it in.” Whereas the emotions hypothesized to drive deontological judgment are far less subtle. They are, as I’ve said, alarm signals that issue simple commands: “Don’t do it!” or “Must do it!” While such commands can be overridden, they are designed to dominate the decision rather than merely influence it.

Second, I am not claiming that deontological judgment can’t be “cognitive.” Indeed, I believe that sometimes it is. (See below.) Rather, my hypothesis is

that deontological judgment is affective at its core, while consequentialist judgment is inescapably “cognitive.” One could, in principle, make a characteristically deontological judgment by thinking explicitly about the categorical imperative and whether the action in question is based on a maxim that could serve as a universal law. And if one were to do that, then the psychological process would be “cognitive.” What I am proposing, however, is that this is not how characteristically deontological conclusions tend to be reached, and that, instead, they tend to be reached on the basis of emotional responses. This contrasts with consequentialist judgments which, according to my hypothesis, simply cannot be implemented in an intuitive, emotional way. The only way to reach a distinctively consequentialist judgment (i.e. one that doesn’t coincide with a deontological judgment) is to actually go through the consequentialist, cost-benefit reasoning using one’s “cognitive” faculties, the ones based in the dorsolateral prefrontal cortex (Greene et al., 2004).

This psychological account of consequentialism and deontology makes sense of certain aspects of their associated phenomenologies. I have often observed that consequentialism strikes students as very appealing, even as tautologically true, when presented in the abstract, but that its appeal is easily undermined by specific counterexamples. (Cf. the discussion above contrasting people’s real-world motivations and abstract justifications for punishment.) When a first-year ethics student asks, “But isn’t it obvious that one should do whatever will produce the most good?” all you have to do is whip out the *footbridge* case and you’ve made your point. Whatever initial “cognitive” appeal consequentialist principles may have is quickly neutralized by a jolt of emotion, and the student is a newly converted deontologist: “Why is it wrong to push the man off the footbridge? Because he has a *right*, an inviolability founded on justice that even the welfare of society as a whole cannot override!” Then, it’s time for a new counter-example: “What if the trolley is headed for a detonator that will set off a nuclear bomb that will kill half a million people?” Suddenly the welfare of society as a whole starts to sound important again. “Cognition” strikes back with a more compelling utilitarian rationale, and the student is appropriately puzzled. As this

familiar dialectic illustrates, the hypothesis that deontology is emotionally based explains the “*NEVER!*—except sometimes” character of rights-based, deontological ethics. An alarm-like emotional response presents itself as unyielding and absolute, until an even more compelling emotional or “cognitive” rationale comes along to override it.

This hypothesis also makes sense of certain deontological anomalies, which I suspect will turn out to be the “exceptions that prove the rule.” I’ve argued that deontology is driven by emotion, but I suspect this is not always the case. Consider, for example, Kant’s infamous claim that it would be wrong to lie to a would-be murderer in order to protect a friend who has taken refuge in one’s home (Kant, 1983). Here, in a dramatic display of true intellectual integrity, Kant sticks to his theory and rejects the intuitive response. (He “bites the bullet,” as philosophers say.) But what’s interesting about this bit of Kantian ethics is that it’s something of an embarrassment to contemporary Kantians, who are very keen to explain how Kant somehow misapplied his own theory in this case (Korsgaard, 1996a). Presumably, the same goes for Kant’s views of sexual morality (Kant, 1930, pp. 169-171; Kant, 1994). Modern academics are no longer so squeamish about lust, masturbation, and homosexuality, and so Kant’s old-fashioned views on these topics have to be explained away, which is not hard, since his arguments were never terribly compelling to begin with (See epigraph). If you want to know which bits of Kant contemporary Kantians will reject, follow the emotions.

III. Normative Implications

Psychological “Is” and Moral “Ought”

The hypotheses advanced above concerning the respective psychological bases of consequentialism and deontology could certainly be wrong. But whether they’re right or wrong cannot be determined from the armchair. Rather, it is an

empirical matter. And although these hypotheses remain open to empirical challenge, I am, from here on, going to assume that they are correct in order to explore their broader philosophical implications. Since most moral philosophers do not regard their views as contingent upon the outcomes of particular debates in experimental psychology, this assumption should not be regarded as unduly restrictive.

Indeed, moral philosophers tend to steer clear of scientific controversies whenever possible on the grounds that scientific details are largely irrelevant to their enterprise: Science is about what *is*, while morality is about what *ought to be*, and never the twain shall meet (Hume, 1978; Moore, 1966). Contrary to this received moral wisdom, I believe that science does matter for ethics, not because one can derive moral truths from scientific truths, but because scientific information can undermine factual assumptions on which moral thinking implicitly depends. The key point of contact between moral philosophy and scientific moral psychology is moral intuition. Moral philosophers, from Plato (Plato, 1987) on down, have relied on their intuitive sense of right and wrong to guide them in their attempts to make sense of morality. The relevance of science, then, is that it can tell us how our moral intuitions work and where they come from. And once we understand our intuitions a bit better we may view them rather differently. This goes not only for moralists who rely explicitly on moral intuitions (Ross, 1930), but also for moralists who are unaware of the extent to which their moral judgments are shaped by intuition.

In recent years, several philosophers and scientists have questioned the reliability of moral intuitions and argued that understanding the psychology of moral intuition has normative implications (Baron, 1994; Greene, 2003; Horowitz, 1998; Sinnott-Armstrong, 2004; Unger, 1996). I will do the same, but in the following more specific way. I will argue that our understanding of moral psychology, as described above, casts doubt on deontology as a school of normative moral thought.

Rationalism, Rationalization, and Deontological Judgment

Your friend Alice goes on many dates, and after each one she reports back to you. When she extols the people she likes and complains about the ones she dislikes, she cites a great many factors. This one is brilliant. That one is self-absorbed. This one has a great sense of humor. That one is a dud. And so on. But then you notice something: All the people she likes are exceptionally tall. Closer inspection reveals that, after scores of dates over several years, she's not given the thumb's up to anyone under six-foot-four, and hasn't turned down anyone over this height. (You plug Alice's dating data into your statistics software and confirm that height is a near perfect predictor of Alice's preferences.) Suddenly it seems that Alice's judgment is not what you had believed, and certainly not what she believes. Alice, of course, believes that her romantic judgments are based on a variety of complicated factors. But, if the numbers are to be believed, she basically has a height fetish, and all of her talk about wit and charm and kindness is mere rationalization.

What this example illustrates is that it's possible to spot a rationalizer without picking apart the rationalizer's reasoning. Instead you need only do two things: First, you have to find a factor that predicts the rationalizer's judgments. Second, you have to show that the factor that predicts the rationalizer's judgments is not plausibly related to the factors that, according to the rationalizer, are the bases for his or her judgments. Using this strategy, I believe that one can make a pretty good case against rationalist versions of deontology such as Kant's, i.e. the ones according to which characteristically deontological moral judgments are justified in terms of abstract theories of rights, duties, etc. The case against such theories is already implicit in the empirical material presented above, but it's worth spelling it out.

The bulk of this essay has been devoted to satisfying the first of the two requirements listed above, i.e. to identifying a factor, namely emotional response, that predicts deontological judgment. Next, we must consider the nature of the relationship between this predictive factor and the factors that, according to

rationalist deontologists, are the bases for their judgments. By definition, a rationalist cannot say that that some action is right or wrong *because* of the emotions we feel in response to it. Nevertheless, as an empirical matter of fact (we are assuming) there is a remarkable correspondence between what rationalist deontological theories tell us to do and what our emotions tell us to do. Thus, in light of these data, there are a series of coincidences for which various rationalist deontologists must account: For example, according to Judith Jarvis Thomson (Thomson, 1986; Thomson, 1990) and Frances Kamm (Kamm, 1993; Kamm, 1996) (both of whom count as rationalists for our purposes), there is a complicated, highly abstract theory of rights that explains why it's okay to sacrifice one life for five in the *trolley* case but not in the *footbridge* case, and it *just so happens* that we have a strong negative emotional response to the latter case but not to the former. Likewise, according to Colin McGinn (McGinn, 1999) and Frances Kamm (Kamm, 1999), there is a theory of duty that explains why we have an obligation to help Singer's drowning child but no comparable obligation to save starving children on the other side of the world, and it *just so happens* that we have strong emotional responses to the former individuals but not to the latter. According to Kant (Kant, 2002) and many other legal theorists (Lacey, 1988; Ten, 1987), there is a complicated abstract theory of punishment that explains why we ought to punish people regardless of whether there are social benefits to be gained in doing so, and it *just so happens* that we have emotional responses that incline us to do exactly that. The categorical imperative prohibits masturbation because it involves using oneself as a means (Kant, 1994), and it *just so happens* that the categorical imperative's chief proponent finds masturbation really, really disgusting. (See epigraph.) And so on.

Kant, as a citizen of Eighteenth Century Europe, has a ready explanation for these sorts of coincidences: God, in his infinite wisdom, endowed people with emotional dispositions designed to encourage them to behave in accordance with the moral law. Kant famously avoided invoking God in his philosophical arguments, but it's plausible to think that his faith prevented him, along with nearly everyone else of his day, from being puzzled by the order and harmony of

the natural world, including its harmony with the moral law. Moreover, in light of his background assumptions, you can't really blame Kant for trying to rationalize his moral intuitions. His intuitions derive from his human nature ("the moral law within" (Kant, 1993)), and ultimately from God. God's a smart guy, Kant must have thought. He wouldn't give people moral intuitions *willy nilly*. Instead, we must have the intuitions we have for *good reasons*. And so Kant set out to discover those reasons, if not by force of reason then by feat of imagination.

Present day rationalist deontologists, as citizens of the Twenty-First Century, can't depend on the notion that God gave us our moral emotions to encourage us to behave in accordance with the rationally discoverable deontological moral truth. Instead, they need some sort of naturalistically respectable explanation for the fact that the conclusions reached by rationalist deontologists, as opposed to those reached by consequentialists, appear to be driven by alarm-like emotional responses. And their explanation needs to compete with the alternative proposed here, namely that rationalist deontological theories are rationalizations for these emotional responses—an explanation that already has a leg up given (a) that so much human behavior appears to be intuitive (Bargh and Chartrand, 1999) and (b) the well-documented tendency for people to rationalize their intuitive behavior (Haidt, 2001; Wilson, 2002).

What sort of explanation can rationalist deontologists give? They will have to say, first, that the correspondence between deontological judgment and emotional engagement is not a coincidence and, second, that our moral emotions somehow track the rationally discoverable deontological moral truth. They can't say that our emotional responses are the *basis* for the moral truth, however, because they are *rationalists*. So they're going to have to explain how some combination of biological and cultural evolution managed to give us emotional dispositions that correspond to an independent, rationally discoverable moral truth that isn't based on emotion.

Those charged with this task face another disadvantage right off the bat, which is the chief point I wish to make here. There are good reasons to think that our distinctively deontological moral intuitions (here, the ones that conflict with

consequentialism) reflect the influence of morally irrelevant factors, and are therefore unlikely to track the moral truth.

Take, for example, the *trolley* and *footbridge* cases. I've argued that we draw an intuitive moral distinction between these two cases because the moral violation in the *footbridge* case is "up close and personal" while the moral violation in the *trolley* case is not. Moreover, I've argued that we respond more emotionally to moral violations that are "up close and personal" because those are the sorts of moral violations that existed in the environment in which we evolved. In other words, I've argued that we have a characteristically deontological intuition regarding the *footbridge* case because of a contingent, non-moral feature of our evolutionary history. Moreover, I've argued that the same "up close and personal" hypothesis makes sense of the puzzling intuitions surrounding Peter Singer's aid cases and the identifiable victim effect, thus adding to its explanatory power.

The key point is that this hypothesis is at odds with any hypothesis according to which our moral intuitions in response to these cases reflect deep, rationally discoverable moral truths. Of course, the hypothesis I've advanced could be wrong. But do rationalist deontologists want to count on it? And do they have any more plausible positive explanations to offer in its place?

A similar hypothesis can explain our inclinations toward retributive punishment. Consequentialists say that punishments should only be inflicted insofar as they are likely to produce good consequences (Bentham, 1982). Deontologists such as Kant (Kant, 2002), along with most people (Baron et al., 1993; Baron and Ritov, 1993), are retributivists. They judge in favor of punishing wrong-doers as an end in itself, even when doing so is unlikely to promote good consequences in the future. Is this a moral insight on their part, or just a by-product of our evolved psychology? The available evidence suggests the latter.

As discussed above, it appears that the emotions that drive us to punish wrong-doers evolved as an efficient mechanism for stabilizing cooperation, both between individuals (Trivers, 1971) and within larger groups (Bowles and Gintis, 2004; Boyd et al., 2003; Fehr and Gächter, 2002). In other words, according to

these models, we are disposed to punish because of this disposition's biological consequences. Moreover, natural selection, in furnishing us with this disposition, had a "choice," so to speak. On the one hand, nature could have given us a disposition to punish by giving us, first, an innate desire to secure the benefits of future cooperation and, second, some means by which to recognize that punishing non-cooperators is often a good way to achieve this end. In other words, nature could have made us punishment consequentialists. Nature's other option was to give us a direct desire to punish non-cooperators as an end in itself, even if punishing in some cases does no (biological) good. As noted above, nature faces this sort of choice every time it generates a behavioral adaptation, and in pretty much every case, Nature takes the more direct approach: Psychologically speaking, we desire things like food, sex, and a comfortable place to rest because they are pleasant (and because their absence is unpleasant) and not because we believe they will enhance our biological fitness. The disposition toward punishment appears to be no exception to this general pattern. Psychologically speaking, we punish primarily because we find punishment satisfying (de Quervain et al., 2004), and find unpunished transgressions distinctly un-satisfying (Carlsmith et al., 2002; Kahneman et al., 1998; Sanfey et al., 2003).

In other words, the emotions that drive us to punish are *blunt biological instruments*. They evolved because they drive us to punish in ways that lead to (biologically) good consequences. But, *as a by-product of their simple and efficient design*, they also lead us to punish in situations in which no (biologically) good consequences can be expected. Thus, it seems that, as an evolutionary matter of fact, we have a taste for retribution, not because wrong-doers truly deserve to be punished regardless of the costs and benefits, but because retributive dispositions are an efficient way of inducing behavior that allows individuals living in social groups to more effectively spread their genes.

Of course it's possible that there's a coincidence here. It could be that it's part of the rationally discoverable moral truth that people really do deserve to be punished as an end in itself. And, at the same time, it could *just so happen* that

natural selection, in devising an efficient means for promoting biologically advantageous consequences, furnished us with emotionally-based dispositions that lead us to this conclusion. But this seems unlikely. Rather, it seems that retributivist theories of punishment are just rationalizations for our retributivist feelings, and that these feelings only exist because of the morally-irrelevant constraints placed on natural selection in designing creatures that behave in fitness-enhancing ways. In other words, the natural history of our retributivist dispositions makes it unlikely that they reflect any sort of deep moral truth.

I should emphasize that I am not claiming that consequentialist theories of punishment are correct because the tendency to punish evolved in order to produce good consequences. These “good consequences” need only be good from a biological point of view, and to assume that our ends must coincide with the ends of natural selection would be to commit the naturalistic fallacy in its original form (Moore, 1966). At the same time, I wish to make it clear that I am not asserting that any tendency that we have as an evolutionary by-product is automatically wrong or misguided. I wouldn’t claim, for example, that it’s wrong to love one’s adopted children (who do not share one’s genes) or to use birth control simply because these behaviors thwart nature’s “intentions.” My claim, at this point, is simply that it’s unlikely that inclinations that evolved as evolutionary by-products correspond to some independent, rationally discoverable moral truth. Instead, it’s more parsimonious to suppose that when we feel the pull of retributivist theories of punishment, we are merely gravitating toward our evolved emotional inclinations and not toward some independent moral truth¹³.

What turn-of-the-millennium science is telling us is that human moral judgment is not a pristine rational enterprise, that our moral judgments are driven by a hodgepodge of emotional dispositions, which themselves were shaped by a hodgepodge of evolutionary forces, both biological and cultural. Because of this, it is exceedingly unlikely that there is any rationally coherent normative moral theory that can accommodate our moral intuitions. Moreover, anyone who

¹³ That is, a truth independent of the details of human moral psychology and natural events that shaped it.

claims to have such a theory, or even part of one, almost certainly doesn't. Instead, what that person probably has is a moral rationalization.

It seems then, that we have somehow crossed the infamous "is"- "ought" divide.¹⁴ How did this happen? Didn't Hume (Hume, 1978) and Moore (Moore, 1966) warn us against trying to derive an "ought" from an "is?" How did we go from descriptive scientific theories concerning moral psychology to skepticism about a whole class of normative moral theories? The answer is that we did not, as Hume and Moore anticipated, attempt to *derive* an "ought" from an "is." That is, our method has been *inductive* rather than *deductive*. We have inferred on the basis of the available evidence that the phenomenon of rationalist deontological philosophy is best explained as a rationalization of evolved emotional intuition (Harman, 1977).

Missing the Deontological Point

I suspect that rationalist deontologists will remain unmoved by the arguments presented here. Instead, I suspect, they will insist that I have simply misunderstood what Kant and like-minded deontologists are all about. Deontology, they will say, isn't about this intuition or that intuition. It's not defined by its normative differences with consequentialism. Rather, deontology is about taking humanity seriously. Above all else, it's about respect for persons. It's about treating others as fellow rational creatures rather than as mere objects, about acting for reasons rational beings can share. And so on (Korsgaard, 1996a; Korsgaard, 1996b).

This is, no doubt, how many deontologists see deontology. But this insider's view, as I've suggested, may be misleading. The problem, more specifically, is that it defines deontology in terms of values that are not

¹⁴ Most agree that the "is"- "ought" divide can be crossed when the "is" amounts to a constraint on what can be done, and is *a fortiori* a constraint on what "ought" to be done. For example, if it is the case that you are dead, then it is not the case that you *ought* to vote. The move from "is" to "ought" discussed below, however, is more substantive, and correspondingly more controversial.

distinctively deontological, though they may appear to be from the inside. Consider the following analogy with religion. When one asks a religious person to explain the essence of his religion, one often gets an answer like this: “It’s about love, really. It’s about looking out for other people, looking beyond oneself. It’s about community, being part of something larger than oneself.” This sort of answer accurately captures the phenomenology of many people’s religion, but it’s nevertheless inadequate for distinguishing religion from other things. This is because many, if not most, non-religious people aspire to love deeply, look out for other people, avoid self-absorption, have a sense of a community, and be connected to things larger than themselves. In other words, secular humanists and atheists can assent to most of what many religious people think religion is all about. From a secular humanist’s point of view, in contrast, what’s distinctive about religion is its commitment to the existence of supernatural entities as well as formal religious institutions and doctrines. And they’re right. These things really do distinguish religious from non-religious practices, though they may appear to be secondary to many people operating from within a religious point of view.

In the same way, I believe that most of the standard deontological/Kantian self-characterizations fail to distinguish deontology from other approaches to ethics. (See also Kagan (Kagan, 1997, pp. 70-78.) on the difficulty of defining deontology.) It seems to me that consequentialists, as much as anyone else, have respect for persons, are against treating people as mere objects, wish to act for reasons that rational creatures can share, etc. A consequentialist respects other persons, and refrains from treating them as mere objects, by counting every person’s well-being in the decision-making process. Likewise, a consequentialist attempts to act according to reasons that rational creatures can share by acting according to principles that give equal weight to everyone’s interests, i.e. that are impartial. This is not to say that consequentialists and deontologists don’t differ. They do. It’s just that the real differences may not be what deontologists often take them to be.

What, then, distinguishes deontology from other kinds of moral thought? A good strategy for answering this question is to start with concrete disagreements between deontologists and others (such as consequentialists) and then work backward in search of deeper principles. This is what I've attempted to do with the *trolley* and *footbridge* cases, and other instances in which deontologists and consequentialists disagree. If you ask a deontologically-minded person why it's wrong to push someone in front of speeding trolley in order to save five others, you will get characteristically deontological answers. Some will be tautological: "Because it's murder!" Others will be more sophisticated: "The ends don't justify the means." "You have to respect people's rights." But, as we know, these answers don't really explain anything, because if you give the same people (on different occasions) the *trolley* case or the *loop* case (See above), they'll make the opposite judgment, even though their initial explanation concerning the *footbridge* case applies equally well to one or both of these cases. Talk about rights, respect for persons, and reasons we can share are natural attempts to explain, in "cognitive" terms, what we feel when we find ourselves having emotionally driven intuitions that are odds with the cold calculus of consequentialism. Although these explanations are inevitably incomplete, there seems to be "something deeply right" about them because they give voice to powerful moral emotions. But, as with many religious people's accounts of what's essential to religion, they don't really explain what's distinctive about the philosophy in question.

In sum, if it seems that I have simply misunderstood what Kant and deontology are all about, it's because I am advancing an alternative hypothesis to the standard Kantian/deontological understanding of what Kant and deontology are all about. I am putting forth an empirical hypothesis about the hidden psychological essence of deontology, and it cannot be dismissed *a priori* for the same reason that tropical islanders cannot know *a priori* whether ice is a form of water.

Evolutionary Moral Psychology and Anthropocentric Morality

Above I made a case against rationalist deontology, the idea that our deontological moral intuitions can be justified via abstract theories of rights, duties, etc. There are however, more modest forms of deontology. Rather than standing by our moral intuitions on the assumption that they can be justified by a rational theory, we might just stand by them because they are *ours*. That is, one might take an *anthropocentric* approach to morality (See Haidt and Bjorklund, this volume), giving up on the Enlightenment dream of deriving moral truths from first principles and settling, instead, for a morality that is contingently human.

This is the direction in which moral philosophy has moved in recent decades. Virtue ethics defines moral goodness in terms of human character (Crisp and Slote, 1997; Hursthouse, 1999). Like-minded “sensibility theorists” regard being moral as a matter of having the right sort of distinctively human sensibility (McDowell, 1988; Wiggins, 1987). Ethicists with a more metaphysical bent speak of moral properties that are “response dependent” (Johnston, 1995), moral sentiments that correspond to “quasi-real” moral properties (Blackburn, 1993), and moral properties that are “homeostatic clusters” of natural properties (Boyd, 1988). And even within the Kantian tradition, many emphasize the “construction” of moral principles that, rather than being true, are “reasonable for us” (Rawls, 1995), or, alternatively, the normative demands that follow from our distinctively human “practical identities” (Korsgaard, 1996b).

In short, moral philosophy these days is decidedly anthropocentric in the sense that very few philosophers are actively challenging anyone’s moral intuitions. They acknowledge that our moral virtues, sensibilities, identities etc. may change over time, but they are not, for the most part, actively trying to change them.

The argument presented above makes trouble for people in search of rationalist theories that can explain and justify their emotionally-driven deontological moral intuitions. But rationalist deontologists may not be the only ones who should think twice. The arguments presented above cast doubt on the

moral intuitions in question regardless of whether one wishes to justify them in abstract theoretical terms. This is, once again, because these intuitions appear to have been shaped by morally irrelevant factors having to do with the constraints and circumstances of our evolutionary history. This is a problem for anyone who is inclined to stand by these intuitions, and that “anyone” includes nearly everyone.

I’ve referred to these intuitions and the judgments they underwrite as “deontological,” but perhaps it would be more accurate to call them non-consequentialist (Baron, 1994). After all, you don’t have to be a card-carrying deontologist to think that it’s okay to eat in restaurants when people in the world are starving, that it’s inherently good that criminals suffer for their crimes, and that it would be wrong to push the guy off the *footbridge*. These judgments are perfectly commonsensical, and it seems that the only people who are inclined to question them are card-carrying consequentialists.

Does that mean that all non-consequentialists need to rethink at least some of their moral commitments? I humbly suggest that the answer is yes. Let us consider, once more, Peter Singer’s argument concerning the moral obligations that come with affluence. Suppose, once again, that the evolutionary and psychological facts are exactly as I’ve said. That is, suppose that the *only* reason we say that it’s wrong to abandon the drowning child but okay to ignore the needs of starving children overseas is that the former pushes our emotional buttons while the latter do not. And let us suppose further that the *only* reason that faraway children fail to push our emotional buttons is that we evolved in an environment in which it was impossible to interact with faraway individuals. Could we then stand by our commonsense intuitions? Can we, in good conscience, say, “I live a life of luxury while ignoring the desperate needs of people far away because I, through an accident of human evolution, am emotionally insensitive to their plight. Nevertheless, my failure to relieve their suffering, when I could easily do otherwise, is perfectly justified.” I don’t know about you, but I find this combination of assertions uncomfortable. This is not to say, of course, that I am comfortable with the idea of giving up most of my

worldly possessions and privileges in order to help strangers. After all, I'm only human. But, for me at least, understanding the source of my moral intuitions shifts the balance, in this case as well as other cases, in a more Singerian, consequentialist direction. As a result of understanding the psychological facts, I'm less complacent about my all-too-human tendency to ignore distant suffering. Likewise, when I understand the roots of my retributive impulses, I'm less likely to afford them moral authority. Likewise for whatever hang-ups I may have about deviant, but harmless sexual behavior.

Taking these arguments seriously, however, threatens to put us on a second slippery slope (in addition to the one leading to altruistic destitution): How far can the empirical debunking of human moral nature go? If science tells me that I love my children more than other children only because they share my genes (Hamilton, 1964), should I feel uneasy about loving them extra? If science tells me that I'm nice to other people only because a disposition to be nice ultimately helped my ancestors spread their genes (Trivers, 1971), should I stop being nice to people? If I care about my self only because I am biologically programmed to carry my genes into the future, should I stop caring about myself? It seems that one who is unwilling to act on human tendencies with amoral evolutionary causes is ultimately unwilling to be human. Where does one draw the line between correcting the near-sightedness of human moral nature and obliterating it completely?

This, I believe, is among the most fundamental moral questions we face in an age of growing scientific self-knowledge, and I will not attempt to address it here. Elsewhere I argue that consequentialist principles, while not true, provide the best available standard for public decision-making and for determining which aspects of human nature it is reasonable to try to change and which ones we would be wise to leave alone (Greene, 2002; Greene and Cohen, 2004).

Adolphs, R. (2002). Neural systems for recognizing emotion. *Curr Opin Neurobiol* 12, 169-177.

Allison, T., Puce, A., and McCarthy, G. (2000). Social perception from visual cues: Role of the sts region. *Trends Cogn Sci* 4, 267-278.

- Allman, J., Hakeem, A., and Watson, K. (2002). Two phylogenetic specializations in the human brain. *Neuroscientist* 8, 335-346.
- Anderson, V. A., Anderson, P., Northam, E., Jacobs, R., and Catroppa, C. (2001). Development of executive functions through late childhood and adolescence in an Australian sample. *Dev Neuropsychol* 20, 385-406.
- Aquinas, T. (1988). Of killing. In *On law, morality, and politics*, W. P. Baumgarth, and R. J. Regan, eds. (Indianapolis/Cambridge, Hackett Publishing Co.), pp. 226-227.
- Bargh, J. A., and Chartrand, T. L. (1999). The unbearable automaticity of being. *American Psychologist* 54, 462-479.
- Baron, J. (1994). Nonconsequentialist decisions. *Behavioral and Brain Sciences* 17, 1-42.
- Baron, J., Gowda, R., and Kunreuther, H. (1993). Attitudes toward managing hazardous waste: What should be cleaned up and who should pay for it? *Risk Analysis*, 183-192.
- Baron, J., and Ritov, I. (1993). Intuitions about penalties and compensation in the context of tort law. *Journal of Risk and Uncertainty* 7.
- Bentham, J. (1982). *An introduction to the principles of morals and legislation* (London, Methuen).
- Blackburn, S. (1993). *Essays in quasi-realism* (New York, Oxford University Press).
- Botvinick, M. M., Braver, T. S., Barch, D. M., Carter, C. S., and Cohen, J. D. (2001). Conflict monitoring and cognitive control. *Psychol Rev* 108, 624-652.
- Bowles, S., and Gintis, H. (2004). The evolution of strong reciprocity: Cooperation in heterogeneous populations. *Theor Popul Biol* 65, 17-28.
- Boyd, R., Gintis, H., Bowles, S., and Richerson, P. J. (2003). The evolution of altruistic punishment. *Proc Natl Acad Sci U S A* 100, 3531-3535.
- Boyd, R. N. (1988). How to be a moral realist. In *Essays on moral realism*, G. Sayre-McCord, ed. (Ithaca, NY, Cornell University Press).
- Carlsmith, K. M., Darley, J. M., and Robinson, P. H. (2002). Why do we punish? Deterrence and just deserts as motives for punishment. *J Pers Soc Psychol* 83, 284-299.
- Chalmers, D. J. (1996). *The conscious mind : In search of a fundamental theory* (New York, Oxford University Press).
- Crisp, R., and Slote, M., eds. (1997). *Virtue ethics* (New York, Oxford University Press).
- de Quervain, D. J., Fischbacher, U., Treyer, V., Schellhammer, M., Schnyder, U., Buck, A., and Fehr, E. (2004). The neural basis of altruistic punishment. *Science* 305, 1254-1258.
- Dutton, D. G., and Aron, A. P. (1974). Some evidence for heightened sexual attraction under conditions of high anxiety. *J Pers Soc Psychol* 30, 510-517.
- Estabrooks, G. H. (1943). *Hypnotism* (New York, E.P. Dutton).
- Fehr, E., and Gächter, S. (2002). Altruistic punishment in humans. *Nature* 415, 137-140.
- Fehr, E., and Rockenbach, B. (2004). Human altruism: Economic, neural, and evolutionary perspectives. *Curr Opin Neurobiol* 14, 784-790.

Fischer, J. M., and Ravizza, M., eds. (1992). *Ethics: Problems and principles* (Fort Worth, TX, Harcourt Brace Jovanovich College Publishers).

Foot, P. (1978). The problem of abortion and the doctrine of double effect. In *Virtues and vices* (Oxford, Blackwell).

Gazzaniga, M. S., and Le Doux, J. E. (1978). *The integrated mind* (New York, Plenum).

Gintis, H. (2000). Strong reciprocity and human sociality. *J Theor Biol* 206, 169-179.

Goel, V., and Dolan, R. J. (2004). Differential involvement of left prefrontal cortex in inductive and deductive reasoning. *Cognition* 93, B109-121.

Greene, J. (2002) The terrible, horrible, no good, very bad truth about morality and what to do about it, Doctoral Thesis, Princeton University.

Greene, J. (2003). From neural 'is' to moral 'ought': What are the moral implications of neuroscientific moral psychology? *Nat Rev Neurosci* 4, 846-849.

Greene, J., and Cohen, J. (2004). For the law, neuroscience changes nothing and everything. *Philos Trans R Soc Lond B Biol Sci* 359, 1775-1785.

Greene, J., and Haidt, J. (2002). How (and where) does moral judgment work? *Trends Cogn Sci* 6, 517-523.

Greene, J. D., Nystrom, L. E., Engell, A. D., Darley, J. M., and Cohen, J. D. (2004). The neural bases of cognitive conflict and control in moral judgment. *Neuron* 44, 389-400.

Greene, J. D., Sommerville, R. B., Nystrom, L. E., Darley, J. M., and Cohen, J. D. (2001). An fmri investigation of emotional engagement in moral judgment. *Science* 293, 2105-2108.

Haidt, J. (2001). The emotional dog and its rational tail: A social intuitionist approach to moral judgment. *Psychological Review* 108, 814-834.

Haidt, J., Bjorklund, F., and Murphy, S. (2000). Moral dumbfounding: When intuition finds no reason. (unpublished manuscript, university of virginia).

Haidt, J., Koller, S. H., and Dias, M. G. (1993). Affect, culture, and morality, or is it wrong to eat your dog? *J Pers Soc Psychol* 65, 613-628.

Hamilton, W. D. (1964). The genetical evolution of social behaviour. *Journal of Theoretical Biology* 7, 1-52.

Harman, G. (1977). *The nature of morality* (New York, Oxford University Press).

Horowitz, T. (1998). Philosophical intuitions and psychological theory. In *Rethinking intuition*, M. DePaul, and W. Ramsey, eds. (Lanham, Maryland, Rowman and Littlefield).

Hume, D. (1978). A treatise of human nature. In, L. A. Selby-Bigge, and P. H. Nidditch, eds. (Oxford, Oxford University Press), pp. xix, 743.

Hursthouse, R. (1999). *On virtue ethics* (New York, Oxford University Press).

Johnston, M. (1995). Dispositional theories of value. In *Meta-ethics*, M. Smith, ed. (Aldershot, Dartmouth).

Kagan, S. (1989). *The limits of morality* (New York, Oxford University Press).

Kagan, S. (1997). *Normative ethics* (Boulder, CO, Westview Press).

Kahneman, D., Schkade, D., and Sunstein, C. R. (1998). Shared outrage and erratic rewards: The psychology of punitive damages. *Journal of Risk and Uncertainty* 16, 49-86.

- Kamm, F. M. (1993). *Morality, mortality vol. 1: Death and whom to save from it* (New York, Oxford University Press).
- Kamm, F. M. (1996). *Morality, mortality vol. II: Rights, duties, and status* (New York, Oxford University Press).
- Kamm, F. M. (1999). Famine ethics: The problem of distance in morality and Singer's ethical theory. In *Singer and his critics* (Oxford, Blackwell).
- Kant, I. (1930). *Lectures on ethics* (Indianapolis/Cambridge, Hackett).
- Kant, I. (1959). *Foundation of the metaphysics of morals* (Indianapolis, Bobbs-Merrill).
- Kant, I. (1983). On a supposed right to lie because of philanthropic concerns. In *(Indianapolis/Cambridge, Hackett)*, pp. 162-166.
- Kant, I. (1993). *Critique of practical reason, 3rd ed.* (Upper Saddle River, NJ, Prentice-Hall).
- Kant, I. (1994). *The metaphysics of morals*. In *Ethical philosophy* (Indianapolis, IA, Hackett).
- Kant, I. (2002). *The philosophy of law: An exposition of the fundamental principles of jurisprudence as the science of right* (Union, NJ, Lawbook Exchange).
- Koechlin, E., Basso, G., Pietrini, P., Panzer, S., and Grafman, J. (1999). The role of the anterior prefrontal cortex in human cognition. *Nature* 399, 148-151.
- Koechlin, E., Ody, C., and Kouneiher, F. (2003). The architecture of cognitive control in the human prefrontal cortex. *Science* 302, 1181-1185.
- Kohlberg, L. (1971). From is to ought: How to commit the naturalistic fallacy and get away with it in the study of moral development. In *Cognitive development and epistemology*, T. Mischel, ed. (New York, Academic Press), pp. 151-235.
- Korsgaard, C. M. (1996a). *Creating the kingdom of ends* (New York, Cambridge University Press).
- Korsgaard, C. M. (1996b). *The sources of normativity* (New York, Cambridge University Press).
- Kripke, S. A. (1980). *Naming and necessity* (Cambridge, MA, Harvard University Press).
- Kuhn, D. (1991). *The skills of argument* (Cambridge, Cambridge University Press).
- Lacey, N. (1988). *State punishment: Political principles and community values* (London; New York, Routledge & Kegan Paul).
- Lieberman, D., Tooby, J., and Cosmides, L. (2003). Does morality have a biological basis? An empirical test of the factors governing moral sentiments relating to incest. *Proc R Soc Lond B Biol Sci* 270, 819-826.
- Maddock, R. J. (1999). The retrosplenial cortex and emotion: New insights from functional neuroimaging of the human brain. *Trends Neurosci* 22, 310-316.
- Maier, N. R. F. (1931). Reasoning in humans: II. The solution of a problem and its appearance in consciousness. *Journal of Comparative Psychology* 12, 181-194.
- McClure, S. M., Laibson, D. I., Loewenstein, G., and Cohen, J. D. (2004). Separate neural systems value immediate and delayed monetary rewards. *Science* 306, 503-507.

McDowell, J. (1988). Values and secondary qualities. In Essays on moral realism, G. Sayre-McCord, ed. (Ithaca, NY, Cornell University Press).

McGinn, C. (1999). Out duties to animals and the poor. In Singer and his critics, D. Jamieson, ed. (Oxford, Blackwell).

Miller, E. K., and Cohen, J. D. (2001). An integrative theory of prefrontal cortex function. *Annu Rev Neurosci* 24, 167-202.

Moore, G. E. (1966). *Principia ethica* (Cambridge, Cambridge University Press).

Nietzsche, F. (1974). *The gay science* (New York, Random House).

Nisbett, R. E., and Wilson, T. D. (1977). Telling more than we can know: Verbal reports on mental processes. *Psychological Review* 84, 231-259.

Paus, T., Zijdenbos, A., Worsley, K., Collins, D. L., Blumenthal, J., Giedd, J. N., Rapoport, J. L., and Evans, A. C. (1999). Structural maturation of neural pathways in children and adolescents: In vivo study. *Science* 283, 1908-1911.

Petrinovich, L., and O'Neill, P. (1996). Influence of wording and framing effects on moral intuitions. *Ethology and Sociobiology* 17, 145-171.

Petrinovich, L., O'Neill, P., and Jorgensen, M. (1993). An empirical study of moral intuitions: Toward an evolutionary ethics. *Journal of Personality and Social Psychology* 64, 467-478.

Phan, K. L., Wager, T., Taylor, S. F., and Liberzon, I. (2002). Functional neuroanatomy of emotion: A meta-analysis of emotion activation studies in pet and fmri. *Neuroimage* 16, 331-348.

Pinker, S. (2002). *The blank slate: The modern denial of human nature* (New York, Viking).

Plato (1987). *The republic* (Penguin Classics).

Putnam, H. (1975). The meaning of 'meaning'. In *Language, mind, and knowledge*, K. Gunderson, ed. (Minneapolis, University of Minnesota Press).

Ramnani, N., and Owen, A. M. (2004). Anterior prefrontal cortex: Insights into function from anatomy and neuroimaging. *Nat Rev Neurosci* 5, 184-194.

Rawls, J. (1971). *A theory of justice* (Cambridge, MA, Harvard University Press).

Rawls, J. (1995). Construction and objectivity. In *Meta-ethics*, M. Smith, ed. (Aldershot, Dartmouth).

Richerson, P. J., and Boyd, R. (2005). *Not by genes alone: How culture transformed human evolution* (Chicago, University of Chicago Press).

Ross, W. D. (1930). *The right and the good* (Oxford, Oxford University Press).

Sanfey, A. G., Rilling, J. K., Aronson, J. A., Nystrom, L. E., and Cohen, J. D. (2003). The neural basis of economic decision-making in the ultimatum game. *Science* 300, 1755-1758.

Saxe, R., Carey, S., and Kanwisher, N. (2004). Understanding other minds: Liking developmental psychology and functional neuroimaging. *Annual Review of Psychology* 55, 87-124.

Schelling, T. C. (1968). The life you save may be your own. In *Problems in public expenditure analysis*, S. B. Chase, ed. (Washington, D.C., Brookings Institute).

Schnall, S., Haidt, J., and Clore, G. (2004). Irrelevant disgust makes moral judgment more severe, for those who listen to their bodies.

Singer, P. (1972). Famine, affluence and morality. *Philosophy and Public Affairs* 1, 229-243.

Sinnott-Armstrong, W. (2004). Moral intuitionism meets empirical psychology.

Small, D. A., and Loewenstein, G. (2003). Helping a victim or helping the victim. *Journal of Risk and Uncertainty* 26, 5-16.

Small, D. A., and Loewenstein, G. (in press). The devil you know: the effects of identifiability on punitiveness. *Journal of Behavioral Decision Making*.

Smith, A. (1976). *The theory of moral sentiments* (Oxford, England, Oxford University Press).

Sober, E., and Wilson, D. S. (1998). *Unto others : The evolution and psychology of unselfish behavior* (Cambridge, Mass., Harvard University Press).

Steinberg, L., and Scott, E. S. (2003). Less guilty by reason of adolescence: Developmental immaturity, diminished responsibility, and the juvenile death penalty. *Am Psychol* 58, 1009-1018.

Stuss, D. T., Alexander, M. P., Lieberman, A., and Levine, H. (1978). An extraordinary form of confabulation. *Neurology* 28, 1166-1172.

Ten, C. L. (1987). *Crime, guilt, and punishment* (Oxford, Clarendon Press).

Thomson, J. J. (1986). *Rights, restitution, and risk : Essays, in moral theory* (Cambridge, Mass., Harvard University Press).

Thomson, J. J. (1990). *The realm of rights* (Cambridge, MA, Harvard University Press).

Trivers, R. L. (1971). The evolution of reciprocal altruism. *Quarterly Review of Biology* 46, 35-57.

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

Unger, P. K. (1996). *Living high and letting die : Our illusion of innocence* (New York, NY, Oxford University Press).

Variety (1989). Tv reviews--network: Everybody's baby.

Weiner, B., Graham, S., and Reyna, C. (1997). An attributional examination of retributive versus utilitarian philosophies of punishment. *Social Justice Research* 10, 431-452.

Wheatley, T., and Haidt, J. (2005). Hypnotically induced disgust makes moral judgments more severe. *Psychological Science*, 16(10), 780-4.

Wiggins, D. (1987). *Needs, values, and truth: Essays in the philosophy of value* (Oxford, Blackwell).

Wilson, T. D. (2002). *Strangers to ourselves: Discovering the adaptive unconscious* (Cambridge, MA, Harvard University Press).

Wrangham, R., and Peterson, D. (1996). *Demonic males: Apes and the origins of human violence* (Boston, Houghton Mifflin).

Wright, R. (1994). *The moral animal: Evolutionary psychology and everyday life* (New York, Pantheon).