

**SENTIMENTALIST PLURALISM:
MORAL PSYCHOLOGY AND PHILOSOPHICAL ETHICS***

Michael B. Gill and Shaun Nichols
University of Arizona

When making moral judgments, people are typically guided by a plurality of moral rules. These rules owe their existence to human emotions but are not simply equivalent to those emotions. And people's moral judgments ought to be guided by a plurality of emotion-based rules.

The view just stated combines three positions on moral judgment: [1] moral sentimentalism, which holds that sentiments play an essential role in moral judgment,¹ [2] descriptive moral pluralism, which holds that commonsense moral judgment is guided by a plurality of moral rules,² and [3] prescriptive moral pluralism, which holds that moral judgment ought to be guided by a plurality of moral rules. In what follows, we will argue for all three positions. We will not present a comprehensive case for these positions nor address many of the arguments philosophers have developed against them. What we will try to show is that recent psychological work supports sentimentalist pluralism in both its descriptive and prescriptive forms.

1. Moral Sentimentalism

Moral judgment is an obvious candidate for psychological investigation, for moral judgment is at least partly a psychological phenomenon.³ We begin by reviewing important recent evidence that indicates that emotions play a crucial role in generating ordinary moral judgments.⁴

Work on the "moral/conventional" distinction provides one source of evidence for the role of emotion in moral judgment (Blair 1995, 1997, Turiel 1983). Drawing inspiration from moral philosophers, developmental psychologists over the last quarter century compared children's responses to "moral" violations, such as unprovoked hitting, to their responses to

“conventional” violations, such as talking in class. Psychologists found that children’s reactions to unprovoked hitting and other “moral” violations differed significantly from their reactions to violations of classroom rules. Children judged hitting to be more seriously wrong than talking in class. Children typically said that hitting would be wrong even if the teacher had no rule against it, but they were less likely to say that about talking in class. And children tended to justify their answers by saying that hitting is wrong because it *hurts* the person but that talking in class is wrong because it’s against the rules. These distinctions may not be altogether surprising. What *is* surprising and illuminating is that emotions apparently play an important role in generating these distinctions. For although psychopaths are apparently normal on standard cognitive and intellectual measures, they have diminished emotional responses to suffering in others, and James Blair has found that psychopaths, and children with psychopathic tendencies, perform atypically on the moral/conventional task (Blair 1995, Blair 1997). This suggests that emotional responsiveness plays an important causal role in generating normal moral judgments.

Further evidence of the importance of emotion in generating the pattern of ordinary moral judgment comes from work on “trolley cases” and other moral dilemmas. In the *bystander* case, a trolley is bound to kill five people on the main track unless it is diverted, but if diverted, it will kill one person on a side track. In the *footbridge* case, a trolley is bound to kill five people unless a man is pushed off a footbridge into the path of the train. Many philosophers have held that it is intuitively permissible to divert the trolley in *bystander* but not to push the man in *footbridge*, and some of these philosophers have also maintained that such cases help reveal the proper set of normative principles by which to live (Thomson 1976). Psychologists have taken this normative thought experiment and turned it into a method for exploring moral judgment (e.g. Petrinovich & O’Neill 1996, Mikhail 2000, Greene et al. 2001). One persistent finding is that lay people draw a moral distinction that matches the philosophical view that it is permissible to divert the trolley in the *bystander* case but impermissible to push the man in the *footbridge* case. And, once again, it seems that emotions are implicated in this distinction. For patients with damage to the ventro-medial pre-frontal cortex, a brain region associated with emotional sensitivity, tend to have abnormal moral reactions. These patients tend to judge that it is permissible not only to divert the train in *bystander* but also to push the man in *footbridge* (Koenig, Young, et al. 2007; see also Valdesolo & DeSteno 2006).

The evidence on psychopaths and patients with ventro-medial damage suggests that emotions play a critical role in normal moral judgment—that without certain emotional responses, a person’s moral judgment will be abnormal or incongruous. In light of this evidence, some prominent theorists have offered pure emotion-based accounts of moral judgment, claiming that there is a certain type of emotion that is identical with a certain type of moral

prohibition (e.g., Blair 1995). It would be a mistake, however, to conclude that emotional response is the sole ingredient in moral judgment—a mistake to conclude that our moral distinctions originate in emotion alone. For the patterns of lay moral judgments probably do not uniformly track any purely emotional responses. For instance, aversive emotional responses to the suffering of others are implicated in judgments about the wrongness of hitting, but many people likely have these same aversive emotional responses to cases of killing in self-defense, harming criminals, and spanking children while still judging it permissible to kill in self-defense, to harm criminals, and to impose corporal punishment (Nichols and Mallon 2006).

What needs to be added to the psychological picture in order to account for the patterns of lay moral judgment are *rules*. The categorizing of moral transgressions is not simply a function of the activation of a certain emotion but rather also depends on internalized rules proscribing certain kinds of actions. The presence of an internally represented rule explains, for instance, why many judge that pushing the man off the footbridge is impermissible but that other forms of killing (such as in cases of self-defense or punishment) are permissible. While people may have similar initial emotional responses to both kinds of killing, they have internalized a rule against one kind of killing but not the other. In addition, by appealing to internalized rules, we can easily explain cross-cultural differences in normative judgments. Some cultures think that polygamy is morally permissible, others regard it as an abomination. The obvious explanation for this difference is that people in the different cultures have internalized different rules about polygamy (see Mallon & Nichols forthcoming).

There is also a more general reason to think moral cognition depends partly on rules. Humans (and other animals) often acquire critical information from a single encounter. If you get shocked once by an electric fence, you typically don't need to repeat the exercise to learn to avoid it. It is a familiar point in cognitive science that this kind of "one-shot" learning is hard to accommodate in prominent models that eschew rules and representations (see Garson 2007). By contrast, rule-based accounts can easily accommodate one-shot learning. Just as adding a single line of code to a computer program can immediately alter the computer's response, so too adding a single rule to a person's mental economy might have such an immediate and enduring effect. And what is crucial for our purposes is that one-shot learning plays a major role in normative judgment as well. Young children have a knack for learning which actions are proscribed, even in nonmoral contexts. Indeed, children's facility for learning rules extends to rules that are completely arbitrary. In one experiment on 4-year olds, the experimenter said, "One day Carol wants to do some painting. Her Mum says if she does some painting she should put her helmet on" (Harris & Nunez 1996, 1581). Although this is obviously an arbitrary rule, children were adept at identifying transgressions and at explaining their choice (Harris & Nunez 1996). This kind of response doesn't

require massive training. Children pick it up almost immediately. Although it has not been studied ecologically, it's likely that much of the knowledge of culturally specific moral rules is also acquired in one-shot learning episodes.⁵ For in moral education, as in other instances of one-shot learning, the context often signals an important learning event. When we tell our children that it's wrong to hit the baby, we ensure that we have their attention, and we transmit the message with gravitas.

While rules are an additional ingredient in moral judgment, they are not entirely independent of the emotions. Rules at the core of commonsense morality—such as rules not to harm others, to care for one's offspring, and to punish the guilty—resonate with our natural emotional repertoire. And work in cultural evolution indicates that rules that resonate with our emotions get differentially preserved through history (Nichols 2004). Once these rules are in place, however, they can lead us to make moral judgments that do not simply track our occurrent emotional states. While emotions help to causally enshrine moral rules, the moral rules do not remain forever yoked to our present-moment feelings.

As we will elaborate more fully in the next section, the moral rules can come into conflict with each other, as in the cases of perceived moral dilemmas. What psychological processes underlie our resolutions of such apparent conflicts? There is no clear answer to that question at present. Undoubtedly many factors contribute (including both occurrent emotional responses and general purpose reasoning), but there does not seem to be any elegant processing account that tells us how these ingredients get weighted to produce outputs. This situation is not unique to moral psychology. Decision making has generally eluded tidy explanations. Hence Fodor's First Law: "the more global ... a cognitive process is, the less anybody understands it. *Very* global processes, like analogical reasoning, aren't understood at all" (Fodor 1983, 107). While Fodor's law might be somewhat hyperbolic, it is clear that we are nowhere near having detailed processing models of the kind of global decision making that is implicated in difficult moral decision-making.

2. Descriptive Moral Pluralism

There is, then, psychological support for the view that ordinary moral judgment depends crucially on rules, and that the contents of those rules are greatly influenced by emotion. In this section, we give a richer characterization of the rules in everyday moral life. We will argue for a pluralist account of the rules implicated in commonsense morality. Note that we are not, in this section, addressing the prescriptive question of whether or how people *ought* to use moral rules, nor the metaphysical question of what properties might underlie moral rules. Our goal in this section is entirely descriptive:

to address how rules actually figure in ordinary moral judgment. This goal contrasts in important ways with the “external” project in normative ethics, which aims to construct a set of rules that *conforms* to the judgments that people make, regardless of whether those rules play a causal role in generating moral judgment (e.g. Thomson 1976). While we have no objection to that project, our focus is on the “internal” project, which attempts to glean the psychological details of rule-based moral judgment.

The most famous morally pluralist philosophical view is W.D. Ross’s (Ross 16-47). Ross’s account combines descriptive, metaphysical, and prescriptive elements. We find it difficult to disentangle his prescriptive recommendations from his metaphysical commitments, and we reject his metaphysical commitments. Specifically, we see no reason to believe in the mind-independent moral properties—nor in our ability to access those properties—that Ross’s intuitionism involves. We also believe that there are some important moral rules that Ross did not identify. But we think that some of Ross’s central pluralist claims are consonant with everyday moral experience and gain support from recent work in psychology and cognitive science. These claims are:

- A plurality of different basic rules contribute to moral judgment.
- Some of these rules are agent-relative.
- The basic rules implicated by commonsense morality can conflict with each other.
- There is no invariable ranking of the basic moral rules.

We will now say a bit to defend each of these claims.

2.1. A plurality of different basic rules contribute to moral judgment

The rule (or set of rules) against harming innocents occupies a central spot in commonsense morality. Elsewhere, one of us has argued that our natural emotional reactions to harming others confer a special force on the rule against harming others, and this connection with emotion likely played a critical role in the cultural success of the rule (Nichols 2004). In this sense, the rule against harming others causally depends on emotions—if we didn’t have those emotions, we likely wouldn’t have the rule. But we come equipped with a variety of different emotions, and it’s plausible that many of the different commonsense moral rules causally depend on these different emotions. Rules about making reparations depend on guilt. Rules about punishment and desert depend on anger. Rules about special obligations to one’s children depend on parental love. Rules about sexual deviance likely depend on disgust. Although each case requires its own defense, we think it’s plausible that for many of these rules the same structure holds: the rules have had cultural success because they resonate with our natural

emotional repertoire, and they continue to have considerable psychological force because we continue to have those emotions.⁶

This picture of rules and emotion gains support from Rozin and colleagues (1999). Building on work by Shweder and colleagues (1997), Rozin and colleagues maintain that three different emotions—contempt, anger, and disgust—align with three different moral codes. Contempt is connected with community violations, such as neglecting community-based duties. Anger is connected with violations of autonomy, such as infringing on an individual's rights. Disgust is connected with violations of purity or sanctity, such as behaving in ways that are religiously forbidden. In addition, recent work by Haidt & Joseph (2004), suggests that rules concerning respect for elders depend on awe or fear, that rules concerning group loyalty depend on group pride and belongingness, and that rules of reciprocity depend on gratitude. We find the work of Shweder, Rozin and Haidt fascinating, but we needn't cleave closely to the details. Our point is simply that this work supports the view that people deploy a plurality of different moral rules when forming moral judgments, and that some of these different rules depend on different emotions.⁷

Moral pluralism also maintains that more than one of the rules implicated in commonsense morality are fundamental or basic. In the philosophical literature, this feature of moral pluralism is often put in terms of the mutual irreducibility or underivability of the moral rules, and in terms of the absence of any single moral measure against which all of the rules can be placed (see Chang 1997). We want to emphasize the psychological aspect of this feature of moral pluralism—the idea that there is a plurality of moral rules that each has its own independent psychological basis. One way to think about this is that if one of these psychologically basic moral rules—along with the emotions that helped sustain it—were completely eliminated, it wouldn't just grow back from the other rules and emotions, nor is there any more basic underlying norm from which all the moral rules could be generated. Different rules have different emotional boosters and constitute different vectors of moral thinking.

This psychological pluralism also provides a way to accommodate recent evidence on subjects' justifications for responses to trolley dilemmas. Most subjects fail to adequately justify why they respond differently to cases like footbridge and bystander, and Hauser and colleagues suggest that the problem might be that we lack conscious access to the underlying justifying principle (e.g., the doctrine of double effect) (Hauser et al. 2007). But if there are independent basic rules, there might not *be* an underlying principle that justifies both responses. People justify their response to *bystander* by pointing out that you save more people, and they justify their response to *footbridge* by saying that you can't intentionally kill someone. Such justifications might well be accurate about the principles that guide their judgments, for in different cases different principles might be driving the judgment.

2.2. *Some rules are agent-relative*

Some of the commonsense moral rules are *agent-relative* (see Ridge 2005). Agent-relative rules command actions that cannot be fully specified without reference to the particular agent who performs them. This is in contrast to agent-neutral rules, which command actions that can be fully specified without reference to the agent who performs them. Some commonsense rules do seem to be agent-neutral—the rule to minimize human suffering, for instance. But many rules are agent-relative. For instance, the rule against neglecting one’s children is agent-relative in that it forbids a parent from neglecting his own child. It is not a command to minimize the neglect of children, regardless of whose children they happen to be. If I neglect my child so I can bring it about that two other children are not neglected, I have still violated the prohibition on neglecting one’s children. In addition to obligations to one’s children, several other commonsense rules are clearly agent-relative, such as rules concerning reparations, lying, and killing.

2.3. *The basic rules implicated by commonsense morality can conflict with each other*

In some situations, the only way to act in accord with one rule will be to violate another. What are commonly taken to be moral dilemmas illustrate this kind of conflict (Sinnott-Armstrong 1988). In Sartre’s famous story, for instance, a student struggles with the decision of whether to tend to his ailing mother or join the resistance. Here we have a conflict between two moral considerations—loyalty to one’s family and duty to country. And the relevant pluralist claim is that commonsense moral thought does at times have to deal with this sort of conflict between different psychologically basic considerations.

One way to try to dissolve this sort of conflict is to add to the moral rules fine-grained “exceptive clauses” that reduce the scope of each rule so that conflict between them becomes impossible (Richardson 1990, Shafer-Landau 1995). It is doubtful, however, that the rules, as represented in the minds of most ordinary humans, have built into them such conflict-dissolving clauses.⁸ More likely is that the moral rules that actually play a role in ordinary persons’ moral judgments have a kind of generality that does make it possible for them to come into conflict. Evidence for this comes from responses to another trolley case. In *catastrophe*, pushing a large stranger off a footbridge in front of an oncoming train will not merely save the lives of five innocent people but will prevent catastrophically bad consequences from occurring. Most subjects said that it was permissible to push the stranger to prevent catastrophically bad consequences and *also* that pushing the stranger

in front of the train violated a moral rule (Nichols and Mallon 2006). People thus seem to think that an action can be morally permissible despite its violating a moral rule (e.g., the rule not to kill an innocent person). The most plausible diagnosis here is that the commonsense rule against killing innocents can conflict with and be overridden by a different rule to minimize suffering.

2.4. There is no invariable ranking of the basic moral rules

In particular situations people can and do rank competing moral considerations. If a person is in a situation in which she can prevent an injury only by lying, she will typically decide that one particular moral consideration (either preventing this particular injury or refraining from this particular lie) is morally more important than the other. What pluralism as a descriptive claim denies is that commonsense morality involves any overarching principle or uber-rule that ranks the rules in an exceptionless ordering.

This pluralist view of resolving moral conflict occupies a middle ground on the issue of how we form difficult moral judgments. On one side is extreme moral monism, which holds that moral judgments require only decisions about how to apply a single rule to morally fraught situations.⁹ On the other side is extreme moral particularism, which holds that moral judgments do not involve the application of general rules but only the assessment of, or sensitivity to, the particulars of morally fraught situations. Pluralism holds that ordinary moral judgment does involve the application of general rules: when making moral judgments, people take into consideration whether an act accords (or conflicts) with certain moral rules and they generally take the fact that an act accords (or conflicts) with certain moral rules to be a significant moral reason in its favor (or disfavor).¹⁰ Because the rules are multiple, basic, and potentially conflicting, however, the application of these rules cannot be all that is involved in moral judgment. Also involved must be some case-by-case prioritizing of these rules, some way of coming to a decision about which rules in particular situations are normatively more powerful than the other rules with which they conflict.

The basic features of ethical pluralism we've so far sketched are all in evidence in the "principles" approach that has dominated medical ethics for several decades. As developed by Beauchamp and Childress, the principles approach holds that decisions in medical ethics ought to be guided by four distinct (clusters of) principles: autonomy, beneficence, nonmaleficence, and justice (Beauchamp and Childress 2001). Other medical ethicists have accepted the basic structure of Beauchamp and Childress's approach but argued that we ought to use a somewhat different set of principles. Such differences are not our concern here. The point we want to call attention to

is that adherents of the principles approach in medical ethics (regardless of the specific sets of principles they advocate) have generally held that there is a plurality of basic principles, that some of these principles are agent-relative, that these principles can come into conflict with each other, and that there is no invariable ranking of these basic principles. In addition, one of us (Gill) has served on medical ethics committees in community hospitals for ten years, and his experience has suggested that the principles approach is popular not merely with academics writing about medical ethics but also with those who are actually engaged with ethically fraught situations in medical settings. The basic features of ethical pluralism seem to characterize a great deal of on-the-ground thinking in actual clinical medical ethics. The principles approach has been criticized by numerous people working in medical ethics, but these criticisms have mainly been normative, not descriptive. Critics have argued that it would be better if the principles approach were not so popular, that it would be better if people working in medical ethics did not use moral principles in the pluralist way we've been describing. But for our purposes in this section, the popularity of the principles approach in medical ethics (independent of its desirability or undesirability) is itself significant, as it constitutes additional evidence that the moral judgments people are inclined to make have the pluralist character we've been describing.

Let us close by underscoring that the pluralist view we've sketched in this section is offered as a description of how people form moral judgments in actual cases. We acknowledge, though, that some ordinary people may have beliefs about morality that imply that there is one supreme moral principle, or that every moral rule can be specified in a way that ensures that moral rules never come into conflict, or that moral rules ought to be expurgated from our decision-making altogether. These beliefs would be aspects of the theories such people hold about morality, and perhaps such beliefs can lead some of these people to form their moral judgments in non-standard (i.e., monist, specificationist, or particularist) ways. Our view, however, is that the actual practice of how most ordinary people form moral judgments in particular cases involves moral rules in a pluralist manner. And this view could be accurate even if it does not harmonize with the theoretical views some people hold about morality, nor with the way some people sometimes form moral judgments.

3. Prescriptive Moral Pluralism

In section 1, we argued that commonsense morality depends critically on the emotions. The pattern of commonsense moral judgments and the moral rules that guide them are greatly influenced by emotional responses. Without the emotions, we would have nothing much like commonsense morality. In section 2 we argued that commonsense morality is pluralistic—that ordinary

moral judgments depend on a plurality of psychologically basic rules. We now want to propose that the sentimentalist account of moral judgment combined with the pluralism of commonsense morality provides a presumptive case in favor of the prescriptive claim that our morality *ought* to include a plurality of basic rules. We are not trying to make a prescriptive case for all of the moral judgments and moral rules that are taken to be commonsensical. For all we say here, many particular commonsense judgments—and some of the rules on which those judgments are based—may be indefensible. Our goal is only to defend the general pluralist structure of commonsense morality.

One quick way to try to make the prescriptive case for pluralism is to deploy the maxim *ought-implies-can*. If our psychological constitution is such that we cannot avoid basing our moral judgments on a plurality of rules, then (if *ought* implies *can*) it cannot be the case that we ought not to base our moral judgments on a plurality of rules. This way of arguing for pluralism fails, however. For the influence of rules on our moral judgments may be something we can consciously override. Even if we have pre-reflective pluralistic tendencies, through reflection we may (as monists prescribe) be able to bring all of our moral judgments into conformity with the dictates of a single overarching rule; or we may (as particularists prescribe) be able to disregard rule-based considerations in the formation of our moral judgments.

Although there is no quick and decisive argument for pluralism, we think that the sentimental basis of moral judgment presents us with a new way to think about the available options in normative ethics. And once we follow these options through, pluralism will come out looking very attractive.

The empirical work suggests that our moral worldview has been largely shaped by our emotional repertoire. Thus, we start with the question: *are the emotions a proper ground for normative ethics?* In trying to determine the ethically right thing to do, is it appropriate to rely on rules that depend on the emotions?

The most prominent rejection of emotion-based morality comes from rationalists who maintain that reason alone is the only proper ground for ultimate moral principles, where a principle is “ultimate” if it enshrines something that is valued for its own sake. Among rationalists, there are different views about how reason delivers these principles. Intuitionist versions hold that we intuit them through some form of a priori reason (Ross 14). Kantian versions hold that we arrive at the principles through pure practical reason (Kant 3).

The evidence of section 1 suggests, however, that a critical core of commonsense moral judgment implicates rules that are based on emotions. The evidence suggests that if the influence of emotion is eliminated, one’s pattern of moral judgment will be incongruous or bizarre to commonsense—one’s pattern of moral judgment will look to be not merely a refinement or revision of commonsense morality but a very different thing altogether. Now

some normative rationalists may be happy to embrace results that seem incongruous and bizarre to commonsense. They may be sanguine about prescribing a wholesale revamping of our patterns of moral judgments. As it happens, however, most intuitionists and Kantians have *not* endorsed moral views that differ very much from commonsense morality (Ross 21; Kant 16). Most intuitionists and Kantians have produced theories that are fairly similar to the pattern of commonsense moral judgment. (Indeed, even Peter Singer who says he is quite willing to give up all commonsense ethical intuitions (2005, 345–6) still wants to retain the fundamental judgment that “it is a bad thing if a person is killed” (350).) The evidence from section 1 thus challenges these commonsense intuitionists and Kantians to justify the claim that their prescriptive results ultimately derive from reason alone. For a hypothesis that at least initially fits better with the evidence is that the supposedly purely rational moral judgments that intuitionists and Kantians (and even Singer) rely on to generate their moral views seem obviously valid to us only because we accept certain moral rules which are themselves emotion-based. That is, it’s likely that many of the central principles would not seem self-evident were it not for the influence of emotion (Nichols 2008). Ironically, the people with the best claim to having intact reasoning capacities but fewer emotional biases are psychopathic.

Thus, while there is certainly an important philosophical position according to which emotions are not proper grounds for ultimate moral principles, the available evidence suggests that this position will involve a radical revision of commonsense ethics. It will have to divest itself of the huge chunk of commonsense moral judgments that draw on the resources of emotion-based moral rules. So if we are right about the role of emotions in commonsense morality, then rationalists face a dilemma: either give up the claim that reason alone is the only proper ultimate ground of moral judgment, or give up the bulk of commonsense morality. And this point generalizes: all those who deny that emotions can be a proper ultimate ground of morality will need to abandon more of commonsense morality than almost anyone is willing to do.

Let’s turn to the view that emotions *are* a proper ground for morality. One version of this view holds that emotions are a proper ground because they track emotion-independent moral truths. Shaftesbury, for instance, seemed to think that goodness was a property that existed independently of humans’ emotional responses but also that humans had a moral sense that led them to have positive emotional responses to that which was good.¹¹ On this Shaftesburean view, that the moral sense leads us to have good-tracking emotions is no coincidence but rather is explained by the fact that God, who is perfectly good, implanted our moral emotions in us. There could be an evolutionary version of this position as well. On this kind of theory, just as the best evolutionary explanation of our having the fear responses we do is that those responses are sensitive to real (emotion-independent) dangers,

so too the best explanation of our having the moral emotions we do is that they are effective at tracking real moral truths that exist independently of our emotions.

We have serious doubts about the tenability of such tracking sentimental accounts. The main problem is that these views seem to require that we have *two* ways of getting at moral truth—through emotions *and* through some emotion-independent means, the latter being necessary to confirm the former. But what might these emotion-independent means be? Shaftesburean-type views require a purely rational apprehension of moral truth. We have already sketched above our doubts that pure rationality will deliver anything close to (largely emotion-based) commonsense morality. The evolutionary view seems to involve non-emotion-based appeals to some kind of regular, substantive connection between morality and evolutionary success. But there are well known worries about linking moral propriety with evolutionary advantage (see e.g. Kitcher 1994; Singer 2005). We are thus pessimistic about the prospects for tracking sentimentalism.¹²

Let us now consider the final option, that the emotions are a proper ground for morality but do not track any emotion-independent moral truths. We'll call this view "normative sentimentalism" to distinguish it from the *psychological* sentimentalism (which is a purely descriptive thesis) that we defended in section 1. Normative sentimentalism is, in broad outline, the view put forward by the classical sentimentalists Francis Hutcheson and David Hume.¹³ Hutcheson and Hume's fundamental claim was that moral justification has to bottom out in sentimentally-grounded ultimate ends, in ends we have that cannot be justified by anything else (such as pure rationality) but are themselves the starting points for moral justification. As Hume explains, if you ask someone why he exercises he may say that he does so to avoid illness. If you ask why he wants to avoid illness he may say that he does so because illness is painful. But if you ask why he wants to avoid what is painful you will get no answer. Avoiding pain is an ultimate end of his (Hume 293). Similarly, according to the classical sentimentalists, if you ask someone why she has made a particular moral judgment (say, *that A ought not be punished*) she may give you a reason (such as, *A did not perform action X, or action X which A performed did not cause any harm*). And she may be able to give further reasons for those reasons. Eventually, however, you will reach a point beyond which there will be no further justificatory answer. That point, according to classical sentimentalism, will be a sentimentally-grounded ultimate end, an end that is not justified by anything else but rather serves as a starting point for justification. One such end, for Hume, could be "benevolence," which grows out of the pleasure we take in what is "useful" to others (Hume 218). Another such end could be "Good Manners or Politeness," which grows out of the pleasure we take in what is "immediately agreeable to others" (where agreeability is clearly distinct from usefulness) (Hume 261). Yet another could be honesty, which

grows out of the pleasure we take in “conversation” and “society” (Hume 283).

On this classical sentimentalist view, there is an important analogy between morality and beauty (Gill 2007). Judgments of beauty are normative judgments, and they are, at least often, grounded in emotions; the aesthetic judgments we make are partly a function of the emotions we have. Most would agree, however, that emotions are a perfectly appropriate ground for aesthetic judgments. We don’t recant our aesthetic judgments upon discovering that they depend on the sentiments. And as in the aesthetic domain, the classical sentimentalist held, emotions are a proper ground for judgment in the ethical domain as well.

This classical sentimentalist position is well known, but we think the recent empirical work on the role of emotions in moral judgment gives us even more reason to adopt it. The empirical work suggests that if we were to restrict ourselves only to ultimate moral principles that could be derived from reason alone, we would be saddled with normative consequences virtually no one is willing to accept. And this provides additional heft to the classical sentimentalist considerations in favor of the idea that the emotions themselves should be accepted as a proper ultimate ground for morality.¹⁴

So if emotions are taken to be a proper ultimate ground for morality (the classical sentimentalist position, which gains support from the considerations of section 1), and if we have a plurality of psychologically basic emotion-grounded rules (as we maintained in section 2), then normative ethics should take as its starting point the plurality of emotion-backed rules that lie at the core of commonsense morality. Commonsense may be a dubious starting point for physics and metaphysics. But if normative sentimentalism is true, normative ethics should not be thought of as being in the same business as physics or metaphysics. If normative sentimentalism is true, morality should be taken to begin, at least partly, from sentimentally-based ends, and the rules of commonsense morality are reflections of those ends.

The combination of morality’s sentimentalist origins and our possessing a plurality of morally-significant sentimentally-based ends is nowhere near a decisive argument for pluralism. We are claiming only that it makes a presumptive case in favor of pluralism. The central question thus becomes: are there successful arguments against that presumption?

Attempts to overturn the presumption in favor of pluralism come from two sides: from particularists who hold that moral rules should not be used when forming moral judgments, and from monists who hold that only one ultimate principle should be used.¹⁵ We will not address the particularist challenge here.¹⁶ Our focus instead will be on arguments purporting to show the superiority of monism over pluralism.

Monists might try to overcome the presumption in favor of pluralism by arguing that in fact we have just one ultimate end after all. One might take Aristotle to be attempting such a thing when he maintains that eudaimonia

is the end that is “most sovereign,” the end that alone “we wish for because of itself” (Aristotle 2002, 95). But Aristotle’s view is only apparently monist, not really so. Eudaimonia is a complex entity constituted by a plurality of ultimate ends, and Aristotle was aware of that. To be eudaimon is to succeed at the design problem of bringing into as much harmony as possible the multiplicity of things we value for their own sake. Now it’s true that some other philosophers have said things that might be taken to suggest that we have only one substantive ultimate end. Bentham can at times be read as arguing that we ultimately value only pleasure.¹⁷ Hobbes at times can be read as arguing that we ultimately value only self-preservation.¹⁸ Berkeley says, “Sensual pleasure is the Summum Bonum.”¹⁹ As descriptive claims about what people actually place ultimate value on, however, these arguments seem highly unlikely to succeed, as the considerations in section 2 should help to show. But this monistic strategy can succeed only if it is true that we have only one ultimate end.

Some might think, however, that commonsense morality is normatively committed to monism because pluralism involves an unacceptable arbitrariness (see Seung and Bonevac 1992). Pluralist theories tell us which factors to consider when making moral judgments, but they cannot adjudicate between those factors when they pull in opposite directions. Monist theories, in contrast, hold out the promise of justificatory finality—of filling the justificatory gap between rules and judgments. If normative sentimentalism is true, however, reason alone does not ensure that all the things we take to be of fundamental moral importance will always harmonize with each other, or that if one moral rule overrides another in one case that the former must override the latter in every case. Now it may be that we take justificatory finality to be a kind of goal; we may seek the most determinate justification we can come up with for our moral judgments. But it’s far from clear that monistic theories themselves provide complete and total justificatory finality; the application of a single moral rule can often involve its own indeterminacy (see Hill 1992). Moreover, pluralist theories don’t leave moral justification in a completely arbitrary state; the plurality of rules almost always significantly narrows the morally legitimate choices and may in many cases uniquely select one (see Wolf 1992). It’s also far from clear that we value justificatory finality enough to make it reasonable to restrict our moral theorizing to monist theories. If we have a true plurality of ultimate ends, then the best approach to normative theorizing may be to leave open the possibility of pluralism rather than restrict ourselves only to monistic options.

We don’t mean to suggest that we have shown that the pluralist aspects of ordinary morality will necessarily be preserved in the best normative theory. The best normative theory will be what we arrive at when we achieve wide reflective equilibrium, and it’s possible (for all we’ve said) that such a view will be monistic. But the starting points for this justificatory process will include a plurality of basic rules. Some might hold that monism should be

taken to be a fixed point in moral theorizing, something that we should take to be non-negotiable from the start. The fact that ordinary moral judgment is pluralist, however, suggests that we shouldn't take monism to be a fixed point.²⁰

A final way for those who adopt a normative sentimentalist view to challenge the presumption in favor of a plurality of different moral rules is to maintain that one emotion (or perhaps one kind of emotion) is better than all the others. Recently, Joshua Greene has presented some extremely interesting arguments of this sort in an attempt to establish the superiority of utilitarian moral theory. If Greene is right, all of the ultimate deontological principles—all the principles that imply that we have an ultimate reason to perform acts that fail to produce the best utilitarian outcomes (such as the rule prohibiting killing innocents)—should be extirpated, depleting commonsense morality of a large chunk of its rules.²¹ But Greene doesn't think his view's distance from commonsense is a problem. He thinks a large chunk of commonsense morality—namely, the deontological chunk—ought to be jettisoned. We will argue in the remainder of this paper, however, that the reasons for giving up the deontological rules are not compelling. Greene's arguments don't suffice to undercut a pluralist account that includes both utilitarian and deontological rules.

Greene sometimes suggests that utilitarian judgment is superior because it is inherently cognitive, whereas deontological judgment is emotional (Greene 2008, 64). In some domains, we certainly do privilege the cognitive over the emotional. We prefer the accountant who does calculations with cold reason rather than under emotional influence. We think the scientist should analyze data with detached calculations rather than affective bias. But our preference for cognitive processes in these domains plausibly derives from the fact that accountants and scientists are attempting to track some emotion-independent way things are. This certainly doesn't imply that we ought to endorse more cognitively-activated thinking about issues that do *not* track emotion-independent truths. That one judgment “tends to be driven by emotion” while another judgment “tends to be driven by ‘cognitive’ processes” (Greene 2008, 40) may be a reason to favor the former over the latter when we are attempting to accurately represent a cold fact about the universe. But once we are in the realm of the emotion-dependent, it is unclear that more cognitive judgments are normatively superior. Of course, emotions sometimes distort our judgments. For instance, subjects who were led to feel angry by a movie clip subsequently attributed more blame to a negligent person in an unrelated scenario (Lerner et al. 1998). In such cases, anger corrupts our better judgment. But the fact that emotions sometimes impair normative judgment doesn't show that they always do so. Again the analogy with aesthetics is instructive. If we have to choose which of two items is more beautiful, our more cognitive reactions might tend toward one item while our more emotional reactions might tend toward the other. It is far

from obvious, however, that the cognitive reactions should be privileged over the emotional ones. For we don't take our judgments about beauty to be concerned exclusively with the tracking of some emotion-independent states of affairs. A correct judgment about what is beautiful is a function in part of our reactions, and some of our emotional reactions may be just as legitimate in this regard as our cognitive ones.

The most direct argument Greene gives against emotion-based deontological rules comes in his response to a commentary by Mark Timmons (2008). Greene writes, "Timmons asks, why can't deontologists embrace the emotive foundations of their judgments? The answer... is GIGO [garbage in, garbage out]. Kant was opposed to emotion-based morality because emotions are fickle and contingent in oh-so-many ways.... About that, he was right" (2008, 117).

But what makes emotion-based input garbage? The above passage says that it's because "emotions are fickle and contingent." However, it's not just deontological rules that causally implicate the emotions. Consider a utilitarian rule like "Minimize human suffering." Why do we disvalue others' suffering at all? Emotions provide the best explanation (contra Nagel 1986). If we had different emotions (if what most thrilled us was torturing children), the utilitarian inputs would be very different (cf. Greene 2008, 64). So if we are to reject deontological considerations because they depend critically on the emotions, then we are similarly bound to reject familiar utilitarian considerations.

At one point, Greene suggests that what makes deontology inferior is the *kind* of emotion that drives deontological judgment. He says that the kind of emotion involved in utilitarian judgments is "currency-like." These emotions say "Such-and-such matters this much. Factor it in." In contrast, the emotions hypothesized to drive deontological judgment are far less subtle. They are, as I have said, alarm signals that issue simple commands: 'Don't do it!' or 'Must do it!'" (Greene 64). But even if the emotions involved in utilitarian judgments are "currency-like," that doesn't show that utilitarian rules are normatively superior to deontological ones. It's true that our emotion-backed deontological (or agent-relative) rules can't be converted into utilitarian currency. But it would beg the question to say that these emotion-backed rules are inferior simply because they can't be converted into utilitarian currency. So the currency-like character of the utilitarian-emotion doesn't provide a direct argument for its normative superiority.

Perhaps, then, the idea is not that the utilitarian emotion is uniquely good, but rather that the emotions involved in deontological judgment are especially *bad*. But why? Presumably because these emotions are "alarm-like." As we've emphasized throughout, however, emotions—alarm-like or otherwise—are only part of the story for deontological judgment. Rules also play a crucial role. And while alarm-like emotion might be fickle

and contingent, the rules needn't be. Once in place, the rules are not so brittle that they disappear when the emotions are absent. Nor are the rules spontaneously created simply by activating an emotion.²² The rules are an independent vector in moral thought. As a result, the fact that emotions are fickle doesn't show that the emotion-backed *rules* are fickle. Once we recognize this, we can see that it's even possible that the very same emotion (alarm-like or otherwise) played a causal role in establishing both utilitarian and deontological rules. In particular, it's likely that compassion played a critical causal role in the cultural success of both the deontological prohibition on harming and the utilitarian rule to minimize suffering. To be sure, it remains a most interesting fact that a given emotion can be causally responsible for both utilitarian and deontological rules. But insofar as the same emotion is causally implicated in both cases, we don't yet see why the role of emotions impugns one type of rule but not the other.

In addition, the deontological rules do not constitute barriers to careful consideration of myriad objective factors that bear on the situation. Applying a deontological rule (such as the rule to punish and reward fairly) may involve a great deal of objective fact-finding (such as trying to determine which acts a person has performed and how others who have performed similar acts have been treated). The effects of deontological rules on moral judgment cannot be broadly assimilated to unreflective, knee-jerk, impulsive reactions. Furthermore, on the pluralist view we are promoting, the deontological rules themselves don't rigidly dictate all-in moral judgment. Even after one identifies that a deontological rule has been violated, one still needs to determine whether the action was, all-things-considered, permissible, and this might involve a consideration of other rules, both deontological and utilitarian.

Our aim in this essay has been to elucidate some of the meta-ethical and normative implications of empirical work in moral psychology. There is now converging evidence that commonsense moral judgment depends critically on our emotional repertoire. There is also good reason to think that commonsense moral judgment implicates a plurality of psychologically basic rules that act as independent vectors in moral thought. This doesn't have any simple implications for philosophical ethics, but we think that the most initially promising response to these findings is sentimentalist pluralism, according to which the plurality of emotion-backed commonsense moral rules provides an appropriate starting point for normative ethics.

Notes

- * We are grateful for helpful comments from Julia Annas, Mark Collier, John Doris, Jerry Gaus, Terry Horgan, Rachana Kamtekar, Walter Sinnott-Armstrong, Tamler Sommers, and Mark Timmons.

1. This is a *psychological* moral sentimentalism, according to which the emotions have played (and continue to play) an important causal role in moral judgment. *Philosophical* sentimentalism is often associated with a stronger conceptual claim that morality is necessarily constituted by the emotions. We do not rely on this stronger claim.
2. Our claim is that moral rules are internalized as mental representations which causally contribute to judgment and behavior. In some cases rules might be consciously accessible, but in other cases rules might operate outside of conscious awareness. Apart from those qualifications, we construe “rules” in a permissive way, so that even the dictates of etiquette count as rules. Thus in the moral domain, we mean “rules” to include various types of moral imperatives, such as the Ten Commandments, Rossian prima facie duties (Ross, 16-47). Hursthouse’s v-rules (Hursthouse 1999), and Beauchamp and Childress’ principles of biomedical ethics (Beauchamp and Childress 2001). For discussion of moral rules as “guides,” see McKeever and Ridge 2006, 7-9.
3. By “moral judgment” we mean to refer to a very wide category that could include moral distinctions, verdicts, beliefs, and assessments.
4. Allman and Woodward (this volume) adduce additional evidence for the role of emotion in intuitions that factor into moral judgment.
5. Some theorists argue that morality is not “learned” at all, but grows from an innate moral grammar (see Mikhail 2007). We are skeptical about such views (see Nichols 2005), but moral grammar nativists would basically agree with our point here—that children acquire culturally local rules with little training.
6. Although we maintain that many central moral rules causally depend on emotions, we do not claim that *all* moral rules depend on emotions in this way.
7. Of course, the rules figure into judgment in different ways. Some moral judgments result from consideration of one rule, other moral judgments result from consideration of a different rule, and still other moral judgments result from consideration of multiple rules.
8. Nor do Richardson and Shafer-Landau present their “specificationist” views as being primarily descriptive; Richardson presents specificationism largely as a prescriptive claim, and Shafer-Landau presents it largely as a metaphysical claim.
9. The monist view that moral judgment requires only the application of a single moral rule should not be equated with the implausibly simplistic view that moral judgment is entirely algorithmic. Monists can and do hold that the application of a rule can be a very complex non-algorithmic process that requires a great deal of practical wisdom.
10. Moral rules do not have to be exceptionless or absolute; they can still be rules even if they have ceteris paribus clauses or are default principles. For discussion, see McKeever and Ridge, 113-137 and Horgan and Timmons (forthcoming).
11. For discussion of this Shaftesburean view, see Gill 2006, chapters 7-9.
12. For discussion of these issues, see the Street 2006 and Copp (this volume).
13. For discussion of the relevant aspects of the classical sentimentalist view, see Gill 2006, chapters 13 and 15.
14. We recognize that some aspects of commonsense morality may presuppose an objectivity that is incompatible with the sentimentalist account we offer. This might lead some to embrace an *error theory* according to which all commonsense

- moral judgments are false (e.g., Mackie 1977). In our view, it remains quite unclear whether error theory is the appropriate response to sentimentalism (see, e.g., Gill 2008; Nichols 2004). But for present purposes this is a side issue, for even error theorists typically endorse principles that are very much like *moral* principles (e.g. Mackie 1977, 105-6; cf Blackburn 1985).
15. Monists can allow the use of multiple subsidiary principles so long as they are all reducible or lexically inferior to the single ultimate principle.
 16. For a comprehensive critique of particularism with which we are in substantial agreement, see McKeever and Ridge 2006.
 17. See Bentham 11. Bentham seemed to believe that only a rigorously monistic theory would enable us to produce principled, systematic ethical prescriptions. But actually, even Bentham acknowledged that we had as ultimate ends both pleasure and the absence of pain. And it's not clear that these *two* ultimate ends will always line up perfectly. A pleasure-and-absence-of-pain pluralism may be where even a seeming monist like Bentham ends up, although Bentham himself might not have been pleased with this result. See Sinnott-Armstrong 77.
 18. See Hobbes 1991, 115 [*De Cive* I.7]. In other places, however, Hobbes says that what is good for a person is what he desires (see, e.g., *Leviathan* 6.7), and this may cohere with the idea that we have more than one ultimate end, insofar as it's possible that we can desire for its own sake more than just our self-preservation.
 19. For discussion of Berkeley's view, see Kail 2007, 198.
 20. That Kant proceeds as though monism is a fixed point is evident from his stating at the outset of the *Groundwork* that his goal is to seek out and establish "the *supreme principle of morality*" (Kant 193) and his claim that "there is only a single categorical imperative" (Kant 31). That Mill also at least sometimes has this monistic restriction in mind is evident from the section of his *Logic* entitled "Necessity of an ultimate standard, of first principle of Teleology," in which he maintains that there must be one and only one ultimate rule of conduct (Mill 951). In contrast to these approaches—which build monism into moral theorizing from the start—is the approach of someone like Kagan (Kagan 1989). Kagan doesn't start from the assumption that the best moral theory will have only one kind of rule. He starts by taking as a serious moral contender a view that includes both consequentialist and deontological rules, and he then goes on to argue that the demands of reflective equilibrium lead us to a view that has only one (consequentialist) rule.
 21. Although Greene typically pits *consequentialism* against deontology, the consequentialism he has in mind is clearly utilitarian rather than egoist or altruist. To better focus our discussion, we will speak more narrowly of utilitarianism.
 22. In the wonderful study by Wheatley & Haidt (2005), hypnotically-induced disgust did change subjects' morality ratings for an action against which there is no rule; however, on a scale from 1 ("not at all morally wrong") to 100 ("extremely morally wrong"), the mean rating was still only 14 (p. 783).

References

- Allman J and Woodward J. 2008. "What are Moral Intuitions and Why Should We Care About Them: a Neurobiological Perspective," *Philosophical Issues*, 18, 164–185.

- Aristotle. 2002. *Nicomachean Ethics*. Translated by Broadie and Rowe. Oxford University Press.
- Beauchamp, T. and Childress, J. 2001. *Principles of Biomedical Ethics*. Oxford University Press.
- Bentham, J. 1996. *An Introduction to the Principles of Morals and Legislation*. Oxford University Press.
- Blackburn, S. 1985. "Errors and the Phenomenology of Value." In T. Honderich (ed.), *Morality and Objectivity*. London: Routledge & Kegan Paul.
- Blair, R. 1995. "A Cognitive Developmental Approach to Morality: Investigating the Psychopath," *Cognition*, 57, 1–29.
- Blair, R. 1997. "Moral Reasoning and the Child with Psychopathic Tendencies," *Personality and Individual Differences*, 26, 731–739.
- Blair, R., L. Jones, F. Clark, M. Smith and L. Jones 1997. "The Psychopathic Individual: A Lack of Responsiveness to Distress Cues?" *Psychophysiology*, 34, 192–198.
- Copp, D. 2008. "Darwinian Skepticism About Moral Realism," *Philosophical Issues*, 18, 186–206.
- Chang, R. 1997 (editor). *Incommensurability, Incomparability, and Practical Reason*. Cambridge, MA: Harvard University Press.
- Fodor, J. 1983. *Modularity of Mind*. MIT Press.
- Garson, James. "Connectionism," *The Stanford Encyclopedia of Philosophy (Spring 2007 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/archives/spr2007/entries/connectionism/>.
- Gill, M. 2006. *The British Moralists on Human Nature and the Birth of Secular Ethics*. Cambridge University Press.
- Gill, M. 2007. "Moral Rationalism vs. Moral Sentimentalism: Is Morality More Like Math or Beauty?" *Philosophy Compass* 2: 16–30.
- Gill, M. 2008. "Meta-ethical Variability, Incoherence, and Error." In W. Sinnott-Armstrong (ed.) *Moral Psychology*, Vol. 2. Cambridge, MA: MIT Press.
- Greene, J. 2008. "The secret joke of Kant's soul." In W. Sinnott-Armstrong, ed., *Moral Psychology*, Vol. 3. MIT Press.
- Greene, J., Sommerville, R., Nystrom, L., Darley, J., & Cohen, J. 2001. "An fMRI investigation of emotional engagement in moral judgment," *Science*, 293, 2105–08.
- Haidt, J., & Joseph, C. 2004. "Intuitive ethics: How innately prepared intuitions generate culturally variable virtues." *Daedalus*, Fall, 55–66
- Harris, P. and M. Núñez 1996. "Understanding of Permission Rules by Preschool Children," *Child Development*, 67, 1572–1591.
- Hauser, M., Cushman, F., Young, L., Jin, R., & Mikhail, J. 2007. "A dissociation between moral judgment and justifications," *Mind & Language*, 22, 1–21.
- Hill, T. 1992. "Kantian Pluralism," *Ethics* 102: 743–762.
- Hobbes, T. 1991. *Man and Citizen*, edited by Bernard Gert. Indianapolis: Hackett Publishing.
- Horgan, T and Timmons, M. Forthcoming. "What Does the Frame Problem Tell Us about Moral Normativity?" *Ethical Theory and Moral Practice*.
- Hume D. 1975. *Enquiries concerning Human Understanding and concerning the Principles of Morals*. Edited by Selby-Bigge, revised by Nidditch. Oxford University Press.
- Hursthouse, R. 1999. *On Virtue Ethics*. Oxford University Press.
- Kagan, S. 1989. *The Limits of Morality*. Oxford University Press.
- Kail, P. 2007. *Projection and Realism in Hume's Philosophy*. Oxford University Press.
- Kant, I. 1997. *Groundwork of the Metaphysics of Morals*, edited by Mary Gregor. Cambridge University Press.
- Kitcher, P. 1994. "Four Ways of Biologizing Ethics." In *Conceptual Issues in Evolutionary Biology*, edited by E. Sober. Cambridge, MA: MIT Press.
- Koenigs, M., Young, L., Adolphs, R., Tranel, D., Cushman, F., Hauser, M., Damasio, A. 2007. "Damage to the prefrontal cortex increases utilitarian moral judgments," *Nature*, 446, 908–911.

- Lerner, J.S., Goldberg, J.H., & Tetlock, P.E. 1998. "Sober second thought: The effects of accountability, anger and authoritarianism on attributions of responsibility," *Personality and Social Psychology Bulletin*, 24, 563–574.
- Mackie, J. 1977. *Inventing Ethics*. New York: Penguin.
- Mallon, R. & Nichols, S. forthcoming. "Rules and moral psychology." In J. Doris *et al.* (eds.) *Handbook of Moral Psychology*. Oxford: Oxford University Press.
- McKeever, J and Ridge, M. 2006. *Principled Ethics: Generalism as a Regulative Ideal*. Oxford University Press.
- Mikhail, J. 2000. *Rawls' Linguistic Analogy*. Unpublished Ph.D. thesis, Cornell University.
- Mikhail, J. 2007. "Universal moral grammar," *Trends in Cognitive Sciences*, 11, 143–152.
- Mill, JS 1974. *A System of Logic Ratiocinative and Inductive*. Edited by JM Robsen. University of Toronto Press and London: Routledge & Kegan Paul.
- Nagel, T. 1986. *The View from Nowhere*. Oxford University Press.
- Nichols, S. 2004. *Sentimental Rules*. New York: Oxford University Press.
- Nichols, S. 2005. "Innateness and Moral Psychology." In *The Innate Mind: Structure and Content*, eds. P. Carruthers, S. Laurence, and S. Stich. New York: Oxford University Press.
- Nichols, S. 2008. "Moral Rationalism and Empirical Immunity." In W. Sinnott-Armstrong (ed.) *Moral Psychology*, Vol. 2. Cambridge, MA: MIT Press.
- Nichols, S. and Mallon, R. 2006. "Moral Dilemmas and Moral Rules," *Cognition*, 100, 530–542.
- Petrinovich, Lewis, and O'Neill, Patricia. 1996. "Influence of Wording and Framing Effects on Moral Intuitions," *Ethology and Sociobiology* 17, 145–171.
- Richardson, H. 1990. "Specifying Norms as a Way to Resolve Concrete Ethical Problems," *Philosophy and Public Affairs* 19: 279–310.
- Ridge, M. 2005. "Reasons for Action: Agent-Neutral vs. Agent-Relative," *The Stanford Encyclopedia of Philosophy (Spring 2007 Edition)*, Edward N. Zalta (ed.), URL = <http://plato.stanford.edu/entries/reasons-agent/>.
- Ross, W. 1930. *The Right and the Good*. Oxford: Clarendon Press.
- Rozin, P., Lowry, L., Imada, S., Haidt, J. 1999. "The CAD Triad Hypothesis," *Journal of Personality and Social Psychology*, 76, 574–586.
- Seung TK and Bonevac D 1992. "Plural Values and Indeterminate Rankings," *Ethics* 102: 799–813.
- Shafer-Landau, R 1995. "Specifying Absolute Rights," *Arizona Law Review* 37: 209–225.
- Shweder, R., Much, N., Mahapatra, M. & Park, L. 1999. "The "Big Three" of morality (autonomy, community, divinity) and the "Big Three" explanations of suffering." In A. Brandt & P. Rozin (Eds.), *Morality and health* (pp. 119–169). New York: Routledge.
- Singer, P. 2005. Ethics and Intuitions. *Journal of Ethics*, 9, 331–352.
- Sinnott-Armstrong, W. 1988. *Moral Dilemmas*. Blackwell.
- Street, S. 2006. "A Darwinian Dilemma for Realism Theories of Value," *Philosophical Studies* 127: 109–66.
- Thomson, J. 1976. "Killing, letting die, and the trolley problem," *The Monist*, 59, 204–217.
- Timmons, M. 2008. "Towards a Sentimentalist Deontology," In W. Sinnott-Armstrong (ed.) *Moral Psychology*, Vol. 3. Cambridge, MA: MIT Press.
- Turiel, E. 1983. *The Development of Social Knowledge*. Cambridge: Cambridge University Press.
- Valdesolo, P., and DeSteno, D. 2006. "Manipulations of Emotional Context Shape Moral Judgment," *Psychological Science* 17: 476–77.
- Wolf, S. 1992. "Two Levels of Pluralism," *Ethics* 102: 785–798.
- Wheatley, T., and Haidt, J. 2005. "Hypnotically Induced Disgust Makes Moral Judgments More Severe," *Psychological Science* 16: 780–84.