

The Normativity Challenge: Why the Empirical Reality of Traits Will Not Save Virtue Ethics

Jesse Prinz

ABSTRACT: Situationists argue that Virtue Ethics empirically untenable, since traditional virtue ethicists postulate broad, efficacious character traits, and social psychology suggests that such traits do not exist. I argue that prominent philosophical replies to this challenge do not succeed. There is, however, empirical reason to postulate character traits, and this undermines the situationist critique. There is, however, another empirical challenge to virtue ethics that is harder to escape. Character traits are culturally informed, as are our ideals of what traits are virtuous, and our ideals of what qualifies as well-being. If virtues and well-being are culturally constructed ideals, then the standard strategy for grounding the normativity of virtue ethics in human nature is undermined.

KEYWORDS: Situationism, Virtue Ethics, Personality, Self Concepts, Well-Being, Culture, Normativity

In recent years, moral philosophers have gotten wind of an old story in social psychology. Human behavior is often influenced by subtle situational factors, and such influences are often considerably better predictors of what we do than character traits as measured by standard personality inventories. Character seems to be causally inert, and circumstance pulls the strings. Some philosophers, most influentially Gilbert Harman and John Doris, have recently argued that these findings undermine certain version of virtue ethics. Contemporary virtue ethics derive from Aristotle and other ancient philosophers who argued that morality principally involves the cultivation of noble character traits. If traits are on shaky empirical footing, these theories may be in trouble. Virtue ethicists have responded to the Situationist Challenge, and the ensuing debate has become something of a cottage industry. In this discussion, I hope to contribute to that industry in four ways. First, I will claim that some recent responses to the Situationist Challenge are not decisive, because they underestimate the extent to which social psychology raises doubts about trait *efficacy*. Second, I will argue that there are decisive responses to the Situationist Challenge; there is evidence for efficacious character traits. But, I will then introduce another empirical objection virtue ethics, pertaining to culture, well-being, and the normative status of virtues. Finally, I will argue that, when traits are properly understood, they do not mitigate the force of this second challenge, but rather exacerbate it. I will not suggest that virtue ethics has been refuted; rather, I will claim that any defensible version of virtue ethics may not be a genuine *alternative* to rule-based ethical theories. Moreover, rules may be more fundamental than virtues. In sum, even if Harman and Doris are wrong to be skeptical about traits, a proper understanding of traits raises some serious doubts about the place of virtue in moral theory.

1. The Situationist Challenge

The impact of situational factors has been well publicized in psychology since the 1930s, at least, when Muzafer Sherif demonstrated social effects on perceptual judgment. Solomon Asch made related discoveries in the 1950s, and, in the 1960s, Stanley Milgram conducted his infamous obedience studies. In the 1970s, Philip Zimbardo, Richard Nisbett, Alice Isen, John Darley, and others authored influential studies that extended these early findings. By this time, social psychologists were taking direct aim at the construct of character. They were directly comparing situational influences to character traits, and concluding that traits had little impact. Lee Ross, building on work by Edward Jones and Keith Davis, argued that the attribution of character traits was a fundamental error (a point B.F. Skinner has made two decades earlier). Defenders of situationist psychology began to see character traits as an illusion.

During the same time period, philosophers were re-discovering virtue ethics. Beginning with Elizabeth Anscombe's "Modern Moral Philosophy" in 1958, some philosophers began to worry that leading approaches in contemporary ethics focused too much on rules and too little on character—too much on what to do, and not enough on what sort of person to be. In the following decades, influential defenses of virtue were written by Philippa Foot, Bernard Williams, Alistair MacIntyre, among others. By the 1990s, virtue ethics was no longer *avant garde*; it was mainstream. College courses were taught on virtue ethics, and publishers issued a steady stream of monographs and anthologies. The emerging enthusiasm for virtue in philosophy was bound to clash with the emerging skepticism about character in psychology. The 1990s was also a time period in which philosophers were increasingly cognizant of the social sciences, and many were beginning to notice that philosophical theories lie in *prima facie* conflict with conclusions that some psychologists regard as axiomatic. Owen Flanagan (1991) published a seminal discussion of this conflict, and papers by Doris (1998) and Harman (1999) brought the issue into the center of debate by arguing that prevailing versions of virtue ethics were empirically untenable. Since then, it's been a publishing frenzy. Virtue theorists think that Harman and Doris misuse empirical results or misunderstand Aristotelian ethics. Doris, in his (2002) monograph, Harman, Peter Vranas, Maria Merritt and others remain undaunted, and continue to marshal psychological findings against certain philosophical theories.

The challenge raised by Harman and Doris can be summarized as follows. On some versions of virtue ethics, especially those that are inspired by Aristotle, morality is presumed to be concerned with the cultivation of character traits that have certain properties. Doris says that Aristotelians postulate "global" character traits, which are presumed to be consistent across a wide range of conditions, stable over repeated trials, and integrated with evaluatively similar traits. Harman says Aristotelians postulate traits that are "broad-based," meaning they pertain to a wide range of behavioral contexts; being talkative is broad-based, while being talkative in the cafeteria at lunch is not. Broadness is like Doris's notion of consistency. In addition, Doris and Harman think that Aristotelians presume that such traits are causally efficacious; Doris talks of "empirical content" and Harman says these traits are presumed to play a role in explaining behavior. Merritt (2000) sharpens this idea by introducing the term Motivational Self-Sufficiency; she says that Aristotelians think traits are sufficient for motivation independent of factors outside the self. Doris and Harman argue that social psychology raises empirical

objections to global efficacious character traits. Experiments show that efficacious traits are not global, and allegedly global traits are not efficacious.

The experiments in question are familiar enough not to review here, but I will mention four that have been especially prominent in recent discussions. First, there is Milgram's (1974) obedience study, in which every single subject administered (what they believed to be) highly dangerous electrical shocks to another individual when instructed to do so by the experimenter. Second, there is Darley and Bateson's (1973) study in which Princeton Theological Seminary students en route to deliver a lecture encountered a person who seemed to be in need of help; subjects tended to help if they were not running late for the lecture (63%), but not if they were running late (10%). Third, there is Isen and Levin's (1974) study in which subjects who found a dime in a phone booth were likely to help a passerby who dropped some papers, while subjects who did not find a dime were unlikely to help. Finally, there is Hartshorne and May's study of honesty in which eleven thousand children were assessed on 33 behaviors in which honesty can be measured, and the authors found only a low correlation (about 0.2) across behaviors.

The first three studies show that a relatively minor situation manipulation with no obvious moral significance exerts a major influence on people's moral behavior, and, importantly, no personality variable seems to exert such influence. These findings are not anachronistic; there are many studies with similar results (for a review, see Ross and Nisbett, 1991). This suggests, that people who might seem to have global character traits, as measured by standard personality inventories, do not act on those traits; otherwise, behavior would vary as a function of traits, rather than situation. The Hartshorne and May study shows that individuals behave in different ways across contexts in which they should behave similarly if they were acting under the influence of global character traits. This suggests that, if people have traits, they may be narrower than the traits postulated by virtue theorists; a person who doesn't steal may nevertheless cheat, which undermines the application of broad labels such as "honest" or "dishonest." Together, the studies suggest that character traits are either narrow and efficacious or broad and inert. Either way, the conception of traits favored by Aristotelian virtue ethics finds little empirical confirmation in these studies. Harman and Doris invite us to conclude that such traits lack psychological plausibility. Global efficacious traits do not seem to exist. Aristotelians believe that such traits can be cultivated, but situationist psychology raises doubts about that. If such traits do not seem to drive behavior, why think that minds like ours could come to be governed by them if we simply had the right moral education? Strictly speaking, social psychology does not prove that such traits *couldn't* exist, but they shift the burden of proof. In the absence of any evidence for global efficacious traits, Aristotelians must show that minds like ours could come to possess or approximate that which we do currently have. This is an initial statement of the Situationist Challenge, and, of course, Aristotelians have as much to say in response.

2. Assessing the Situationist Challenge

2.1 Inconclusive Objections

Ever since Harman and Doris published their formulations of the Situationist Challenge, fans of virtue ethics have been lining up to respond. The rapidly growing literature is

already too large to summarize responsively. What I offer here is a selective sampling of some of the recurring themes. I will not attempt to list every response, and I will not do justice to the responses that I do discuss. Rather, I will simply offer caricatures of five response strategies that are, in my view, unpromising.

Reply 1. Personality psychologists have proven that character traits exist. This is a point that is not frequently pursued in detail by philosophers, but it features centrally in psychological discussions of situationism. These days, it is widely agreed that certain traits are quite robust. They are broad-based, culturally universal, and relatively stable across the lifespan (John and Srivastava, 1999). There are five personality dimensions have been studied most extensively. Each of these dimensions is a continuum along which personality can vary. These dimensions emerge with great regularity and consistency in study after study. They are so experimentally robust that they have become known as the Big Five. The first dimension is *openness*, which involves imagination, insight, and receptivity to new experiences. The second dimension is *conscientiousness*, which is a measure of how thorough and organized a person is. Third comes *extroversion*; some people are very outgoing and some have a more inward, or introverted, orientation. The fourth dimension is *agreeableness*, which encompasses a person's capacity to have affection and sympathy for others. And finally, there is *neuroticism*, which is a measure of anxiety and emotional instability. Psychology 101 students are taught the acronym OCEAN to remember these five dimensions. Some theories define or label them in different ways, but there are a number of standard personality tests that are used to assess how each of us comes out with respect to the Big Five. Personality profiles, including clinical syndromes such as depression, anxiety, and psychopathy, have been conceptualized as point in the Big Five personality space—locations in the OCEAN. Given the overwhelming empirical support for this model, it is completely misleading to claim that character traits do not exist.

For personality psychologists, research on the Big Five has done much to deflect the radical conclusions of situationists. If one were to read Mischel (1968) or Ross and Nisbett (1991) one might think that the field of personality psychology is a sham, driven entirely by our tendency to make the fundamental attribution error in both folk psychology and scientific psychology. Work on the Big Five has shown that personality is a legitimate subject of scientific study. Nevertheless, I don't think this research program offers any comfort to defenders of virtue ethics. Here are some reasons for doubt.

First, there continues to be considerable debate about the efficacy of Big Five traits in influencing behavior. Most of the research on the Big Five uses questionnaires rather than behavioral measures. When behavior is tested, Big Five traits may have limited impact. For example, Newcombe (1929) shows that there is very little cross-situational behavioral consistency in extroversion, and Mischel and Peake (1982) showed that there is also very limited behavioral consistency for conscientiousness. Mischel and Peake did show higher behavioral consistency for people who regarded themselves as consistently conscientious, but, importantly, many subjects claimed to be highly variable on this basic dimension of personality. Questionnaire methods not only fail to adequately test for behavioral effects, they may also fail to measure self-assessment of behavioral consistency, because they ask subjects about the degree to which they have certain traits

rather than the frequency with which such traits are manifest. I might regard myself as an extrovert even if I am only extraverted in certain social setting. In a recent unpublished paper, English and Griffith (2003) showed that subjects do not generally regard themselves as having highly consistent Big Five traits.

Second, it's far from clear whether the Big Five traits are morally significant; none of them look like items on a standard list of virtues. "Conscientiousness" sounds like it might be a virtuous trait, but the name is misleading. Conscientious people are organized and tidy, not necessarily governed by the dictates of conscience. Agreeableness has somewhat better claim to being a virtue. It is said to encompass friendliness and sympathy, traits that have been traditionally important in the virtue ethics of Aristotle and Hume, respectively. But the actual questions used to measure agreeableness are only vaguely related to traits emphasized in such ethical theories. Essentially, people are rated as agreeable if they have a positive regard for others rather than a tendency to be spiteful. That's an appealing trait, but it may fall short of being virtuous. In any case, agreeableness is the only one of the Big Five dimensions that even resembles traits that moral philosophers care about.

Third, it is important to virtue ethics that character traits be amenable to cultivation of a certain kind. We build character through moral education, and strive to improve ourselves. The Big Five personality traits are difficult to cultivate. They are estimated to be between 42% and 57% heritable (Jang et al. 1996), which means that they probably owe a great deal to our genes. That is already a concern for Aristolian ethics, because no trait can be regarded as more natural than others. The remaining variance that is not genetic may owe to organic factors, such as health, diet, climate, as well as life experiences. Alternation through life experience suggests that personality is mutable, but studies of the environmental factors that impact personality tend to suggest that traits do not change in the way that virtue theory requires. For example, personality can be influenced by gender and age. Women tend to score higher on neuroticism, which may indicate that women are socialized in a way that tends to erode confidence (Costa et al. 2001). As we grow older, we may become less neurotic and extraverted, but conscientious (McCrae et al., 2000). This might be driven by the fact that we acquire more responsibilities as we age (hence more conscientiousness), we become more practiced in our abilities (hence less neurotic), and we have less need to build social networks (hence less extroversion). Studies also suggest that personality can be affected in the early years of life, through, for example, the parenting styles of our caregivers (Roelofs et al. 2006). None of these biographical effects on personality looks like Aristotle's suggestions about moral education. We do not alter personality by seeking to better ourselves or the people we care for. Rather personality adapts in an involuntary way to life circumstances.

Fourth, the Big Five traits are too rudimentary to vindicate virtue ethics. They are probably constituted by fairly simple affective dispositions. Neuroticism may involve dispositional fear, openness may involve dispositional curiosity, agreeableness may involve dispositional social affection (liking people), extroversion may involve dispositional social joy (liking attention), and conscientiousness may involve the disposition to be agitated by disorder. There are researchers who think that these traits are implemented by rudimentary brain systems, perhaps distinguished by varying levels or neurotransmitters (Plomin and Caspi, 1999). As we will see below, defenders of virtue

theory conceive the virtues as rational capacities or practical reasoning skills. They are not mere emotional dispositions, but rather action-guiding ways of thinking about the world. Research on the Big Five offers little support for the existence of traits like that.

Reply 2. Experiments conducted by Situationists are confounded because they introduce factors that can rationally override virtuous traits. Virtues are not supposed to be automatic or inflexible response dispositions. Whether a virtue guides behavior depends on two cognitive factors at least: whether an agent construes a situation as germane to the virtue in question and, if so, whether there are other facts that weigh against behaving in accordance with that virtue. For example, if a courageous person is mugged, she may not construe it as a situation in which courage matters. Why even consider putting up a fight if all one has to lose is a few dollars and a cheap watch? And if she considers exercising her courage in this situation, she might have a change of heart when she realizes that the risks are too high. In order to test for courage, it wouldn't necessarily be a good idea to examine how people behave when they are mugged, because other factors may rationally override acts of resistance under those circumstances. Likewise, critics complain that standard experiments in the situationist literature introduce factors that may rationally override behavior that accords with virtue. This point has been pressed by Sreenivasan (2002).

To see if this objection can be sustained, let's consider two of the classic experiments. First, consider Milgram's obedience study. A virtuous person might be expected to refrain from administering an electric shock to a stranger, but, in Milgram's experiment, all people administer dangerously high shocks and about 65 percent administer shocks of the highest possible voltage (labeled XXX). But of course none of these subjects would do that under certain other circumstances. For example, Milgram found considerably lower levels of compliance when subjects were being instructed to administer shocks by a student rather than a scientist. Obedience seems to depend on the presence of an authority figure. If a scientist instructs you to shock someone, you assume the scientist understands the risks involved and would not ask you to put someone in serious danger. You also assume that your participation in the study may have some benefit. So, although it's a difficult moral decision, you decide that there is good reason to comply. Now consider Darley and Bateson's seminary study. If you see someone in need, virtue demands that you offer assistance, but suppose that you are running late to give a lecture. If you are late, it will be rude to the instructors who asked you to give that lecture, and so you decide not to help. Like the Milgram case, one norm (the obligation to comply with legitimate authorities) is pitted against another norm (concern for people in distress). When two norms are pitted against each other, we are forced to violate one of them, and, in these experiments, subjects may decide that it is rational not to aid a person in agony.

I find this defense deeply troubling. If subjects allow their concern for authority to trump their concern for a person in pain, they are making a grievous moral mistake. The desire to follow orders and the desire to be on time may be admirable in other contexts, but here the salient distress of another human being should trump. Sreenivasan implies that there is reason to abandon virtue in these cases, but I think he mistakes rationalization for reason. Is it really reasonable to leave a moaning stranger slouched in doorway simply because you are in a hurry? What appointment could be so important?

Satisfying your teacher's request is not an excuse for abandoning a person who may depend on you and you alone. In any case, any seminary teacher would understand if a student was delayed for such a worthy cause. Things are equally clear in the Milgram case. In other versions of the experiment, Milgram found that subjects comply even if the "scientist" is working for corporate interests in a seedy office building. He also found that subjects comply when they learn that the victim of their shocks has a serious heart condition. By the end of the experiment, they have reason to think the victim has been seriously hurt or even killed, yet they continue to up the voltage. It's important to remember that Milgram's study was meant to shed light on why ordinary people participate in incredibly atrocities, such as the holocaust. The suggestion that it is rational to obey authorities even when they ask you to do terrible things misses the point of these studies. We may delude ourselves into thinking it is rational, but that it is because the situational variables overwhelm our capacity to make decent decisions.

Reply 3. Situational influences are highly circumscribed. I don't think Sreenivasan's reply saves virtue theory. Virtues seem to be impotent under conditions where they should be applied. Sabini and Silver (2005) offer a related reply to Doris and Harman, which is slightly more plausible, but ultimately unsuccessful. They do not suggest that subjects regard their behavior as rational, but they point out that all these experiments involve influences of a certain kind. In the classic situationists studies, subjects cave in under social pressures. Harman and Doris sometimes imply that everything we do is driven by an untold variety of external factors that operate outside our awareness. Sabini and Silver say that this is an exaggeration. We are not always driven by external factors; we are just driven by one special kinds of external factor under special circumstances. If they are right, situational effects may be sufficiently restricted to leave plenty of room for the exercise of virtue.

I have three responses. First, if social pressures were the only factors that could override virtues, that would still be cause for alarm. Social pressures are ubiquitous. We are always exposed to the demands, opinions, and expectations of others. It's hard to imagine a condition under which the attitude of a social group, whether real or imagined, would not come readily to mind when making a morally relevant decision. Second, social pressures are not monolithic. We bow under authority, we conform to the majority, and we loaf when we think other individuals can do our work for us. Each of these social influences obeys a different logic, and, consequently, we should not assume that situational influences can be easily circumscribed. Third, Sabini and Silver are too quick to dismiss cases of situational influence that are not social in nature. Recall, Isen and Levin's phone booth study: subjects tend to help a passerby pick up papers if they find a dime just before, and otherwise they tend no to help. A measly dime has an enormous effect. Isen explains this by saying that positive emotions promote prosocial behavior. She has also shown that induction of positive affect influences risk taking (Isen and Patrick, 1983). Notice that these effects depend on emotion induction rather than social influence. Emotions are swayed by minor factors on a regular basis: weather, diet, sleep, sound, and smell can all influence how we feel. These influences are ubiquitous and they are difficult to prevent. Thus, there is little reason to think that situational influences on behavior are unusual as Sabini and Silver suggest.

In response to the phone booth study, Sabini and Silver say that helping someone pick up papers is not a very significant act, from a moral perspective, and being in a bad mood may be a legitimate excuse not to help. But this reply is insufficient. First of all, common human decency is just the kind of thing virtue ethicists tend to care about. Morality isn't always about saving lives; it's also about being kind to people around us. Second, the reply fails to reckon with the fact that emotional influences can be quite powerful. For example, Scandinavian suicide rates are linked to the emotional effects of weather and darkness, and violent road rage is often associated with minor provocation when compounded by the stresses of driving. Emotion can also influence our moral judgments. Schnall et. al (2008) asked subjects filled out a questionnaire assessing the wrongness of various scenarios (e.g., eating your pet dog after killing it accidentally). Some subjects were seated at a clean desk, and others were seated at a dirty desk. Those at the dirty desk gave higher wrongness ratings for the very same scenarios. It seems that our values shift with mild fluctuations of mood. This can have dramatic effects when exploited by savvy pundits and politicians. Consider how leaders muster support for extremist policies by promoting feelings of patriotism or using vague innuendos to promote climate of xenophobia, suspicion, and fear. In Serbia and Rwanda murderous hatred was fuelled by leaders who described certain ethnic groups as insects or animals. These cases may be extreme, but the political values that guide our own lives may be shaped by a mixture of social inculcation and emotion induction.

Reply 4. The fact that most people are swayed by situational factors is no surprise to virtue ethicist, because they have always claimed that it's hard to be virtuous. In Hellenistic philosophy, it was often suggested that virtue is almost impossible to achieve. Only true sages, who dedicate their lives to character development, possess true virtue. On this conception of virtue ethics, the rest of us can just try in our feeble ways to emulate these noble masters. Therefore, we should not expect many people in a randomly chosen sample to do what virtue demands. Virtuous behavior is rare.

This point has been made by several commentators (Annas, 2005; Athanassoulis, 2000; Sreenivasan, 2002), but I think it has two serious shortcomings. First, underestimates the ramifications of situationist psychology. The claim that virtue is rare presupposes that people can become virtuous. And that presupposes that behavior can be driven by character traits. But why should we think this is true? What evidence is there for thinking that minds like ours are capable of being driven by inner traits? It would be perfectly reasonable to say that virtue is rare if most people were driven by character traits that were not virtuous. But situationist psychology purports to show that people are not ordinarily driven by character traits at all. In the face of this experimental evidence, it is seriously question begging to assume that human beings, who don't ordinarily possess efficacious character traits, can come to possess such traits, whether virtuous or not. If situationists are right about human psychology, the acquisition of virtuous traits is not merely the acquisition of nobler versions of the minds we currently possess, but rather the acquisition of minds of an entirely different kind. Become virtues would require a whole cognitive architecture. There is absolutely no reason to think that moral education could give us new mental machinery. If situationists are right about average minds, then virtuous minds are not merely hard to attain; they may be nomologically impossible.

Second, those who insist that virtue is very hard to achieve may be raising the bar too high for a moral theory. In contemporary moral philosophy, virtue ethicists often pride themselves on having theories that are less demanding than consequentialist and Kantian ethical theories. Consequentialists suggest that we should choose our actions by performing cold utility calculations rather than reflecting on our personal relationships with other human beings. Kantians think we should obey the law-like dictates of reason, even if the net result is great suffering. Opponents of these views point out that we wouldn't want our friends to reason like Kant or Mill (Williams, 1981; Wolf, 1982). Virtue ethics is sometimes seen as a remedy for this problem, because virtue is not supposed to require super-human methods of practical reasoning. If contemporary virtue theorists return to the demanding Hellenistic conception of virtue, according to which only the sage can be moral, then virtue theory will lose much of its interest. Morality takes care and vigilance, but a moral theory that demands powers that are out of reach for most of us cannot serve as guide to daily life. Such a morality is not what we're really after in trying to lead good lives.

Reply 5. Situationists experiments threaten a notion of character that has nothing to do with the notion defended by virtue ethicists. Annas (2005) and Kamtekar (2004) accuse Harman and Doris of assuming that character traits are supposed to be behavioral dispositions that function like habits, or unthinking automatic responses. On Aristotle's conception, traits are more cognitive than that. Having a trait involves possession of the ability to engage in certain kinds of practical reasoning. Traits are rational capacities, not reflexes.

On the face of it, situationist psychology has little bearing on the existence of such capacities. Situationists seem to reason as follows. If character traits exist, they would be like automatic reflexes. If behavior were controlled by reflexes, then they should be automatically triggered under certain conditions, regardless of subtle variations in situational context. Reflexes are robust across a wide range of conditions, and they are certainly not affected by reasoning. If we see someone being tortured, and we have a justice-reflex, we should intervene no matter what. Milgram's obedience study shows that we do not behave this way. Situational context can lead to dramatic differences in how we behave. If we believe that an authority figure wants us to cause harm to another person, we comply. In general, situationist experiments suggest that behavior is not controlled by reflex-like traits.

Virtue ethics has no stake in the existence of reflex-like traits, so experiments that cast doubt on such traits leave virtue ethics unscathed. Virtues are practical reasoning capacities, and reasoning can be flexible and context sensitive. When we reason, we take facts about the current situation into consideration. Perhaps the subjects in Milgram's experiment lack a justice-reflex, but some of them may have a capacity to reason carefully about justice, and they simply decide that, in the context of the experiment, it would be inappropriate to aid the victim of their electrical shocks.

I don't see this response is decisive. It's important to remember that subjects in situationist experiments often engage in deliberation—they exercise their rational capacities. Thus, if character traits are practical reasoning skills, then they should be available when the subjects decide what to do. Suppose that caring about justice involves a subtle understanding of what actions are unjust and an appreciation of the conditions

under which intervention in the name of justice are called for. If such a concern were an enduring trait of character for some individuals, we should expect it to be a factor in their deliberations during the Milgram experiment. Yet, none of the subjects in the initial version of that experiment behave in accordance with that disposition. There are three possible explanations. Either, no one has that disposition, or that disposition is present but rationally overruled by other beliefs, or the disposition is present but swamped in a way that renders it effectively inert. None, of these options help save virtue theory. The first two options correspond to replies 3 and 4, which I just considered and rejected. The third option is, in my view, the most likely. A lot of people have genuine concern about justice, and are able to reason about the injustice of torture under certain circumstances, as when they are watching the news. But, when people are put into a situation where they are asked to serve as torturers by an authority figure, their ability to reason clearly about justice is seriously impaired. Zimbardo's infamous prison experiments have a similar moral: when people are endowed with the authority to torture, they often do, even if they would condemn torture when they reflect on it in the abstract. This was sadly reconfirmed in Abu Ghraib. The experiments in situationist psychology suggest that *reasoning*, not just behavioral reflexes, is heavily influenced by situational variables. People who have perfectly noble reasoning dispositions are incapable of deploying them effectively when certain situational pressures arise. This is the most chilling and important moral of situationist psychology.

2.2 Evidence for Traits

I am inclined to reject each of the foregoing strategies for replying to Harman and Doris. None of these worries demonstrates that behavior can be guided by efficacious, morally significant global traits. Without proving that efficacious global traits are possible, defenders of virtue ethics fail to adequately answer the Situationist Challenge. The experiments suggest that widespread situational variables exert a strong influence on practical reasoning and behavior. In the presence of such variables, individual differences don't seem to matter very much. People who come out differently on personality tests end up acting the same way. On the face of it, then, the Situationist Challenge continues to threaten virtue ethics. In this section, I want to consider some direct evidence for efficacious global traits. The evidence I will review goes some way towards quelling the situationist attack, but, as we will see, virtue theory may still be in trouble.

Eliminativism about efficacious global character traits is an extremely radical view. It is radically opposed to the conception we have of ourselves in Western folk psychology. Situationists want us to accept that folk psychology is grossly mistaken; they refer to many of our trait-based explanations as a "fundamental attribution error." But, on brief reflection, there are obvious examples of traits that have an impact on behavior. Let me begin with some mundane examples. Suppose you like modern jazz. That personal characteristic will certainly affect what you do. Jazz enthusiasts buy certain music, listen to certain radio stations, and attend certain concerts. Or consider the attribute of being a "foody" or an "art lover." Foodies are more likely than others to seek out good restaurants, and art lovers are more likely to visit museums. Personal tastes affect behavior. So do self-conceptions. If you see yourself as a hippy, you might wear

certain clothing and advocate particular causes. Your preferences would be different if you defined yourself as a redneck. If you see yourself as gay, or as a Jew, or as preppy, you might adopt patterns of behavior that associated with each of those categories. We often live in accordance with prescribed social roles, and, when we do, we often see these roles as expressions of our identity—they are part of our character, as that term is ordinarily used.

Self-conceptions influence a wide range of behaviors over a wide range of circumstances. They are global and efficacious. They are often highly integrated as well. If you are a hippy, you probably drive a non-ostentatious foreign car, have an environmentalist sticker on your bumper, have a Grateful Dead cassette in your glove compartment, and shop at an organic supermarket. We see these behaviors as a coherent package. Thus, self-conceptions have many of the properties that virtues are supposed to have.

These examples are anecdotal, of course, but they are consistent with finding from the lab. It's well established in social psychology that self-conceptions are efficacious. In one classic demonstration of this, Miller et al. (1975) compared two methods of increasing tidiness among inner city school children. With one group of children, they tried persuasion, urging them not to litter. With a second group, they used an attribution technique, telling the children that they were good at cleaning up after themselves, and giving them an award for tidiness. In the short term after the intervention, both groups became more tidy, littering less and cleaning up after others, but the attribution group showed a much greater increase than the persuasion group. After two weeks, the persuasion group was beginning to return to its earlier pattern of littering, but the attribution group remained tidy. The children who were (falsely) labeled as tidy were more than five times less likely to litter two weeks later. Clearly self-conceptions have an effect.

Of course, virtue ethicists need to show that there are character traits with moral significance. Perhaps our personal preferences and self-conceptions are too personal to count. Hippies and preppies may behave in exactly the same way in all the social psychology experiments that we have been looking at. Tidiness is arguably a morally relevant trait, but it is hardly important enough to support the contention that character should have a central place in ethical theory. Is there evidence for personal traits that influence more important morally relevant behaviors? One obvious example is political identity. Liberals and conservatives behave different in at least one context that counts: the voting booth. And this behavioral difference is certainly morally significant. People usually don't think of political liberalism and conservatism as character traits, but notice that these attitudes have much in common with character; they are dispositions to reason in specific ways across a range of issues under a range of different conditions. Moreover, people identify with their political views. They feel passionate about them, and they often refuse to associate with people whose political views are very different. Parents also try hard to cultivate specific political perspectives in their offspring, and this process of inculcation can be likened to moral education. For the moment, let's agree that political identity is *like* character.

Other morally relevant, behaviorally efficacious, character-like traits can be found when we make cross-cultural comparisons. For example, it seems reasonable to say that members of some cultures are more aggressive or violent than others. Consider the

Yanomamö of the Amazon basin, who engage in endless raids against their neighbors (Chagnon, 1968). Or consider the Ik of West Africa, who after a period of profound economic hardship, adopted patterns of behavior that observers saw as selfish and cruel (Turnbull, 1972). Or consider the Illongot of Luzon in the Philippines, who were inveterate headhunters (Rosaldo, 1980). Some people think that the violence of these cultures has been exaggerated by anthropologists, who are shocked by cultural rituals that differ from our own. But the data suggest that many small-scale societies really are much more violent than our own. For example, in pre-colonial New Guinea, many tribes had a male homicide rate exceeding 30 percent (Wrangham, 2004). Numbers like that are far greater than what we see in the post-industrial world. The incidence of death by violence was far greater in New Guinea than it was in Europe during the height of the Second World War. It is reasonable to presume that these cultures cultivated violence. Chagnon describes Yanomamö men as “fierce people,” who beat their wives, engage in brutal games, kill each other, and explode into violent rages at minor provocations.

Some social scientists have tried to identify cultural variables that determine the prevalence of violence. Edgerton (1971), for example, hypothesized that economic variables may play a major role. Agriculturalists, he reasoned, are likely to be peaceful because farming depends on cooperation and emerges in ecological settings that have rich natural resources. Herding, in contrast, tends to emerge in ecological settings with rough terrain, and herders, unlike farmers, work in isolation. Moreover, herders are always at risk, because their livelihood depends on animals that can be abducted by competitors. So Edgerton predicted that herders would be more violent in their attitudes than agriculturalists and, in a study of several African groups, that is exactly what he found. Recently, Nisbett and Cohen (1976) tested this hypothesis in the United States. They reasoned that Southern white Americans are descended from poor Scotch-Irish herders and Northern white Americans tend to be descended from farmers from wealthier parts of Britain and other regions in Europe. Nisbett and Cohen reasoned that, even today, Southern and Northern whites may be heirs to different cultural value systems as a result of their ancestry, and this may affect behavior. They use this theory to explain the fact that violent crimes are more prevalent in the South, and often arise as the result of minor provocation. Southern whites also have more tolerant attitudes towards violence (including capital punishment, gun control, physical discipline of children, and self-defense), and, in laboratory settings, they show more stress and aggression than Northerners when they are insulted. The differences between Southern and Northern whites certainly look like a difference in character.

Another cultural difference that seems to affect character is the degree of individualism or collectivism. Individualists value autonomy and personal achievement, and collectivists value interpersonal dependencies and group harmony. The cultures of Western Europe and people with Western European heritage tend to be individualists, and cultures in the Mediterranean, South America, South Asia, and East Asia tend to be collectivist (Hofstede, 2001). Members of individualist and collectivist cultures differ in attitudes and behavioral tendencies (Triandis, 1995). For example, modern individualists send aging relatives to nursing homes, while collectivists tend to care for those relatives themselves. Individualists tend to be frank about their attitudes, highly expressive, and comparatively unconcerned about embarrassment, whereas collectivists tend to be more guarded in their expression of emotions, and highly concerned with saving face in the

presence of others. Individualists tend to be self-promoting, and collectivists tend to be motivated by a sense of duty to others. These differences correlate and are reflected in personality traits. For example, Hofstede and McCrae (2004) found a .64 correlation between individualism and extroversion.

Individualists and collectivists also differ in behaviors that have obvious moral significance. In one series of studies, psychologists examined helping behavior in a number of different countries (Levine et al., 2001). For example, they watched to see if people would retrieve a pen for a stranger who “accidentally” dropped it. People were found to be especially helpful in cultures that rate high for collectivism, including Brazil, Malawi, India, and China. Many individualist cultures were less helpful, including the United States (New York City), the Netherlands, and Italy. The correlation between helpfulness and collectivism was not perfect (there are some very unhelpful collectivist cultures, such as Malaysia), but it was positive. There was a very strong correlation between helpfulness and economic variables. People in poor countries tend to be more helpful than people in rich countries.

The research on helping is striking in the context of the present discussion, because it relates to the behavior studied by Isen and Levin in their phone booth study. Isen and Levin found that a situational variable (the presence of a dime) was the major predictor of whether someone would help a stranger pick up some papers. The cross-cultural study suggests that another major factor is nationality. 100 percent of the Brazilians in Rio helped retrieve a pen when a stranger dropped it, but only 31 percent of the New Yorkers helped. That’s a big difference. Perhaps all the Brazilians would have helped had they been subjects in Isen and Levin’s study, regardless of whether they found a dime. This suggests that enculturation can lead to predictable patterns of behavior—including behavioral dispositions that render situational variables ineffective or irrelevant.

Some of the classic situationist experiments have actually been conducted cross-culturally, and the results have been striking. Consider the Milgram study. When conducted in the United States, 65 percent of the subjects were fully obedient: they continued to administer electric shocks up to the highest voltage request by the experimenter. German subjects in the same experimental set-up were considerably more obedient, following orders 85 percent of the time (Mantell, 1971). The least obedient were the Australians, who had a 28 percent compliance rate; 40 percent of the men and 16 percent of the women were fully obedient (Kilham and Mann, 1974). This range is spectacular, and it aligns stereotypical views about cultural differences. Americans are hyper-individualistic (every man for himself!), and Germans are extremely obedient, so neither tended to assist the victim in the experiment. Australians are resolutely anti-authoritarian, so they readily resisted the experimenter’s demands. These stereotypes may be pseudo-explanations, of course. It’s an open empirical question why people in different countries perform differently on the same experiments, but the fact is that they do. This brings some legitimacy to the notion of “national character.” That term has fallen out of favor because it invites negative and unfounded generalizations, but it may be a scientifically legitimate construct. Indeed, performance on the Milgram studies suggests that cross-national variance may be greater than within culture variance. Remember, Germans were more than five times as likely to be fully obedient than women in Australia! Character was invisible when Milgram first did these studies, because he used an exclusively American sample. Within a nation, situational variables drive most of the

variance. But, once we look cross-nationally, character seems to re-appear.

All of the examples that I have been considering tend to support the idea that there are efficacious character traits. The empirical evidence presented by Doris and Harman is incomplete. Once we begin to look into self-conceptions, political attitudes, and national origins, we see systematic variation across individuals. This seems to vindicate the postulation of efficacious character traits.

Doris and Harman might resist this assessment. I will consider two objections on their behalf. First, they might point out that the examples I have been considering actually confirm their thesis that people are swayed by social pressures, rather than counting against that thesis. For notice that self-conceptions, political attitudes, and national character are all forged under the influence of social groups. Becoming a hippy or a preppy is a matter of adopting an identity that has been established by others. Political affiliations are (disturbingly) influenced by demographic factors such as age, sex, gender, religion, income, and urban vs. rural habitation. One of the landmark studies in situationist psychology is Newcomb's (1943) longitudinal analysis of women at Bennington College whose were enculturated to become politically liberal; upon entering Bennington, 60 percent of the class supported the Republican presidential candidate, and, upon leaving, support dropped to 15 percent. Studies of national character are obviously illustrations of social influence as well. Our tendency to be individualists or to be obedient to authority is influenced by the ethos of the nations in which we are reared. Thus, the evidence adduced in this section is resounding confirmation of the situationist perspective.

In response, I want to draw an important distinction. Factors external to a person can influence behavior in two different ways: synchronically or diachronically. In the studies by Milgram, Darley, and Isen, behavior is affected by the current situational context, but, in the cases I have been considering, the external influence has already occurred before the subjects enter the lab. This difference matters because the two forms of influence have different implications. If we were swayed only by synchronic factors, then all people would be the same: put two people in the same situation, and they will probably do the same thing. But, if diachronic influences are possible, then people can internalize social norms, and, as a result, people with different backgrounds will behave differently in the exact same situations. If all people behaved alike in the same situations, character based ethical theories would be in trouble: it would be impossible to *cultivate* character. At best, we could do what Doris recommends: try to put ourselves in situations that promote good behavior. But, if diachronic influence is possible, the cultivation of character *is* possible. This is enough to get virtue ethics back in the door. Virtue theorists do not insist that character initiates from an internal source. We may need to avail ourselves of accumulated social wisdom, relationships with others, and disciplined moral education to achieve ideal traits.

A second objection that Doris and Harman might raise questions whether my examples have anything to do with character traits. I've mentioned things such as being a jazz enthusiast, being a hippy, being a Republican, and being a collectivist. These categories look very different from traditional examples of character traits, such as being loyal or being courageous (see also Ross and Nisbett, 1991: 202). Haven't I just changed the topic? I don't think so. Recall that the traits I am talking about here have all the features that character traits are supposed to have: they are long-term, global, dispositions

that influence practical reasoning and behavior. If there is a principled distinction here, I don't see it. Furthermore, the traits I have been discussing can affect traditional traits such as courage or loyalty. Perhaps the Yanomamö are more courageous than some peaceful agriculturalists. Perhaps members of collectivist cultures are more loyal than members of individualist cultures. Perhaps Germans are more obedient than Australians.

In sum, I think there is an empirical case to be made for the existence of efficacious global character traits. Making such a case is, in my view, the best way to address the Situationist Challenge. If such traits exist, then perhaps they can be cultivated to serve the demands of morality. Cultivating virtuous traits that do not get overwhelmed by synchronic situational variables may be difficult, but there is no reason to think it's impossible. Dramatic differences in behavior can be found across cultural borders.

3. Normativity: Another Empirical Challenge to Virtue Ethics

I have just been arguing that there are efficacious global character traits, and should come as good news to fans of virtue ethics. But my ultimate goal is not to defend virtue theory. On the contrary, I think virtue theory faces another empirical objection, which may prove more damaging than the Situationist Challenge. I will introduce the objection in this section, and I will call it the Normative Challenge. In the final section, I will argue that character traits of the kind I have just been discussing cannot point us towards a solution. In fact, such traits may exacerbate the Normative Challenge.

One of the central questions facing any ethical theory concerns the source of normativity. Virtue ethics is a normative theory. Its proponents say that we *should* cultivate the virtues. But why? What is the source of this obligation? Within moral philosophy, there tend to be four strategies for addressing such questions. Some normative theories invoke divine command: we should do thus and such because it the will of God. Other normative theories purport to have a rational foundation: we are told that the norms they recommend are demanded by reason. Kantian ethics falls into this category. Still other normative theories make reference to passion, rather than reason: norms get their normative force from our own preferences and desires. Utilitarians fall into this camp, and so do subjectivists, emotivists, and sensibility theorists. Finally, there are theories that ground normativity in nature: we should obey norms that are dictated by natural teleology—act to fulfill the ends towards which humans, as a particular kind of animals, are naturally aimed. Some evolutionary ethicists go this route. (Evolutionary error theorists do not, because they think morality lacks normativity.) Theory that explain normativity in terms of passion or nature are said to be naturalist theories, because they explain the source of normativity by appeal to some natural feature of the world, such as our psychological states or biological nature. On these views, the basis of normativity is open to scientific investigation, and norms are usually thought to have their force contingently: they depend on contingent facts about us, rather than, say general facts about agency. There may be other candidate sources of normativity, but this taxonomy captures many of the major ethical theories.

Where does virtue ethics fit in? On some versions, God is the source of normativity. Virtue ethics has long been popular among Christian philosophers, and some Christians and other theists think morality issues from divine command. Aristotelians tend to go another route. They argue that morality derives from natural

teleology. The virtues are the proper teleological ends of creatures like us. In the *Nicomachean Ethics*, Aristotle pursues this idea by suggesting that human beings are by nature rational, and that we can fulfill our telos if we cultivate habits of practical reason that dispose us to behave in accordance with the doctrine of the mean: we should acquire traits that steer between deficiency and excess. Such traits bring emotions in line with our rational capacities, and thus constitute a way of being that exemplifies our rational nature. Virtues endow us with characteristically human lives. The normativity of virtues derives from the fact that they are the end to which our nature directs us, and thus constitute human flourishing.

Contemporary virtue ethicists usually shy away from Aristotle's doctrine of the mean, but they do invoke the notion of flourishing. They say that creatures like us will naturally flourish to the extent that we possess virtuous traits. Virtuous traits can be identified as those that promote flourishing, and because they allow us to flourish, we ought to cultivate them. Notice that on this formulation, virtue ethics appears to straddle the line between the view that normativity derives from nature and the view that normativity derives from passions. For flourishing is widely presumed to involve certain affective states. Most notably, flourishing is related to well-being, and well-being is an affective construct. If virtues gain their normative force from well-being, then there is a sense in which normativity depends on the passions. But I suspect that most virtue ethicists would say that this way of putting things is a bit misleading. Virtues are good not simply because they make us feel good, but because they make us feel good in a way that is indicative or constitutive of having fulfilled our natural ends as a species. Well-being confers normative status not because of its hedonic qualities, but because of its teleological status. Indeed, well-being is not reducible to any particular feeling, nor is it simply the fulfillment of ends. Rather, it is a state we recognize in ourselves on the basis of a wide range of different factors, including feelings (both positive and negative), ends (both achieved and pursued), relationships, capacities, and so on. Well-being is a subjective state, but when serving as the normative foundation for virtue theory, it is thought to depend on criteria that are in some sense objective: well-being is a measure of whether we our lives exemplify our full potential as humans.

There are philosophical objections one might raise to this account of how the virtues get their normative grounding. One might worry that virtue ethics depends on an unwarranted conflation of the natural and the good. Natural certainly doesn't seem to entail good. For example, human beings might be naturally violent (Keeley, 1997). As already remarked, many hunter-gatherer societies are exceedingly violent (Wrangham, 2004), and evidence suggests that extreme violence and warfare have always been a factor in human life (Keeley, 1997). The tendency to be violently territorial certainly seems to be typical of our species, even in industrial societies, yet we hesitate to call territorial violence a moral good. To avoid conflating the good with the natural, virtue ethicists are forced to draw distinctions between natural tendencies that are noble and those that are ignoble. It's hard to do this without circularity. One cannot define noble natural tendencies as those that accord with virtue, and then argue that virtues derive normativity from their status as natural.

Faced with this worry, virtue theorists might be inclined to beef up their notion of well-being. If a substantive account of well-being can be offered that makes no reference to virtue, then we can define the virtues as the natural behaviors that promote well-being

without any circularity. But this is tricky business. To offer a substantive account of well-being, one might embark on the empirical project of identifying the dimensions that are most determinative of positive life assessments. One might investigate the conditions under which we judge our lives to be going well. If a list of conditions can be identified, then the virtues can be defined as character traits that promote those conditions.

This is a feasible empirical project, and indeed there has been a huge effort in recent psychology to identify the factors that contribute to well-being. That effort has been successful insofar as factors have been identified, but the results offer little comfort to the virtue theorist. The problem is that different factors emerge for different people, and there are well-documented cultural differences in which factors matter (for philosophical discussions of this literature, see Tiberius, 2003). One of the most extensively investigated variables of cultural difference is the contrast between individualism and collectivism. As noted above, people in the West tend to be individualists, and people in the Far East tend to be collectivists. It turns out that this cultural difference has a very significant impact on conceptions of well-being. In the West, self-esteem is very important to well-being and in the East it's less important or, in some studies, totally unimportant; for example, for American women, there is a .60 correlation between life satisfaction and self-esteem, but for Indian women, the correlation is .08, which is not statistically significant (Diener and Diener, 1995). In the West, personal enjoyment is important to well-being, and in the East it is more important to please others (Diener et al., 2003). In the West, well-being is largely a function of one's present state, and in the East well-being is correlated with working towards future goals (Asakawa and Csikszentmihalyi, 1998). In the West, well-being is associated with certain emotions, and in the East it is more associated with the fulfillment of obligations (Suh et al., 1998). And significantly, in the West, well-being is correlated with perceived consistency of identity, but such consistency is less important to people in the East (Shuh, 2002). These five contrasts are probably the tip of the iceberg. There are likely to be considerable cultural variation in the conditions of life satisfaction. For example, Nussbaum (2000) mentions that members of some cultures may find exposure to the natural environment important more important than others. There will also be individual differences in the value placed on solitude versus social interaction, physical versus intellectual activity, and the importance of various pleasures such as sex, art, and novelty.

This variation has serious implications for virtue ethics. Virtue ethicists have traditionally assumed that there is a universal set of virtues; indeed many virtue ethicists try to list them. And, they assume that these virtues are universal precisely because they are all part of a universal human nature. But that supposition is untenable. What leads to fulfillment in life is neither universal nor entirely natural. Culture can shape our conception of the good life. The virtues that facilitate well-being in the West (if any list can be given) will not be exactly the same as the virtues that facilitate well-being in the East. For example, Aristotle includes pride, ambition, and wit among the virtues. These strike me as particularly Western. Perhaps one could define them so vaguely and abstractly as to obtain traits that most people would regard as worthwhile, but, in so doing, one would render them vacuous. The kind of ambition valued in the East may differ substantively from the kind that is valued in the West. Indeed, I would venture that the kind of ambition, pride, and wit valued in Aristotle's culture differ from the traits that come to mind when we read these English translations of his Greek terms. Scholars work

hard to render Aristotle's list in contemporary English, but they are thwarted by nuances of meaning that may be forever lost to history.

When we look to the history of philosophy, we see authors in different traditions offering lists of virtues that differ from Aristotle's. Roman virtue ethicists promoted *virtus*, a kind of military valor, associated with courage in the face of death. This made sense in an imperialist society that needed to constantly secure its borders. For Christian virtue ethicists, like Augustine, chastity was added to the list of virtues. The church promoted a conception of well-being linked to faith and directed away from carnal pleasures. For us, the good life might involve recreational sex, but, for the medieval Christian, this may have been a source of guilt and shame. Arguably, the Christian attitude toward sex was culturally constructed to reduce family sizes, which increased the relative power of the church (Goody, 1983). For Confucius, who also defended a version of virtue ethics, the most important trait was filial piety: a deep respect for family, which is still prevalent in Chinese culture. Then as now, people in China seem to have found well-being in interpersonal relationships. Those relationships are markedly underrepresented in Western lists of virtues.

One might go so far as to argue that virtue theory is itself a byproduct of a certain historically constructed conceptions of well-being. The idea that virtues are necessary for (or even conducive to) flourishing may derive from the Western tendency to value consistent identity over time. As remarked above, members of Eastern cultures do not place equal value on identity. In the East, well-being is not necessarily attained by the exercise of enduring character traits but may instead benefit from (or even require) a kind of fluidity of character that would diminish from well-being in the West. Perhaps Aristotelians believe that virtues are normatively grounded because they belong to cultures that inculcate the idea that we need to be virtuous to thrive. For us, virtues may be essential to a maximally fulfilling life, but we cannot presume that this is true everywhere. Virtue ethics has been less prevalent in the East, where well-being does not presuppose personal consistency. Confucius is an exception, but his version of virtue theory differs in a crucial respect from its Western counterparts. Filial piety can be construed as a pattern of interpersonal connectedness, and as such it actually diminishes the extent to which its possessor can be analyzed as a self-propelled or autonomous entity. To be virtuous for Confucius is to have traits that render trait-based explanations of behavior inadequate on their own. It is interesting to compare this to some Indian conceptions of Nirvana, according to which transcendence comes from the elimination of self. Nirvana may be the height of well-being, but it is not one for which cultivation of traits would be appropriate. Likewise for Confucius, traits may be a ladder that can be kicked away after fluid interconnectedness is achieved. Westerners coming out of the Hellenistic tradition may have special claim on the claim that traits are purely internal dispositions that should be the main aim or morality and the main ingredients of well-being. In the West, having certain consistent internal traits may be constitutive of well-being, whereas in the East, traits are more likely to be construed as means to forms of well-being that are fundamentally relational.

In summary, the cultural variation in the sources of well-being pose a serious threat to the idea that there is a single set of virtues that are conducive to human flourishing. Such variation may even cast doubt on the claim that flourishing depends on the cultivation of character traits; this goal may be more conducive to flourishing in the

West, where individual achievement and characterological consistency are highly prized. To address this objection, virtue ethicists must argue that one set of virtues promotes all forms of well-being, or they must argue that there is a universal form of well-being. I don't see much promise in either strategy. It seems to be empirically true that conditions of well-being vary, and the variations run deep. There is no obvious common denominator. We cannot demonstrate that one conception of well-being is better than any other, because such an argument would inevitably hinge on one of two mistakes. On the one hand, such a comparative assessment might be made from within an evaluative framework, in which case we would inevitably impose our own conception of well-being on others, when determining which is best. On the other hand, we might try to find a neutral position from which to compare conceptions of well-being; for example, we might argue that some conception is more natural. But this assumes a conception of the human species that is profoundly false. We are a cultural species, and it is part of our nature that our values should be forged in the context of human interactions. To the extent that those interactions engender different conceptions of well-being, as the empirical literature shows, we cannot pretend that there is some pre-social, purely natural conception. That is the myth of the noble savage. Indeed, if we could find such a pre-cultural conception, it would have scant normative force, for conceptions of well-being that did not emerge through cultural processes would hardly be applicable or conducive to thriving once we find ourselves situated in a cultural context. Faith in a universal form of well-being teeters between cultural chauvinism and a form of naturalism rivaling Spencer's in its naivety and vulgarity.

At this juncture, virtue ethicists might opt for another strategy. They might argue that we can dispense with a single notion of well-being and, for that matter, with a single list of virtues. But this proposal will do little to preserve the standard virtue theoretic answer to the question of normative grounding. Recall, that Aristotelian virtue theorists often endorse the view that virtues are good because they are natural, where that means that virtues are the key to flourishing in accordance with our full potential as human beings. The argument that I have just sketched undermines this approach to normative grounding. There is nothing purely natural about the notion of well-being, and, as a result, if there is nothing purely natural about the notion of flourishing. Therefore, virtues cannot gain the normativity from the status as instruments serving natural ends.

As far as I can see there is only one plausible way out. The virtue theorist is forced, for empirical reasons, to abandon Aristotelian naturalism about the source or normativity. Having abandoned Aristotelian naturalism, there is one of two options: either virtues lose their normative status, in which case virtue theory must be abandoned, or their normative status must derive from another source. Let's assume that the theological source is untenable (either because Plato was right in the Euthyphro or because God doesn't exist). That leaves reasons and the passions. I will not here explore the idea that virtues may have their basis in reason (I don't think any behavioral norm has its basis in reason), and instead focus on the plausible suggestion that virtues are normatively grounded in the passions.

The suggestion I want to entertain is that norms of character hold, if at all, in virtue of our desires, preferences, or other affective/motivational states. More specifically, I want to suggest that such norms hold in virtue of our sentiments, where sentiments are dispositions to have emotions of approbation or disapprobation. Consider

a norm such as: you ought to be a generous person. Why does this norm hold? One possibility is that it holds because we approve of the trait of generosity and disapprove of misers. From within the value system that we have internalized, the trait of generosity is commendable. It has normative force over us because we value it. It may not have normative force from within some other system of values—some other form of life—but such systems are neither mentally available to us in any deep sense, nor are they relevant to us, when we are deciding what we most deeply care about. I have elsewhere defended the claim that morality derives from sentiments at great length (Prinz, 2007). I will not present that case here. I present this suggestion to make a point. On the one hand, if norms about character can hold in virtue of our sentiments, then virtue theorists are right that such norms exist. Although it's an empirical question, I would bet a small fortune that people have sentiments about character traits. On the other hand, if such norms attain their normative status from our sentiments, then Aristotelian virtue theory is mistaken. They are mistaken about the source of normativity, but also about the place of the virtues in moral theory. Aristotelians argue that the central ethical question is, what sort of person should I be? But if norms about character derive from our sentiments, then our sentiments determine what questions are central. There is good reason (intuitive and empirical) to think that we have many sentiments pertaining to questions of how we should act. We disapprove of stealing and approve of charity. If sentiments are the source of normativity, then norms of character have no privileged place. They are exactly like norms of action. They all derive from our sentiments.

Of course, not everyone believes that normativity derives from sentiments. Not much hangs on this claim, because parallel morals follow from other ethical theories. Suppose you are a consequentialist. Then rules derive their normative status from their contribution to net utility (or some other commodity). Once again, action-directed rules and character-directed rules will be on all fours. Character and action can both increase utility. If you are Kantian about the source of normativity, it will be even harder to show that character has precedence over action, because Kantian norms are usually understood as action oriented: it's not clear whether the moral law has direct implications for character. As long as your actions do not undermine your will, you can have any kind of personality that you like (though see Kant, 1785, on the "Doctrine of Virtue").

Now let's take stock. The Normativity Challenge can be succinctly stated as follows: empirical evidence suggests that there is no purely natural conception of well-being, and that undercuts the account of normativity central to leading versions of virtue ethics. The Challenge can be elaborated more fully as follows. First, empirical findings about variation in conceptions of well-being threaten the idea that there is one set of virtues. Second, such findings thereby undermine the Aristotelian conception of where virtues get their normative status. Third, if virtues have normative status, then they may get that status from another source, such as our sentiments, and if that's the case then norms pertaining to virtues may not be privileged over norms pertaining to actions (we have sentiments about both). Each of these points undercuts a central element of Aristotelian virtue ethics. In the final section, I will raise one more worry to add to this list.

4. Do Empirically Demonstrable Traits Answer The Normativity Challenge?

I just argued that the Normative Challenge threatens Aristotelian ethics by suggesting that norms about character may not be privileged. If normativity issues from our sentiments (or some other source), then there may be normative rules pertaining to both how we should act and to what sort of people we should be. Thus, virtue theory is not an alternative to action-based ethical theories. It merely captures one dimension of human moralizing. In this section I want to push the objection a little harder. I want to suggest that action-directed rules may actually be more fundamental than character-directed rules. If I am right, then virtues are normatively parasitic on rules of action. This would be a major blow to virtue theory.

To make this case, I will build on the assumption that normativity derives from our sentiments, though parallel arguments might be developed for other theories of the source of normativity. If sentiments are the source, then something is morally good if you have an attitude of approbation towards it (see, Prinz, 2007, for an analysis of what such sentiments consist in). It seems very plausible that we have a sentiment of approbation towards certain character traits, such as honesty. But ask yourself: why do you praise honest people? If you're like me, you praise them because you consider truth-telling valuable. If you didn't care about truth, you wouldn't care about honesty. Thus, there is an asymmetric relation between our attitudes towards traits and our attitudes towards actions; the former depend on the latter, but not conversely. If Aristotelians could defend the claim that norms about virtue have their source in natural human flourishing, the logic of our preferences would not matter. If the good life is the normative bedrock, then character traits that are conducive to (or constitutive of) the good life would not necessarily depend normatively on acts. On that approach, the normative status of an act many depend on whether it issues from a good character trait. But, if sentiments are the normative bedrock, then character traits get their normative force from our sentiments, and, I suspect that sentiments towards traits are parasitic on the actions associated with those traits.

This, of course, is speculation. We would need to investigate the structure of human preferences to establish that character norms are asymmetrically dependent on action norms. I don't know of any research that directly looks at this, so I offer my assessment as a tentative empirical conjecture. I suspect that our sentiments are first and foremost about actions, and only secondarily about character traits, insofar as traits promote actions that favor or disfavor. If so, virtues cannot be the foundation of morality.

Is there any evidence that we value virtuous character traits in a way that is not parasitic on our action preferences? I will conclude by briefly considering three empirical arguments that might be used to argue for the claim that we cherish virtue in a non-parasitic way.

First, virtue theorists might call on some of Kohlberg's (1984) work on moral development. According to Kohlberg, there is a stage of moral development in which people are very concerned with character. Kohlberg calls the Good Boy/Nice Girl stage, where children seem to become preoccupied with how they are regarded as people. They do good things because they want to be good people, and, thus, they regard virtue as primary. Does this finding establish that virtue is not parasitic on other norms? I don't think so. For one thing, I doubt that kids ever pursue virtue as an end in itself. Rather, they want approval from others, and being a good is a way to get it. For another thing,

the Good Boy/Nice Girl stage passes quickly. When older children and adults reason about morality, they rarely say that an action is wrong because it's what a bad girl would do. This gets things backwards, for the mature moralizer. People are good if they perform good actions. Kohlberg actually thinks that virtue has little role in mature moral psychology. I don't want to rely on Kohlberg too heavily. His theory has been criticized on various grounds, and it's unclear whether it holds up cross-culturally (see, e.g., Flanagan, 1982; Snarey, 1985). But, I certainly don't think Kohlberg's data provide reason to think that virtues are central to moral life. If we were to look cross culturally, we might even find that people in other societies are even less concerned with virtue than we are.

Second, virtue theorists might try to establish the psychological importance of virtue by citing the empirical literature on reputation. There is evidence that people care about their reputations, and that they use a good reputation as a tool to earn the trust of others. Much of this research comes out of behavior economics (e.g., Frank, 1988). Economists construe human relations as attempts to resolve various coordination problems. Consider the prisoner's dilemma. Defecting against another person is more attractive than cooperating, because, defecting dominates (i.e., no matter what the other player does, defecting comes out ahead of comes out ahead); but that means that both players are likely to defect, and mutual defection is has a power payoff that mutual cooperation. In iterated prisoner's dilemmas, people figure this out, and can end up cooperating. But what if you are playing a one-shot game? What incentive would you ever have to cooperate? One answer is that, by cooperating, one can establish a positive reputation, and, as a result more people will want to enter into ventures with you, because they know you won't cheat them. In other words, we have reason to pursue positive reputations; they earn us trust, which is an extremely valuable form of "social capital." This theory is supported by data. People will cooperate in games that aren't iterated, and they will do good things for others when there is little chance of direct reciprocation.

Does research on reputation show that character traits are sometimes normatively prior to actions? Absolutely not. First, it's not clear that reputation is a character trait. Rather, it might just be a statistical pattern. A vicious person might have a reputation for always cooperating, because he recognizes that cooperation is a winning strategy. In this context, it's worth nothing that psychopaths often cooperate on prisoner's dilemmas (Widom, 1976). Second, reputation is not an intrinsic good, but rather an instrumental good. If we didn't desire successful coordination with others, we wouldn't value good reputation. But the reverse is not true: we would continue to value cooperation even if we stopped valuing reputation. Finally, there is experimental evidence that suggests reputation is not a reliable predictor of cooperation. Liberman et al. (2004) asked dormitory RAs (resident assistance) to rate the likelihood of cooperation of their dorm residence based on their reputations. They then had those residents play a prisoner's dilemma game, labeled either "The Wall Street Game" or "The Community Game." Reputation had absolutely no predictive value on the likelihood of cooperation, but, in line with situationist psychology, the label had a dramatic impact. Mutual cooperation occurred four times as often when the game was labeled with the word "Community" rather than "Wall Street." This suggests that even if we do value reputation, perhaps we oughtn't.

Let me turn to a final empirical argument that virtue ethicists might advance to

show that we often value character in a way that does not depend on valuing behavior. Ostensibly, if behavior were our primary concern, there would be little pressure to develop consistent character traits. Character would be truly unusual, just as situationists claim. But, earlier I argued against the situationists that character traits are quite common. I gave various examples: being a jazz enthusiast, being a hippy, being an individualist, being aggressive, being liberal, being obedient, and so on. Each of these cases seems to qualify as an efficacious global disposition. Character traits are both real and pervasive. Doesn't this show that we value character as much as we value behavior? If we didn't value character, wouldn't behavior be much more erratic?

I don't think that the examples under consideration help to establish the thesis that we have non-derivative sentiments towards character traits. On the contrary, I think examination of these examples actually supports the thesis that our attitudes towards character traits depend on our attitudes towards action. If I am right, then these traits actually undermine the thesis that character has primacy in our moral psychology, and, thus, these traits can actual figure into an argument against virtue ethics, rather than proving support.

To make this case, let me draw a terminological distinction. In cultivating character traits, our ultimate goal, whether explicit or implicit, can take one of two forms. We can describe the goal as "first order" if the main focus of our efforts in the behavior associated with a trait. The goal is "second order" if having the disposition to produce that behavior is the main focus of our efforts. Typically, second order goals are achieved by targeting the will and first order goals are achieved by targeting behavior directly. Socialization can operate at either of these two levels. Our teachers, peers, and role models may tell us that we ought to certain kinds of people, or they may tell us that we ought to engage in certain kinds of behavior. The net result may be the same. If we are successfully socialized to be honest *people* (e.g., to operate from honest motives), then we acquire the disposition to tell the truth, and that is exactly the disposition that would arise if we were directly socialized to exhibit honest *behavior*. In either scenario we end up with a character trait, but, in the first scenario, that was target of socialization, and in the second scenario, it was merely a consequence. It is a central tenet of virtue theory that character has priority over action. I want to suggest that our ultimate goal in cultivating character is first-order rather than second order.

Under what kind of pressures, do people end up having the character traits that they do? To answer this question, let's consider the kind of character traits that I have been defending—traits that derive from taste, self-conceptions, and cultural norms. I suspect that the social mechanisms that promote these traits are first order, rather than second order. If I become a jazz enthusiast, it's probably because I heard some jazz and enjoyed it. That is a first-order induction into the trait. If I become an individualist, it's probably because I live in a culture that rewards me for being self-reliant. If I develop a hair trigger or a bellicose personality, it's probably because I am often under threat from others who are disposed to violent. If I become obedient, it is probably because I have been encouraged to obey. If I become a liberal, it may be because I am enculturated to passionately support several items on the liberal platform. Even the decision to become a hippy may be driven at first by certain behavioral ends, rather than characterological ends: I want to have membership in a group whose behavior I envy or esteem. It is plausible in each case that the factors driving us towards culturally based character traits

are implicitly or explicitly first order: behavior is shaped and character follows.

These are just conjectures. I may be wrong about how real character traits develop. My claims here are tentative. But suppose I'm right. If real character traits are developed as a result of pressures on behavior, then character is really an emergent phenomenon. People don't ordinarily strive to have the character that they have; character is just the result of behavioral regularities. If you love sweets, you can be labeled as having a sweet-tooth, but that personal attribute is not one you ever sought. Virtue theory requires that character traits have normative status, and, I have proposed, that normativity must issue from our own sentiments and preferences. Typically, character traits are not things we seek, but rather they are things we possess in virtue of seeking other ends. If this assessment is right, then virtue theorists should find no comfort in the fact that character traits are commonplace. Traits emerge from first-order mechanisms, and, therefore, lack the normative status that virtue theory requires.

I do not mean to suggest that we never value character traits as ends in themselves. Perhaps we do. I want only to say that this may be a rare phenomenon. Most of the time, the traits we have are the result of first-order preferences. Thus, if there are norms of character, they may be the exception, rather than the rule. That's a blow to virtue theory. At the beginning of this section, I also endorsed a stronger thesis: when we do value character traits, those values are derivative. They depend on the fact that we value the behaviors that the traits reliably bring about. I have not presented empirical evidence for this stronger claim about the structure of our preferences, but I have raise doubts about three arguments that might be marshaled against it. In the last section, I presented a Normative Challenge for virtue ethics. It was a upshot of that challenge, that character traits attain normative status as a result of human preferences or sentiments. If I am right to say that our attitudes towards norms hinge on our attitudes towards actions, then Aristotelian Virtue theory faces a final and fatal objection. If I am right, then the normative status of rules prescribing virtuous character traits is neither prior nor equal to the normative status of action-directed rules. Rather, virtues have their normative status parasitically. The problem with virtue ethics is not that character traits don't exist, but rather that they cannot form the foundation of a normative theory.

References

- Annas, J. (2005). Comments on John Doris's Lack of Character. *Philosophy and Phenomenological Research*, 71, 636-642
- Asakawa, K., and Csikszentmihalyi, M. (1998). The Quality of Experience of Asian American Adolescents in Activities Related to Future Goals. *Journal of Youth and Adolescence*, 27, 141-163.
- Athanassoulis, N. (2000). A Response to Harman: Virtue Ethics and Character Traits. *Proceedings of the Aristotelian Society*, 100, 215-221.
- Chagnon, N. A. (1968). *Yanomamö: The Fierce People*. New York, NY: Holt, Rinehart and Winston.
- Costa, P.T. Jr., Terracciano, A., and McCrae, R.R. (2001). Gender Differences in Personality Traits Across Cultures: Robust and Surprising Findings. *Journal of Personality and Social Psychology*, 81, 322-331.

- Darley, J. M., and Batson, C. D. (1973). From Jerusalem to Jericho: A Study of Situational and Dispositional Variables In Helping Behavior. *Journal of Personality and Social Psychology*, 27, 100-108.
- Diener, E., and Diener, M. (1995). Cross-Cultural Correlates of Life Satisfaction and Self-Esteem. *Journal of Personality and Social Psychology*, 68, 653-663.
- Diener, E., Oishi, S., and Lucas, R. E. (2003). Personality, Culture, and Subjective Well-Being: Emotional and Cognitive Evaluations of Life. *Annual Review of Psychology*, 54, 403-425.
- Doris, J. M. (1998). Persons, Situations, and Virtue Ethics. *Noûs*, 32, 504-530.
- Doris, J. M. (2002). *Lack of Character: Personality and Moral Behavior*. New York, NY: Cambridge University Press.
- Edgerton, R. B. (1971). *The Individual in Cultural Adaptation: A Study of Four East African Peoples*. Berkeley, CA: University of California Press.
- English, A. and Griffith, R.L. (2003) Trait Consistency and The “Big Five.” Paper presented at the annual meeting of the Society for Industrial Organizational Psychology, Chicago, Illinois.
- Flanagan, O. (1991). *Varieties of Moral Personality*. Cambridge, MA: Harvard University Press.
- Flanagan, O. J. (1982). Virtue, Sex and Gender: Some Philosophical Reflections on the Moral Psychology Debate. *Ethics*, 92, 499-512.
- Frank, R. H. (1988). *Passion within Reason: The Strategic Role of the Emotions*. New York, NY: Norton.
- Goody, J. (1983). *The Development of the Family and Marriage in Europe*. Cambridge: Cambridge University Press.
- Harman, G. (1999). Moral Philosophy Meets Social Psychology: Virtue Ethics and the Fundamental Attribution Error. *Proceedings of the Aristotelian Society*, 99, 315-331.
- Hofstede, G. (2001). *Culture’s Consequences: Comparing Values, Behaviors, Institutions and Organizations across Nations*. Thousand Oaks CA: Sage.
- Hofstede, G. and McCrae, R.R. (2004). Culture and personality revisited: Linking traits and dimensions of culture. *Cross-cultural Research*, 38, 52-88.
- Isen, A. M. and Levin, P. F. (1972). Effect of Feeling Good on Helping: Cookies and Kindness. *Journal of Personality and Social Psychology*, 21, 384-388.
- Isen, A.M., and Patrick, R. (1983). The Influence of Positive Feelings on Risk Taking: When the Chips Are Down. *Organizational Behavior and Human Performance*, 31, 194-202.
- Jang, K., Livesley, W. J., and Vernon, P. A. (1996). Heritability of the Big Five Personality Dimensions and Their Facets: A Twin Study. *Journal of Personality*, 64, 577-591.
- John, O. P., and Srivastava, S. (1999). The Big-Five trait taxonomy: History, measurement, and theoretical perspectives. In L. A. Pervin and O. P. John (Eds.), *Handbook of Personality: Theory and Research* (Vol. 2, pp. 102–138). New York, NY: Guilford Press.
- Kamtekar, R. (2004). Situationism and Virtue Ethics on the Content of Our Character. *Ethics*, 114, 458-491.
- Kant, I. (1885/1996). *The Metaphysics of Morals*. M. Gregor (Trans.). Cambridge,

- Cambridge University Press.
- Keeley, L. H. (1997). *War before Civilization: the Myth of the Peaceful Savage*. New York, NY: Oxford University Press.
- Kilham, W. and Mann, L. (1974). Level of Destructive Obedience as a Junction of Transmitter and Executant Roles in the Milgram Obedience Paradigm. *Journal of Personality and Social Psychology*, 29, 696-702.
- Kohlberg, L. (1984). *The Psychology of Moral Development: Moral Stages and the Life Cycle*. San Francisco, CA: Harper & Row.
- Levine, R., Norenzayan, A., and Philbrick, K. (2001). Cross-Cultural Differences in Helping Strangers. *Journal of Cross-Cultural Psychology*, 32, 543-560.
- Liberman, V., Samuels, S. M., and Ross, L. (2004). The Name of the Game: Predictive Power of Reputations Versus Situational Labels in Determining Prisoner's Dilemma Game Moves. *Personality and Social Psychology Bulletin*, 30, 1175-1185.
- Mantell, D. M. (1971). The Potential for Violence in Germany. *Journal of Social Issues*, 27, 101-112.
- McCrae, R.R., Costa, P.T., Ostendorf, F., Angleitner, A., Hrebickova, M., Avia, M. D., Sanz, J., Sánchez-Bernardos, M. L., Kusdil, M. E., Woodfield, R., Saunders, P. R., and Smith, P. B. (2000). Nature over Nurture: Temperament, Personality, and Lifespan Development. *Journal of Personality and Social Psychology*, 78, 173– 186.
- Merritt, M. (2000). Virtue Ethics and Situationist Personality Psychology. *Ethical Theory and Moral Practice*, 3, 365- 383.
- Milgram, S. (1974). *Obedience to Authority*. New York, NY: Harper and Row.
- Miller, R. L., Brickman, P. and Bolen, D. (1975). Attribution Versus Persuasion as a Means of Modifying Behavior. *Journal of Personality and Social Psychology*, 31, 430-441.
- Mischel, W. (1968). *Personality and Assessment*. New York, NY: John J. Wiley and Sons.
- Mischel, W., and Peake, P. K. (1982). Beyond Déjà Vu in the Search for Cross-Situational Consistency. *Psychological Review*, 89, 730-755.
- Newcomb, T M. (1943). *Personality and Social Change: Attitude Formation in a Student Community*. New York: Holt.
- Newcomb, T. M. (1929). *Consistency of Certain Extrovert-Introvert Behavior Patterns in 51 Problem Boys*. New York, NY: Columbia University, Teachers College, Bureau of Publications.
- Nisbett, R. E. and Cohen, D. (1996). *Culture of Honor: The Psychology of Violence in the South*. Boulder, CO: Westview Press.
- Nussbaum, M. (2000). *Women and Human Development: The Capabilities Approach*. Cambridge: Cambridge University Press.
- Plomin, R. and Caspi, A. (1999). Behavioral Genetics and Personality. In L.A. Pervin and O.P. John (Eds.), *Handbook of Personality: Theory and Research*. New York, NY: Guilford Press.
- Prinz, J. J. (2007). *The Emotional Construction of Morals*. Oxford: Oxford University Press.
- Roelofs, J., Meesters, C., Ter Huurne, M., Bamelis, L., and Muris, P. (2006). On the Links Between Attachment Style, Parental Rearing Behaviors, and Internalizing and Externalizing Problems in Non-Clinical Children. *Journal of Child and*

- Family Studies*, 15, 331-344.
- Rosaldo, R. (1980). *Illongot headhunting, 1883-1974: A Study in Society and History*. Stanford, CA: Stanford University Press.
- Ross, L. and Nisbett, R. E. (1991). *The Person and the Situation: Perspectives of Social Psychology*. Philadelphia, PA: Temple University Press.
- Sabini, J. and Silver, M. (2005). Lack of Character? Situationism Critiqued. *Ethics*, 115, 535-562.
- Schnall, S., Haidt, J., Clore, G. L., and Jordan, A. H. (2008). Disgust as Embodied Moral Judgment. *Personality and Social Psychology Bulletin*, 34, 1096-1109.
- Snarey, J. (1985). Cross-Cultural Universality of Social-Moral Development: A Critical Review of Kohlbergian Research. *Psychological Bulletin*, 97, 202-232.
- Sreenivasan, G. (2002). Errors about Errors: Virtue Theory and Trait Attribution. *Mind*, 111, 47-68.
- Suh, E. M. (2002). Culture, Identity Consistency, and Subjective Well-Being. *Journal of Personality and Social Psychology*, 83, 1378-1391.
- Suh, M., Diener, E., Oishi, S., and Triandis, H. C. (1998). The Shifting Basis of Life Satisfaction Judgments across Cultures: Emotions Versus Norms. *Journal of Personality and Social Psychology*, 74, 482-493.
- Tiberius, V. (2003). Cultural Differences and Philosophical Accounts of Well-Being. *Journal of Happiness Studies*, 5, 293-314.
- Triandis, H.C. (1995). *Individualism and Collectivism*. Boulder CO: Westview.
- Turnbull, C. (1972). *The Mountain Meople*. New York, NY: Simon and Schuster.
- Widom, C. S. (1976). Interpersonal Conflict and Cooperation in Psychopaths. *Journal of Abnormal Psychology*, 85, 330-334.
- Williams, B. (1981). Persons, Character, and Morality. In *Moral Luck* (pp. 1-19). New York, NY: Cambridge University Press.
- Wolf, S. (1982). Moral Saints. *The Journal of Philosophy*, 79, 419-439.
- Wrangham, R. (2004). Killer Species. *Daedalus*, 133, 25-35.