

DRAFT: Please do not quote, cite, or distribute.

## Ordinary Virtue

Susan Stark  
Bates College

Ethics should not be merely an analysis of ordinary mediocre conduct, it should be a hypothesis about good conduct and about how this can be achieved. How can we make ourselves better? Is a question moral philosophers should attempt to answer.

Iris Murdoch, "The Sovereignty of Good Over Other Concepts," 76.

1. Most of us concern ourselves with the moral goodness of our own lives. We try to live rightly and try to make ourselves better. But this ordinary disposition can be hard to square with reflection on human psychology. One need only think of the many massacres, slaughters, and genocides perpetrated by humanity during our long and disturbing history to get a sobering picture of human cruelty. As if the horror of these killings is not enough, what many psychological studies conclude, partly in explanation of how ordinary people can commit such overwhelming cruelty, is that situations determine actions. This view stands in contrast to the received view of human agency, namely that actions are explained by an agent's character. When someone, thus, acts rightly, this is due to their virtuous character; when someone fails to act rightly, this is due to a lack of virtue. But if it is the situation and its unique pressures that determines which action someone does (largely independent of character), it looks as though the only thing to keep the so-called "normal" life on the rails is the good fortune to avoid

situations that result in the commission of extraordinary cruelty. But that good moral luck provides only the most uncomfortably thin buffer between goodness and cruelty.

As a result of these many studies in human psychology (both academic and real life), some philosophers have limited their work to the study of the ordinary person's moral beliefs, finding perhaps that only this is justified. On this view, the work of moral philosophy consists in understanding the moral judgments that people make and the moral (and immoral) actions that people commit (and so moral philosophy becomes primarily a descriptive, rather than a normative, undertaking). Recommending moral ideals or arguing for moral requirements has no place in the moral philosopher's repertoire. But there is another and perhaps more important consequence of this psychological data: a number of philosophers have cast serious doubt on one version of moral philosophy, namely, virtue theory. Quite obviously, these philosophers argue, if it is the situation that determines the action, there would seem to be no value in studying, or proffering theories based upon, character. Indeed, taken to its extreme, one might conclude that there is little point even in trying to do better in one's own moral life: if actions are determined by circumstances that are beyond our control, there would seem to be nothing that one could do, as Iris Murdoch suggests, "to make oneself better."

Both of these conclusions are mistaken. We can, and should, attempt to make ourselves better. And, in a somewhat surprising twist, it may turn out that virtue theory offers a compelling explanation for how we can do so. Or, so I will argue.

2. An impressive body of psychological data casts serious doubt on our ordinary conception of the human agent, as someone who is responsible for her actions and an

appropriate subject of praise and blame. Consider Stanley Milgram's well-known obedience experiments. Milgram and his colleagues found that they could repeatedly induce ordinary individuals from all walks of life to administer (fabricated, but fully credible) electric shocks to an eventually unwilling test-taker. The shocks began at 15 volts and increased in a step-wise progression to 450 volts. Once the recipient of the shocks began to protest that the shocks were painful and "bothering his heart," the individuals giving the shocks were observed sweating profusely, wringing their hands, and pacing nervously prior to giving each shock. The subjects clearly perceived that they were harming the victim, yet even the victim's terrible screams and then his seemingly unconscious silence, did not induce them to stop. Two-thirds of the subjects fully obeyed the experimenter's polite but firm requests to continue administering shocks, and even repeatedly gave shocks of 450 volts until a second lab-coated experimenter interrupted the experiment and allowed the agonized subject to stop.

This experiment has been performed in different settings, with different population groups, and in different cultures; in each case, the findings are consistent: a substantial portion (typically more than two-thirds) of ordinary, decent individuals can be induced with a polite request to administer potentially-fatal electric shocks to an unwilling victim. (Citation)

Additional studies have examined other aspects of human moral behavior. Regarding helping behavior, for instance, researchers have found that finding or not finding a dime in a pay phone coin return appears to be the determinant of whether or not someone helps a person in need; and again, researchers have found that being in a hurry (or not) is a significant factor in determining whether someone helps a person in need; and again, the presence of other people (confederates of the experiment) who either help or do not help is a determining

factor in whether a naïve subject helps someone in need (she helps when the confederates help and does not help when they do not help). (Citations)

Many psychologists conclude, from these and countless other studies, that it is the situation and its unique set of psychological pressures, which determine what action an individual takes. This view has come to be known as situationism. These studies thus offer us a way to make sense of how ordinary individuals in Rwanda, Nazi Germany and the Confederate South (among many other places) can commit horrific acts of cruelty toward people who had once been their neighbors and perhaps even their friends: the individuals who commit the cruelty are not bad characters, not cruel people. Rather, they are ordinary individuals, people like you and me, who are overcome by their situations. Their kindness, their decency, their good character is overwhelmed by their situations, and like the subjects of Milgram's obedience experiments, they find themselves doing things they know to be wrong and could never have imagined that they would do.

On the basis of this data, some philosophers have concluded that there is no such thing as enduring character traits that themselves determine an individual's actions. John Doris, for instance, has argued that character traits are local, rather than global. An example of a global trait would be compassion. And evidence of its existence in someone's character, according to Doris, would be that the individual performs compassionate, helpful actions across a wide range of situations: the individual would be helpful when she is not in a good mood, and when she is in a hurry, and when others around her do not help. But it is precisely these conditions that are not met. Most individuals do not help when in a bad (or neutral) mood (24 out of 25 people who failed to find a dime in a pay phone coin return did not help a person who dropped

an armful of papers right in front of them – indeed, some of them stepped right over her papers as they went on their way), they do not help when they are in a hurry (only 10% of hurried seminarians helped an ill person in need as compared to 63% who helped when not in a hurry), and they do not help when others around them are complacent (across a whole variety of psychological studies, individuals almost never help a person in need when confederates of the experiment do not help; and, perhaps, the most infamous case of the failure to help is the 38 people who watched and did nothing while Kitty Genovese was stabbed three times (fatally, the third time) over the course of 35 minutes.). On the basis of these data, Doris concludes that, at best, human beings exhibit local traits: “dime-finding- compassionate,” “helpful-when-not-in-a-hurry,” or “helpful-when-others-around-are-helpful.” But, Doris, concludes, this is hardly what we mean when we attribute compassion or helpfulness to someone’s character. In addition to this critique of individuals, Doris and others have concluded, on the basis of these critiques, that virtue theory relies on an empirically unsustainable view of character.

Some virtue theorists have attempted a response to this critique. One response argues that the situationist’s critique employs a non-Aristotelian conception of character. So we need not be surprised that the individuals in these many studies did not act virtuously: they simply were not virtuous individuals. The truly virtuous person, the argument continues, would be able to think through the many factors in all these situations to arrive at the right thing to do. She would be able to consider the pressures to be obedient, the mood-elevating effect of finding a dime, or the lack thereof, the blinding effect of being in a hurry, the diffusion of responsibility of being in a group, and still always arrive at the correct judgment about what to do. (Cite Merritt, Kamtekar, Annas) Virtue, the argument goes, is hard; and we should not be

surprised that ordinary people often fail to do what is virtuous. But this failure does not impugn the ideal to which we can all nonetheless strive.

But how, we might ask, does the virtuous person do it? By what psychological mechanism does she arrive at the correct understanding of what is required in each situation? The virtue theorist cannot simply stipulate that the virtuous person does this. It cannot just be true by definition that she does it. Indeed, part of the situationist's critique of virtue theory is that it is precisely *surprising* factors, seemingly morally irrelevant factors—such as the finding (or not finding) of a dime—that support virtuous action (or undermine it). (Cite Doris) So, the situationist might urge, all of us, even the rare virtuous person, will sometimes be blind-sided by a seemingly morally irrelevant factor that turns out, in some situation, to be relevant, and thus influence or determine the agent's behavior. Because of this, it seems that it will be impossible for anyone to always get things perfectly morally right. Human psychology is too complex; the influences on our moral choices are too varied, indeed, too random, for any person, even a morally skillful one, to get things right all of the time. Thus, there can be no true virtue.

But the problem goes further still. Even if there were a few rare individuals who could get things just right all of the time, how would that help the ordinary person to do so? Moral theorists want, I presume, not only an account that explains how the morally perfect person gets things right, but we also want to know how the ordinary person can become perfect, or at least better, as Iris Murdoch suggests. The gulf between the ordinary person and the virtuous one seems to have become an unbridgeable chasm. No understanding of how the virtuous person arrives at the correct decisions can help the ordinary person to do so. Even in precisely

similar situations, their thinking, their judgment is so dissimilar, that no reflection on the virtuous person can help the ordinary one to become better. (Citation: Simon Keller)

The virtue theorist has, as yet, given us no story, no psychological process through which the ordinary person can develop morally so as to become (more) virtuous. So the situationist might conclude, what is the use of such a rarified notion of virtue, that it cannot be employed by the ordinary person who wants both to do the right thing and also to become good?

Before we can attempt to respond to these serious objections, there is an additional set of psychological data for which we must account.

3. It has often been thought that our moral attitudes are shaped by some combination of beliefs and desires. Some philosophers, non-cognitivists for instance, have emphasized the role of desires in our moral psychology, while others, cognitivists, have emphasized the role of beliefs. And still others hold the view that it is evaluative beliefs (sometimes pejoratively referred to as 'besires') that shape our understanding of the moral world. My present concern is not with which of these competing views is correct, but rather whether any of them, individually or as a group can explain all of our moral motivations. Psychological researchers are beginning to understand that some of our morally significant actions are motivated by implicit associations, or attitudes about which we are unaware, which we do not reflectively adopt and which conflict with many of the explicit beliefs we do hold. Because these motivational attitudes are subconscious they resist easy analysis into any of these three groups of motivations (beliefs, desires, or besires).

Implicit associations exist between evaluative and descriptive concepts and we are typically unaware of them, we do not choose them, and we cannot explicitly control them. In a manner reminiscent of Hume's association of ideas, someone has an implicit association when her mind links two distinct traits, because the mind perceives that the traits have often occurred together. Psychologists thus measure the ease with which subjects link the evaluative and descriptive concepts: if a subject links them more quickly, the researcher explains this by appeal to the subject's implicit association between the two concepts. Since the late 1990s, psychological researchers have been studying implicit associations, and gathering substantial amounts of information about them, especially as they concern age, race, gender, ethnicity, sexual orientation, ability, and other traits as well.

To study implicit associations, researchers have designed a number of tests, known collectively as Implicit Association Tests,<sup>1</sup> which measure how quickly a subject sorts words into categories. In one such test, researchers study a participant's normative associations (good and bad) with the race-concepts, 'Black' and 'White.' The Implicit Association Test (IAT) asks participants to sort words that flash on the computer screen by pressing one of two buttons on the computer keyboard. So in the first of several sequential tasks, a participant will sort words that appear in the center of the computer screen (words like joy, love, peace, agony, terrible, horrible, etc.) into one of two categories: "Good" (on the left hand side of the screen) and "Bad" (on the right). Then, in the next task, the participant will sort pictures of faces, some African-American, and others European-American, into two categories, this time labeled, "African-American" or "European-American." (And it should be noted that whether "Good"

---

<sup>1</sup> The Implicit Association Test (IAT) can be found at: <https://implicit.harvard.edu/>.



appears on the left or right and whether “African-American” appears on the left or right – and all the different possible combinations – is randomly selected by a computer program for each participant.) After these two preliminary tasks, the next task randomly pairs the evaluative concepts with the descriptive ones. So “African-American” might first be paired with “Good” and “European-American” with “Bad”; the choice on the left would read: “African-American or Good” and the one on the right would read: “European-American or Bad.” Now the same set of words and images would appear serially in the center of the screen. After the appearance of each word or image the participant would sort it into its appropriate category. And then in the final task, the pairings would switch: so “African-American” would be paired with “Bad” and “European-American” with “Good.” And once again the participant would sort the words and images in the center of the screen into their appropriate categories. For each sorting task, the individual is instructed to go as fast as possible but make as few mistakes as possible.

Because individuals can take IATs on the web, there have been hundreds of thousands of participants. These participants have been from all demographic groups in more than 35 countries. Together participants have completed more than 4.5 million sorting tasks.<sup>2</sup> It thus appears to be a powerful measure of the implicit associations that people hold.

In some cases these implicit associations also amount to biases, as when individuals consistently associate “African-American” with bad, or “European-American” with good. The IAT measures prejudice or bias because it hypothesizes that a subject who strongly associates “European-American” with “Good” will more quickly sort words and pictures into the category “European-American or Good.” Furthermore, in a subject who has an implicit association

---

<sup>2</sup> See <https://implicit.harvard.edu/implicit/>

between “European-American” and “Good” it will be more difficult, and thus take more time, to sort words and into the categories, “European-American or Bad,” and “African-American or Good.” What researchers have found is that most test-takers, regardless of race, display some implicit association between “European-American” and “Good” and between “African-American” and “Bad.” For some people these associations are stronger than others. Specifically, researchers have found that 70% of people who completed the race-based IAT exhibit some automatic or implicit association between European American and good/African-American and bad. (Only 12% of subjects exhibit an automatic preference for African American and good/European American and bad; and 17% display little to no automatic preference for either race.) Between 75 and 80% of self-identified whites and Asians exhibit some automatic preference for European American and good/African American and bad.<sup>3</sup> Furthermore, researchers have found that people from all social backgrounds have implicit associations or biases: “from young to old, male to female, Black to White, and conservative to liberal, implicit biases are not held by a select few but are readily observed among all social groups.”<sup>4</sup>

In addition to measuring implicit attitudes about race, the IAT measures other implicit attitudes as well. Another IAT, for instance, asks participants to sort names (“Ben,” “John,” “Daniel,” “Julia,” “Michelle,” “Anna,” etc.) associated with gender into two categories, “Male” and “Female.” Participants then sorted words associated with family (“home,” “parents,” “children,” etc.) and career (“management,” “professional,” “corporation,” etc.) into two categories, “Family” and “Career.” Researchers then combined the categories: so participants

---

<sup>3</sup> See <https://implicit.harvard.edu/implicit/>

<sup>4</sup> See Brian A. Nosek, Mahzarin R. Banaji, and Anthony G. Greenwald, “Harvesting Implicit Group Attitudes and Beliefs From a Demonstration Web Site,” *Group Dynamics: Theory, Research and Practice* 6: 1 (2002) p. 112

sorted the words and names into one of two categories: “Male or Family” and “Female or Career.” And then as with the other IATs, the pairings would then switch to “Female or Family,” and “Male or Career.” Researchers have found that many participants exhibit an implicit association between male names and “Career” and female names and “Family.” More specifically, in this sorting task, 24% of subjects exhibit a strong association between female and family and between male and career; 32% of subjects exhibit a moderate association between female and family and male and career; and 20% of subjects exhibit a slight association between female and family and male and career (a total of 76% of subjects exhibit some association between female and family and male and career). 17% of subjects exhibited little to no association at all, while only a total of 6.3% of subjects combined exhibited an association, of any strength (strong, moderate or slight) between male and family and female and career.<sup>5</sup>

To some it may be surprising to learn that such vast percentages of participants in these studies associate white race with good, black race with bad, female with home, male with career. And it can be even more surprising to learn that the individuals with these associations often lack awareness that they have these associations at all. Indeed, in many cases, the IAT is preceded by a questionnaire designed to measure explicit beliefs about the implicit attitudes that the test is about to measure. What researchers have found is that most individuals are not biased in their explicit beliefs even when the IAT then reveals a bias at the level of implicit attitudes. Indeed, in many cases, researchers find that the associations conflict with explicit beliefs that these individuals hold.

---

<sup>5</sup> All statistics here are from the IAT website: <https://implicit.harvard.edu/implicit/>

4. These findings are important on their own. Also important, however, is the correlation between a person's implicit associations and the actions she commits. Despite the fact that individuals are not aware of their implicit associations and do not choose to have them, these associations nonetheless seem to motivate actions. In other words, the IAT turns out to be a better predictor of certain sorts of actions than are a person's explicit associations or beliefs. One study, for instance, examined a long-existing disparity in the treatment of myocardial infarction (or heart attack) for black and white Americans.<sup>6</sup> White Americans are up to twice as likely to receive thrombolytic therapy for heart attack as are blacks. In this IAT study, 393 internal medicine and emergency medicine residents completed three tasks: first they read a vignette describing a 50-year-old male patient who presented to the emergency department with chest pain. Part of the vignette included an electrocardiogram indicating myocardial infarction. It also included a photograph: some of the physicians received a photograph of a black face and others received a photograph of a white face. But apart from the different photographs, the information sent to all the physicians was identical. After reading the vignette, participants rated the likelihood that the chest pain was caused by coronary artery disease, and said whether they would treat the patient with thrombolysis. The second task was designed to assess explicit bias. Participants answered several questions about whether they preferred black or white Americans, and whether they found black or white patients to be more cooperative with medical treatment. In the third and final task, participants completed three IATs designed to assess implicit racial bias.

---

<sup>6</sup> See Green, et al., "Implicit Bias among Physicians and its Prediction of Thrombolysis Decisions for Black and White Patients."

The researchers found that in their explicit beliefs, the participants expressed an equal preference for black and white Americans; they felt equally warmly toward black and white Americans; and they found black and white patients to be equally cooperative with medical treatment. Regarding their implicit preferences, however, participants showed moderate to strong implicit preferences for white over black; or in other words, participants showed a moderate to strong association between white and good, and between black and bad. Finally, regarding their actions toward the patients represented in the vignettes, physicians prescribed thrombolysis substantially more often for the white patients than for the black patients: “58.2% of physicians were very likely to offer white patients thrombolysis versus 42.7% for black patients.”<sup>7</sup>

What this study shows is that the physician’s explicit beliefs do not predict whether the physician will be more likely to prescribe thrombolysis. However the physician’s implicit associations do: “as the degree of antiblack bias on the race preference IAT increased, recommendations for thrombolysis for black patients decreased.”<sup>8</sup> The greater an individual’s association was between “white” and “good” and between “black” and “bad” the less likely was the individual to prescribe thrombolysis to the black patient. These prejudicial actions thus seem to be better predicted by the IAT than by the subject’s explicit beliefs.

A second study<sup>9</sup>, this time outside of the context of medicine, examined whether either performance on several race-bias IATs or answers to explicit questionnaires about race were correlated with discriminatory actions (one arm of the study examined an individual’s reports

---

<sup>7</sup> See Green, et al, p. 1235.

<sup>8</sup> Ibid.

<sup>9</sup> See Rudman and Ashmore, “Discrimination and the Implicit Association Test.”

about past discriminatory actions, such as telling a racist joke, or using verbal slurs, or even physically harming someone because of their race; the other arm of the study examined making potential budget cuts to Jewish, Asian and Black student organizations at a University as opposed to cutting the budgets of organizations like the drama club and the school band). Researchers again found that little or no prejudice was evident in the explicit measures. Performance on the IAT, by contrast, showed a strong correlation between “White and Good”/“Black and Bad;” and moreover, the stronger the association between “White and Good”/“Black and Bad” the more likely was the individual to report a history of past discriminatory actions, and the more likely the individual was to engage in economic discrimination (cut the budget of the Jewish, Asian or Black student organizations rather than the race-neutral organizations).<sup>10</sup> So again we see that it is the performance on the IAT rather than the response to a questionnaire that predicts the discriminatory actions an individual may take. Implicit attitudes seem, then, to motivate actions that are morally significant.

5. Now, recall the problem with which we began: we saw that situationists doubt that character-based moral theories can sustain themselves in the face of overwhelming psychological data, data that point to the conclusion that actions are the result of situations, not characters. People commit cruel actions when situations are such as to make cruelty the overwhelming option; people commit kind actions when situations give rise to kindness. But individuals are neither kind nor cruel: they simply do the actions that are dictated by their

---

<sup>10</sup> See Rudman and Ashmore.

circumstances. Thomas Nagel aptly called this kind of moral luck circumstantial luck, or luck in the sorts of moral situations that we face.

But we can now see that the problem goes further still. For not only is it our actions that are co-opted by our situations. But it is also our implicit associations, our motivations, which are determined by our environments and by the milieu of social norms and values in which we live. When the social norms and values of a society are not just, when that is, an individual lives in a racist, or sexist, or homophobic society it is very difficult to avoid absorbing racist, or sexist, or homophobic associations. Nagel referred to this sort of moral luck as constitutive luck, or luck in the sort of person you are. Constitutive luck is luck in both one's genetics and one's environment and here we can see that one's social environment may determine the sorts of moral associations or implicit attitudes that an individual has. This, in turn, is a reminder of Aristotle's point that an individual cannot be fully virtuous in the context of an unjust society. A lack of social justice harms us all, oppressors and oppressed alike, though it certainly does not harm us all equally.

So we can see that it is both our actions and also our unconscious moral attitudes that are subject to moral luck. This set of data thus suggest that many of our morally important actions look to be outside of our control: either because the actions are determined by the situations we find ourselves in or because they are motivated by attitudes about which we lack awareness and over which we lack control. This appears to make trouble both for our ordinary conception of the moral agent, as one who is responsible for her actions; it also appears to undermine moral conceptions that emphasize character, since our actions seem not to be the

result of our characters, but rather seem to be the collective result of our situations and our implicit attitudes.

One reply to these worries, which we encountered earlier, was that the subjects in the relevant psychological studies, including now the IAT, are not virtuous. They do not do the right actions nor have the right attitudes precisely because they are not virtuous.

This is a plausible reply. Indeed, some virtue theorists argue that the individuals in these studies are not virtuous in part because these individuals lack practical wisdom. Practical wisdom, according to Aristotle, is a sort of know-how: knowledge of how to act and feel and be motivated that is in accordance with virtue. So these individuals who fail to act correctly, and fail to feel the right emotions, and fail to be motivated correctly, it is argued, lack knowledge of what to do and how to feel that is characteristic of moral virtue. They lack practical wisdom because they individuals have not been able to practice virtuous actions and thereby cultivate both a love of the good and virtuous implicit associations.

But even if this is true, explaining the moral failure of these individuals by appeal to practical wisdom leaves it unclear how the moral novice can develop so as to emulate the virtuous person. (Maria Merritt 2009) In other words, even if the moral novice knows that she ought to cultivate practical wisdom, how can simply knowing that she ought to do this help her determine how to do it? The problem for the moral novice stems from the fact that many orthodox virtue theories ground “the criterion of right action in some conception of what an idealized yet fully human agent would do.” (47) In other words, the right-making feature of an action is ‘that the virtuous person would do it.’ Merritt argues that this right-making feature is so removed from the moral considerations that it is extremely difficult for the novice to



deliberate about what to do. Because most moral novices have no idea what a virtuous person is like, thinking abstractly about the virtuous person and what she might do in a situation will offer little or no guidance to the moral novice (that is, to most of us) about what to do in that situation.

Merritt's antidote to this problem is to suggest that the moral novice use a 'cognitive shortcut.' So rather than think about what the rarefied virtuous person would do, instead the novice ought to aim to do "what you expect your admired associates to approve of, using them as stand-ins for 'the person of practical wisdom.'" (47) But, as Merritt points out, this will not reliably lead to right action: sometimes even your admired associates will fail to do what is virtuous. So instead, Merritt thinks aspiring moral agents need systematic decision-making guidance. Perhaps Merritt thinks the moral novice needs a list of defeasible moral rules to follow, or a list of Kantian-inspired rules of moral salience (cite Barbara Herman).

According to this suggestion, moral agents internalize a set of moral rules throughout the process of moral development. These moral rules codify what is morally significant: the rules alert moral agents to the salient features of moral situations that the individuals can then act upon. Some rules of moral salience might include: "Suffering and pain ought to be avoided," "lying requires a moral justification," and so on. Kantians for example, argue that this set of moral rules enables agents to know when to deploy the categorical imperative and what information to include in a maxim. The virtue theorist, by analogy, could argue that these rules of moral salience enable the novice to see which actions to do, in order to cultivate that know-how that otherwise seems to remain so elusive to the moral learner.

But why think that the moral novice needs a list of, even defeasible, moral rules? Certainly moral agents must learn to be attuned to moral significances, things like suffering, pain, lying, and so on. But if the rules of moral salience are too specific, as in “lying is always wrong,” the rules risk being false or giving inaccurate advice about the moral landscape; on the other hand, if the rules are too general, as in “lying is morally significant,” then they are true, but neither useful nor edifying to the moral novice.

To address this problem, some have argued that virtue theory requires an account of moral commitment (Cite Bakhurst). On this view, a process of moral development enables a moral novice to know that some things like pleasure and pain, suffering and death are morally significant. But the way in which they are significant can vary from context to context. Sometime the fact that an action would cause suffering makes it to be avoided; but other times it does not. More generally, the fact that some action causes suffering will typically be significant; but it need not always be significant nor need it always lead towards or away from action apart from a specific context.

The virtue theorist need not, then, make recourse to rules, understood as exceptionless principles that always lead towards or away from action (the view known as atomism about reasons). She can instead embrace the idea of moral commitments: the idea that certain features of the moral landscape are typically morally significant, though they do not always require the same sort of action (the fact that something is a lie, for instance, is not always a reason not to do it).

Furthermore, we need not hold, as Merritt does, that the virtuous person herself, represents or embodies the criterion of rightness. That is, Merritt argues that what determines

which actions or feelings are morally right is the fact that they are what the virtuous person would do or feel. On this view, what makes an action or feeling or motivation morally required is the fact that the virtuous person would do it, or feel it, or be motivated by it. But an Aristotelian virtue theory has another option: Aristotle holds that “virtue is a state of character, concerned with choice, lying in a mean, determined by *orthos logos*” (cite) that is, determined by the right reasons “and by those reasons by which the person of practical wisdom would determine it” (cite). So it is important to see that these are separate criteria of rightness. The first criterion Aristotle mentions is the *orthos logos*, the right reasons. The second criteria he mentions is the virtuous person him or herself. Now of course, the feelings and actions each criterion recommends will always be co-extensive. But the point is that we need not look only to the person of practical wisdom to determine what to do and how to feel. And ironically, a non-character-based criterion of rightness may enable the virtue theorist to resist the move to a theory that employs a systematized Kantian-inspired decision-procedure in the form of rules of moral salience.

But now how do we determine what the *orthos logos* are? In response, here is one of Aristotle’s descriptions of the Doctrine of the Mean:

Moral virtue... is concerned with passions and actions and in this there is excess, defect, and the intermediate. For instance, both fear and confidence and appetite and anger and pity and in general pleasure and pain may be felt both too much and too little, and in both cases not well; but to feel them at the right times, with reference to the right objects, towards the right people, with the right motive, and in the right way, is both intermediate and best and this is characteristic of virtue. Similarly, with regard to actions also there is excess, defect and the intermediate. (38)

Virtue requires that we feel our emotions and perform our actions in the right amount, at the right times, with reference to the right objects, toward the right people, with the right motive and in the right way. And if we do so correctly, meeting all these criteria, we have hit the mean. What makes a certain path the correct one is not that the path of action and feeling lies in some absolute mean between the excesses and deficiencies. What makes a certain path of feeling and action correct also need not be that it is the path the virtuous person, the person of practical wisdom, would follow. Though, because she is virtuous, the virtuous person will follow that path. Rather what makes a certain path the correct one is that it is the path that corresponds to the proper answer to each of Aristotle's questions. A path of emotion is correct when the individual feels the right emotions, in the right amounts, at the right times, with reference to the right object, toward the right people and in the right way. Her action is correct in the same way, when she does the right action, at the right time, with the right motive, in reference to the right object, toward the right person and in the right way. The right making features of the moral situation, then, are neither the rules of moral salience, nor the virtuous person herself; rather the right making features are the details of the situation, the particulars that correspond to each of these injunctions. To find the mean in each situation is to find what is called for regarding each of these *logos*.

So for example, how much sadness at the death of a loved one ought someone to feel? It depends on the magnitude of the loss that he has suffered. What sort of relationship did he have and how will the loss of it affect his life? It also depends on the details of his psychology and social situation: if he is a resilient person with a lot of supportive friends, some of his sadness may be mitigated. As Aristotle tells us, the mean is relative, or non-absolute in two

senses: it is relative both to the details of the moral situation and it is relative to the moral agent. If someone tends to get overly emotional at such losses, perhaps he should aim for less emotional expression and feeling so as to cultivate the appropriate amount of emotion.

Or again, consider whether someone should stop and help a person in need or continue on to her destination. Which action is virtuous will again depend on the details of the situation she is in: to what destination is she headed? To perform surgery? To pick up her sick child at school? To eat lunch with a friend? And again, it will depend on how grave is the need for help. Is another person at hand who is ready and able to provide the help as well or better than she can? What she ought to do will depend on all these details and will depend again on her psychology. If she tends to offer help too often, and the need is not grave, perhaps she ought to hold back, so as to avoid burnout. Again, in doing so, she cultivates the proper amount of helpfulness in her character.

But suppose, as was the case for some of the subjects in the studies we have considered, the individual is in a hurry or in a bad mood (or a non-elevated mood, because she did not find a dime in the coin return) and so she does not even notice that there is a person here in need of help? Or suppose the person notices vaguely that someone is in need but does not fully internalize their need: their need does not touch her or move her. Aristotle reminds us that, moral matters “depend on particular facts and the decision rests with perception.” What this oft-quoted passage tells us in this context is that we cannot act virtuously unless we first recognize the moral features of situations that call for a response. No set of rules of moral salience, no amount of systematic decision-making guidance can compensate for poor moral perception. This is, of course, because we cannot employ the decision-making guidance until

we have seen what is called for. The same is true of reflection on the person of practical wisdom: no amount of such reflection can determine for us those features of situations to which we ought to be attuned. But once we have cultivated excellent, fully accurate moral vision, the need for systematic decision-making guidance drops away, as does the need to reflect on what the virtuous person would do: we will simply see which features of the situation call for a response and which do not. To see what is called for is to know how to act and feel. (Citation: others have made this point.)

But how do we cultivate this excellent perception? Unfortunately, there are no shortcuts, cognitive or otherwise. The only way to cultivate excellent moral perception is to learn, one by one, to be sensitive to features of situations that do or do not make a moral difference. Sometimes being in a hurry means that we have no moral obligation to help; and sometimes it means just the opposite. Learning to discern the difference is precisely what is involved in the skill of morality. Sometimes being alone means that we are obligated to provide help, and sometimes it means that we are not obligated to provide the help. And there is no way to tell which it is until we find ourselves in the moral situation itself. Now certainly, prior to being in such a situation, the novice can and should learn about diffusion of responsibility and all the other moral particulars. But she will not be able to know what precise affect any of them will have until she faces a concrete moral situation.

6. To this point, I have argued for a kind of particularism. Specifically, I have argued that the virtue theorist need not make recourse to rules of moral salience, or rules that codify precisely how certain moral features will be significant. Instead, I have argued that a virtue

theorist can make use of the idea of moral commitment, the idea that moral agents should be attuned to certain moral significances (suffering, death, pain, deception, etc) but that the precise way in which these particulars will be significant cannot be known until the concrete reality of a moral situation is at hand.

At this point in the particularist's argument, many remain unconvinced, indeed baffled at the prospect of learning about morality one particular at a time. Recall that we said above that human psychology was enormously complex, and that psychologists are only at the beginning of coming to an understanding of it. How can we learn about all the details of human psychology, all the moral requirements, one by one, with no rules, no patterns to be discerned at the non-moral level to help us?

Despite the fact that this process looks overwhelming in its complexity, nevertheless, human beings do just this. We do it not only with morality, but in many contexts. Consider, for example, what is involved in learning to read (a process that begins shortly after birth and continues for, perhaps, a couple of decades). The young school child already well into the process, learns to recognize letters, learns that letters have distinctive sounds and learns to associate each letter with its characteristic sound or sounds. When the child first grapples with printed words, she recognizes that letters constitute words: 'c', 'a', 't' make up the word 'cat'. She learns to sound out these letters until she associates these letters, in just this arrangement, with the word 'cat'. After many instances of encountering just this string of letters, she finally learns that 'c', 'a', 't' represent the word 'cat'. Subsequently she just sees 'cat'. She goes on to do this for the thousands and thousands of words in her language. No doubt she will become more proficient at this over time and do it more quickly. Learning to read is an enormously

complex task. Yet, human beings discern each word, one by one, and remember it. And we do this thousands and thousands of times; we do it for each word in the language, learning one by one to associate the string of letters in that word with the word itself (and the word with the object or concept to which it corresponds).

Moreover, there are no exceptionless rules to help us with this task of enormous complexity. Consider for example just one particular with which the beginning reader must grapple. Typically 'g' at the beginning of a word is pronounced. So the beginning reader will learn that 'g' in 'girl,' 'great' and 'goose' is pronounced. But then he sees the word 'gnome' and initially cannot recognize this word because he tries to pronounce the 'g' at the beginning. But then he learns that *this* 'g' is silent. When he then encounters the word 'gnu' he might again think that this 'g' is silent. In one pronunciation of the word (referring to the elk-like animal) the 'g' is not pronounced; but in another pronunciation of the word (referring to a computer operating system) it is. So there is no exceptionless rule that will help the beginning reader to know what to do with a 'g' at the beginning of a word. Fortunately, the human brain is so complex that he does not need a rule: he simply learns the words one by one, and does this for every word in his language. He may even go on to do it for multiple languages.

I want to suggest that this is not unlike how we learn the moral significances of countless moral particulars. It may be that morality is more complex or again, it may be that there are fewer moral particulars to learn than there are words in a language; but whichever way it turns out, we know that we need more studies in psychology to help us more fully discern the particulars to which we must be attentive. Practical wisdom or know how just is the



virtuous person's skillful response to the myriad, yet ordinary, moral particulars we encounter every day and to which we cultivate a sensitivity.

How can the novice cultivate this sensitivity? There are no shortcuts. We learn about morality situation by situation, and much of this learning is implicit now in the new sense that we notice, watch, and observe moral significances from a very young age. From infancy we notice pleasure and pain, hurt feelings and healed ones, both in ourselves and in others. As toddlers and preschoolers, and hopefully continuing on into old age, we learn what sorts of actions bring pleasure and what sorts of actions cause pain. We learn what hurts feelings and what heals them. And depending on our teachers, our role models, our parents, and our peers, and what they all teach us, and model for us, we can become more or less sensitive to these important moral features of situations. And as we grow, we learn that sometimes, something we thought would bring happiness, actually caused pain, and so we see that the moral landscape is complex. But one by one we learn to be sensitive to all these nuances and complexities. That is the path, and I would say, the only path to virtue and practical wisdom, or the know-how that we need to respond appropriately to moral situations, one at a time.

We are now in a position to respond to the situationist's critique with which we began. We saw that the situationist argues that it is not characters, but situations, with their unique set of pressures, that determine an individual's actions. As a result, the situationist believes that virtue theory is in an especially weak position, because virtue theory relies on a notion of character that is simply not empirically supported. But we can now see that the virtue theorist has, paradoxically, a powerful way to resist the situationist's critique. The virtue theorist holds that individuals can cultivate their characters to be responsive to all of the data that empirical

psychology discovers about human agency. So when psychologists learn about the affects of finding a dime on helping behavior, the moral novice can learn to be sensitive to this fact; when psychologists learn about the diffusion of responsibility, the novice can learn to be sensitive to this fact; and when psychologists learn about the way in which implicit associations motivate biased actions, the novice can learn to be sensitive to this fact. In short, as psychologists reveal morally significant facts, one by one, we moral novices must learn about them and come to be sensitive to them. Virtue theory, precisely because it emphasizes character and moral development has a ready way to accommodate the findings of empirical psychology: individuals can develop their characters so as to be responsive to the relevant facts about human moral agency.

7. Learning to be virtuous, indeed, learning about morality, is a process that begins at birth and grows as our awareness of the world grows. It is a process of learning the significances of the things that matter to human beings: hurt feelings, skinned knees, losses through death, injustice, personal accomplishment, and so on. And it is a process of learning how to respond to each of these significances: learning how to feel and act with respect to each of them and with them in combination. It is complex and detailed and nuanced. But it also ordinary: morality is implicit in our daily lives. We simply do notice these significances, for some we do this right from birth and for others we cultivate a sensitivity through our moral development. And if we are fortunate to have skillful parents, teachers, and mentors, as well as a desire to introspect and grow, our moral sensitivity will continue to develop throughout our lives.

Morality is an ordinary practice, as ordinary as learning to read. And like reading, it is a skill that requires time and effort to acquire. Just as we learn to read one word at a time, we learn about morality one moral significance at a time. But unlike reading, we need more studies in psychology to reveal the precise contours of the moral landscape, the precise nature of the particulars to which we must be attuned. And only when we have a rich and detailed understanding of human psychology can we fully understand all the ways in which we can make ourselves better.

But despite these advantages of virtue theory over theories that emphasize action over character, the response of the virtue theorist to the situationist critique can only go so far. Even though virtue theory is in some ways better situated to respond to the situationist's critique, moral luck remains a recalcitrant problem. Even if individuals cultivate their characters so as to respond to psychological facts about human agency, still our actions will be largely or perhaps wholly determined by circumstantial and constitutive moral luck. Our constitution, both genetic and environmental, alongside the particular sorts of moral problems we face will continue to determine the actions that we do. This problem threatens to erode not only the view of the virtue theorist, but more sweepingly the view of the *moral* theorist. So while I have argued that the virtue theorist's view is better situated to respond to the situationist's critique, precisely because it emphasizes character and moral development, it remains to be seen whether any moral theory is adequately situated to respond to the larger problem of moral luck. But that is a question for another time.