

## 5

## Virtue Theory Against Situationism

*Ernest Sosa*

Why did Tom give up his seat on the bus to someone frail and elderly? Perhaps only to impress his girlfriend, perhaps rather to be considerate, out of concern for the elder's welfare. And if the latter, maybe it was just a random act, entirely out of character, and due more to his being in a good mood. Alternatively, it might be quite in character for him to act kindly in that way. Kindness may be one of his character traits, manifested in that act.<sup>1</sup>

Or so one might think without a second thought. Any such scenario would seem a commonplace example of how we constantly try to understand people's conduct. Based on a body of troubling results in social psychology, however, an intriguing critique has been pressed in recent years against such virtue ethics, raising doubts both about its moral psychology and about its normative content. Here we shall review some of the most striking, best-known results and the arguments based on them.

Similar discoveries have been made by social psychologists about our belief management, moreover, so that a similar critique can be pressed against virtue epistemology. I here come to this topic through my earlier virtue-theoretic approach to epistemology, allied to similar approaches to ethics and even to parts of aesthetics. We focus first on virtue ethics and its correlated virtue psychology, and then turn to virtue epistemology. The logical structure of our response to the critique of virtue ethics is closely replicated by a response available to the virtue epistemologist. Accordingly, rather than detail that obvious analogy, the space available in this chapter will be devoted to an account of epistemic virtues or competences. That view is then defended briefly against its situationist critique by close analogy with how virtue theory prevails in ethics against a similar attack.

<sup>1</sup> Virtue ethics goes back to Aristotle and to Ancient Greek philosophy more generally. Long neglected in the shadow of deontological and utilitarian approaches despite its powerful advocacy by Hume and others, it has in recent decades regained much of its former luster and influence. The approach has both an empirical side and a normative side. It appeals to virtues or to virtuous traits of character in both the explanation and the assessment of human action. That an action manifests a virtue bears positively in its overall assessment, that it manifests a vice bears negatively.

## 5.1 The Situationist Critique of Virtue Ethics

### 5.1.1 *The Milgram experiments*

In the early sixties,<sup>2</sup> experiments conducted by the psychologist Stanley Milgram at Yale University had disturbing results.<sup>3</sup> In multiple replications, moreover, the results have held up with impressive consistency. Milgram's subjects believe themselves to be playing the role of "teachers" in a study of the effects of punishment on learning. Here is the scenario. The participants are "teachers" expected to administer electric shocks to "learners" (who are in fact Milgram confederates). In one version, the teacher sees the learner/confederate strapped down to an electric chair in a separate room. The learner/confederate asks about the shocks and is told that they are not harmful but can be painful. The learner/confederate says that he suffers from a slight heart condition, asks whether the shocks would be dangerous to him, and is assured that they would not be dangerous, although, again, they could be quite painful. When the experiment begins the teacher is given a sample 45-volt shock from the machine, just to add realism. Then he is taken to a position in another room from which the learner/confederate is no longer visible. As the experiment proceeds, the learner/confederate repeatedly fails to answer the questions correctly, so that the teacher is bound to keep increasing the strength of the shocks in 15-volt increments. What were the results?

At 300 volts, the learner/confederate would pound on the wall, scream, and then, at 330 volts, would stop responding. Yet most teachers continued to intensify the shocks in 15-volt increments all the way up to 450 volts. That was done by twenty-six of the forty teachers, or about sixty-five percent. That means *ten* further voltage boosts *after* the pounding and screaming! This is what was done by twenty-six of the forty teachers. As for the remaining fourteen, these all went up to at least 300 volts, and stopped somewhere between 300 volts and 450 volts. If a teacher protested to the experimenter, he got one or more of a standard set of responses, in the following sequence: "Please continue"; "the experiment requires that you continue"; "it is absolutely essential that you continue"; "you have no other choice, you must go on." Again, *all* teachers administered shocks up to the severe 300 volts, and *sixty-five percent* went beyond that to shocks of the maximum 450-volt severity.

### 5.1.2 *The Good Samaritan experiment*

A second much-cited experiment was conducted at the Princeton Theological Seminary.<sup>4</sup> Seminarians were read either something about vocational choices available

<sup>2</sup> Sections 5.1–3 derive (with CUP permission) from my (2009) *Situations Against Virtues: the Situationist Attack Against Virtue Theory*. In C. Mantzavinos (ed.) *Philosophy of the Social Sciences*. Cambridge: Cambridge University Press, 274–91. The contents of the remaining sections are previously unpublished.

<sup>3</sup> Milgram, S. (1963). Behavioral Study of Obedience. *Journal of Abnormal and Social Psychology* 67: 371–8. Also his (1974) *Obedience to Authority: An Experimental View*. New York: Harper and Row.

<sup>4</sup> Darley, J. M. & 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: 100–8.

to them, or else the Parable of the Good Samaritan. They were then told to go to a certain nearby building, where they were expected to give a talk. Some were asked to hurry, since they were already late, others were just told to proceed without delay, and the rest were told that they had a few minutes to spare. On the way to the next building they all came across a figure slumped over in a doorway, groaning and coughing. What mattered to whether they stopped to help was essentially just how much of a hurry they were in. Subjects' offers to help were positively correlated with their beliefs about how pressed they were for time. Sixteen out of the forty subjects offered to help. Of the eight with some time to spare, five stopped to help. Of the twenty-two who had been told to go right over, ten stopped. Finally, of the ten who had been told they were already late, only one stopped.

True, the percentage of subjects who offered to help had some positive correlation with the passage they had just read, but the correlation was moderate: of those who had just read the parable, fifty-three percent offered to help, but of those who had just read the vocational passage, only twenty-nine percent offered to help.

In conclusion, the experimenters and authors suggest, "only hurry was a significant predictor of whether one will help or not." For most subjects, punctuality trumped the evident need of someone in distress.

1. Based on the results of these experiments and several others, an attack has been launched on virtue psychology and virtue ethics.<sup>5</sup> The critics have raised questions of two sorts. First, they have challenged the notion that humans vary significantly in possessing traits of character—some virtuous, some vicious—*important for the explanation and prediction of human action*; and they have challenged also the normative ideal of human virtue held up by virtue ethics, for the reason that humans are very unlikely to guide their conduct by any such ideal.

Leading the situationist charge within philosophy are Gilbert Harman and John Doris, whose views we consider next.<sup>6</sup>

#### 5.1.2.1 HARMAN'S CASE

According to Harman, empirical testing has found no relevantly different character traits to account for behavioral differences. Yes, ordinarily it is supposed that people do differ relevantly in traits and virtues.

<sup>5</sup> To appreciate the extensive relevant literature, see the masterful (2002) *Lack of Character: Personality and Moral Behaviour*. Cambridge: Cambridge University Press by John Doris. The three studies I cite (two here and one in the concluding footnote of section 1) are the best known and the ones I have found most striking; the philosophical issues emerge fully, as I see it, on the basis of these three studies, which are also the most cited in the relevant philosophical literature.

<sup>6</sup> Nisbett, R. E. & Ross, L. (1980). *Human Inference: Strategies and Shortcomings of Social Judgment*. Englewood-Cliffs: Prentice-Hall. Ross, L., & Nisbett, R. E. (1991). *The Person and the Situation: Perspectives of Social Psychology*. New York: McGraw-Hill. Doris, J. (1998). Persons, situations, and virtue ethics. *Notus* 32: 504–30. Harman, G. (1998). Moral Philosophy Meets Social Psychology: Virtue Ethics and the Fundamental Attribution Error. *Proceedings of the Aristotelian Society* 99: 315–31. An earlier, softer challenge is due to Owen Flanagan, in his (1991) *Varieties of Moral Personality*. Cambridge, MA: Harvard University Press.

We ordinarily suppose that a person's character traits help to explain at least some things that the person does. The honest person tries to return the wallet because he or she is honest. The person who pockets the contents of the wallet and throws the rest of the wallet away does so because he or she is dishonest.<sup>7</sup>

However, people might behave differently, regularly so, without differing in character traits. The difference in behavior might of course derive rather from situational differences. In order to differ in character traits, people must be disposed to act differently though similarly enough situated (or similar enough in how they view their respective situations). *As ordinarily conceived*, moreover, traits are dispositions to issue the trait-relevant conduct across a *broad* range of relevant situations. True honesty, for example, requires honest conduct across a broad-enough range of relevant situations. It will not be enough that one be honest in forbearing to shoplift although one cheats on tests, on one's income tax returns, and in returning change.

Harman joins Nisbett and Ross, moreover, in distinguishing traits from sustained goals or strategies, and also in finding us too often guilty of the "fundamental attribution error," the error of attributing a trait based on too paltry an evidential basis. According to Nisbett and Ross, "individuals may behave in consistent ways that distinguish them from their peers not because of their enduring predispositions to be friendly, dependent, aggressive, or the like, but rather because they are pursuing consistent goals using consistent strategies, in the light of consistent ways of interpreting their social world."<sup>8</sup>

Harman comments as follows on our two striking experiments:

The fundamental attribution error in [the Milgram]... case consists in "how readily the observer makes erroneous inferences about the actor's destructive obedience (or foolish conformity) by taking the behavior at face value and presuming that extreme personal dispositions are at fault."<sup>9</sup>

Standard interpretations of the Good Samaritan Experiment commit the fundamental attribution error of overlooking the situational factors, in this case overlooking how much of a hurry the various agents might be in.<sup>10</sup>

And he sums up his brief:

We very confidently attribute character traits to other people in order to explain their behavior. But our attributions tend to be wildly incorrect and, in fact, there is no evidence that people differ in character traits. They differ in their situations and in their perceptions of their situations. They differ in their goals, strategies, neuroses, optimism, etc. But character traits do not explain what differences there are.<sup>11</sup>

<sup>7</sup> Harman, *Moral Philosophy*, section 2.

<sup>8</sup> Nisbett and Ross, *Human Inference*: 20.

<sup>9</sup> Harman, *Moral Philosophy*, section 5.1. Here Harman is agreeing with Nisbett and Ross.

<sup>10</sup> Harman, *Moral Philosophy*, section 5.2.

<sup>11</sup> Harman, *Moral Philosophy*, section 8.

## 5.1.1.2.2 DORIS'S CASE

Situationism for Doris involves three main commitments, concerning behavioral variation, the nature of traits, and trait organization in personality structure:

- (i) Behavioral variation across a population owes more to situational differences than dispositional differences among persons. Individual dispositional differences are not as strongly behaviorally individuating as we might have supposed: to a surprising extent we are safest predicting, for a particular situation, that a person will behave pretty much as most others would.
- (ii) Empirical evidence problematizes the attribution of robust traits. Whatever behavioral reliability we do observe may be readily short-circuited by situational variation: in a run of trait-relevant situations with diverse features, an individual to whom we have attributed a given trait will often behave inconsistently with regard to the behavior expected on attribution of that trait . . .
- (iii) Personality structure is not typically evaluatively consistent. For a given person, the dispositions operative in one situation may have a very different evaluative status than those manifested in another situation—evaluatively inconsistent dispositions may “cohabit” in a single personality.<sup>12</sup>

Situationism is not a Skinnerian evisceration of the person. While rejecting cross-situationally robust traits, the situationist admits local, situationally specific traits that distinguish people from one another. These traits are “local” rather than global and frail rather than “robust”: they do not reliably result in the same trait-relevant conduct across a variety of different situations.

At bottom, the question is whether the behavioral regularity we observe is to be primarily explained by reference to robust dispositional structures or situational regularity. The situationist insists that the striking variability of behavior with situational variation favors the latter hypothesis.<sup>13</sup>

Doris sums up as follows.

To summarize: According to the first situationist thesis, behavioral variation among individuals often owes more to distinct circumstances than distinct personalities; the difference between the person who behaves honestly and the one who fails to do so, for example, may be more a function of situation than character. Moreover, behavior may vary quite radically when compared with that expected on the postulation of a given trait. We have little assurance that a person to whom we attributed a trait will consistently behave in a trait-relevant fashion across a run of trait-relevant situations with variable pressures to such behavior: the putatively “honest” person may very well not consistently display honest behavior across a diversity of situations where honesty is appropriate. This is just what we would expect on the second situationist thesis, which rejects notions of robust traits. Finally, as the third thesis suggests, expectations

<sup>12</sup> Doris, *Persons, Situations*: 507.

<sup>13</sup> Doris, *Persons, Situations*: 508.

of evaluative consistency are likely to be disappointed. Behavioral evidence suggests that personality is comprised of evaluatively fragmented trait associations rather than evaluatively integrated ones: e.g., for a given person, a local disposition to honesty will often be found together with local dispositions to dishonesty.<sup>14</sup>

In brief: *First*, behavioral variation is due more to situational variation than to trait variation. *Second*, traits are frail across situational variations, not robust. *Third*, traits do *not* integrate into coherent characters.

2. Attentive reading reveals that both Harman and Doris reject Skinnerian nihilism on behavioral dispositions. Indeed, both of them *believe that there are traits*, dispositional traits, operative in human conduct generally. *What they deny is that these are traits as conceived of by the folk, or by the Aristotelian tradition of virtue ethics.* Harman does consider the possibility of rejecting traits altogether, even those that are local and frail as opposed to the global and robust. On this more extreme view, human conduct is to be explained not by traits, but perhaps through goals, policies, or strategies. But how are we to distinguish *learned* traits generally from such goals, policies, or strategies? The supposed alternative does not clearly differ more than verbally. Suppose have a firm goal to treat others politely, and I give substance to that goal through my knowledge of what politeness requires in a broad range of situations. How importantly does this differ, if at all, from possessing a trait of treating others politely?

Among people, traits can be rare and distinctive or, alternatively, vulgar and widely shared. They can be local (or narrow, or frail), moreover, or else global (or broad, or robust). *Our traits, insofar as we have any, are said to be vulgar and local, or at least much more so than is usually supposed.* This is what situationism seems to boil down to, apart from the claim that human personality is normally fragmented and falls far short of the integration proper to Aristotelian practical wisdom.

3. Situationists do agree with the folk, and with the tradition of virtue theory, on one important point:

#### *Variation*

There is substantial evaluatively interesting variation in human behavior. People can and do behave variably regarding honesty, kindness, courage, temperance, etc.

This much is in keeping with the experiments, and in line with situationist writings. What explains such variation? For the explanation of *cross-personal* variation we must invoke situational differences and distinctive traits; vulgar traits will not help. Nor can *cross-situational* variation be explained by *stably* vulgar traits. If character traits are stably vulgar, then evaluatively relevant behavioral variation must be situationally explained.

Consider now this question: What explains evaluatively interesting human behavior? This is *not* the question broached a paragraph ago. What explains cross-personal or

<sup>14</sup> Doris, *Persons, Situations*: 508–9.

cross-temporal variation in behavior need not be the same as what explains the behaviors severally. Compare this: The *differences* in the rolls of two round balls cannot be explained (at all) by appeal to their roundness, since they are the same in that respect. Yet, either roll might still be explicable (largely, importantly) through the roundness of the rolling object. Similarly, behavioral *variation* may not be explicable by appeal to traits, while still the behavior itself *is* so explicable, *even if* the traits are vulgar.

So, stable vulgarity will spoil a trait for the explanation of behavioral variation, but not for the explanation of behavior itself, whether individual behaviors or behavior patterns.

What about narrowness? How if at all does the narrowness of a trait impair its explanatory efficacy? A narrow or local trait is one that yields its evaluatively relevant behavioral outputs in a relatively narrow or local set of circumstances. Dispositions come of course in degrees: not only simple dispositions such as fragility and flexibility, but also those more relevant to ethics such as honesty and kindness. Accordingly, it is possible to explain the breaking of a vase by appeal both to its impact and its fragility, even if a fine wine glass is *more* fragile and would have broken not only with that impact but *also* with others that would *not* have affected the vase. Moreover, narrow, local traits of honesty (in returning change, say) may amount to ways in which one can have a low degree of honesty (since one is not also honest in filling out one's income tax, in taking tests, etc.). Consider the virtue psychology accepted as common sense by the folk. This is content to postulate varying degrees of its recognized virtues, though it may be surprised to see just how much we nearly all fall short, and also the ways in which we fall short.<sup>15</sup>

4. Thus far we have focused on the situationist attack on virtue ethics and psychology, led in philosophy by Harman and Doris, and on the proposed situationist alternative. For their part, advocates of virtue theory have converged on a response to that attack. Several authors have now accused situationists of adopting a crude, external, behaviorist conception of virtue psychology—one that virtue theorists reject as a caricature. Situationists are said to ignore the inner deliberative complexity so important to sophisticated virtue ethics and psychology.<sup>16</sup>

We are thus presented with two conceptions of virtue psychology: (a) Crude virtue psychology (CVP) focuses directly on situation/behavior dispositions. (b) Sophisticated virtue psychology (SVP) interposes situation/attitude dispositions between situations and behavior.

<sup>15</sup> This material was presented at the Herdecke conference on philosophy and the social sciences where the commentator, Steven Lukes, pointed out how surprising such surprises would be given ethnic cleansing and other familiar horrors of recent history and current affairs.

<sup>16</sup> Relevant here are four illuminating articles: Kamtekar, R. (2004). Situationism and Virtue Ethics on the Content of Our Character. *Ethics* 114: 458–91. Webber, J. (2006). Virtue, Character, and Situation. *Journal of Moral Philosophy* 3(2): 193–213. Sreenivasan, G. (2002). Errors About Errors: Virtue Theory and Trait Attribution. *Mind* 111: 47–68. Hursthouse, R. Virtue Ethics. *Stanford Encyclopedia of Philosophy*.

According to SVP, it is the agent's character that holistically explains his conduct. In contrast with the CVP attacked by Doris, Harman, and Nisbett & Ross, the character of interest to SVP is a broader whole that includes fundamental motives, desires, and even goals. The rational agent works to integrate these into a coherent whole. Given how often we are evaluatively conflicted—how often we must make hard choices—it can hardly be a surprise that we fail to be cross-situationally consistent at the level of external situations and attendant behaviors. The consistency is found, rather, internally, in the complex inner structure that is one's relevant character.

In order to possess the virtue of kindness, for example, one need not behave kindly whenever one is in a kindness-relevant situation. And the same goes for honesty and other traits. When values conflict in a situation, a value other than kindness may take priority. The virtues of SVP are rational virtues manifest primarily in right choices made through proper rational deliberation.

Here now is a way to develop this alternative picture: When deliberating on a yes-or-no choice, one faces a rational structure of pros and cons, of reasons for and reasons against. Here I mean *good, factive* reasons. These constitute the rational structure of the situation. One could think of this as a one-dimensional vector space with positive and negative vectors as the pro and con reasons. Additional options, beyond our simple yes-or-no case, will of course import a more complex vector space.

Corresponding to such a rational structure more or less well is one's motivational structure constituted by positive and negative motivating reasons, reasons that psychologically attract one to a certain choice or repel one from that choice, to various degrees.

Take our young man comfortably seated when the elder approaches on the bus. Among the factors that structure his situation rationally are her risk of falling and her evident physical and emotional suffering, whereas he is about as well off standing as sitting.<sup>17</sup> Indeed, we can surely shape the scenario so that the balance of reasons *strongly* favors his ceding his seat. Practical wisdom would then require that his motivating reasons reflect the objective reasons, so that the relative weights of his motivating reasons correspond to the relative weights of the relevant objective reasons. If the reasons to cede constituted by the risk and suffering of the oldster outweigh the reason not to cede constituted by his very slightly greater comfort, then the motivating force of those weightier reasons should also psychologically outweigh the motivating force of the less weighty reason, and should do so by a corresponding margin. Thus, the motivational structure in the mind of the agent should reflect the rational structure of the situation faced.

One manifests practical wisdom in any given situation to the degree that one's motivational structure reflects the relevant rational structure in that situation.<sup>18</sup>

<sup>17</sup> Better off, in fact, according to recent research on the steep health cost of our sedentary ways.

<sup>18</sup> Here I am assuming that the subject has access to the relevant plain facts. But it is a nice question of how extensively factual perceptiveness constitutes practical competence, and is detachable from value perceptiveness. Practical competence is a function not so much of mere factual perceptual acuity (of the ability to perceive sharply) as of the foregrounding of facts that are relevant good reasons for what is objectively

Of course, one cannot thereby *manifest* practical wisdom unless one *possesses* such wisdom, which is not something one can do only ephemerally and locally. Rather, one has true practical wisdom (a) to the degree that one has a stable disposition to reflect in one's motivational structures the pertinent rational structures of the various situations that one enters in the course of human relations and other events, and (b) to the degree that this disposition is robust and global. One is practically wise in proportion to how well one appreciates the rational force of the pros and cons by giving them motivationally the respective weights that they deserve.

5. Does this correction of crude virtue ethics give to sophisticated virtue ethics and psychology what they need in order to repel the situationist attack? Not fully. It does so, surely, only to the extent that subjects in situationist experiments reflect in their motivational structures the rational structures of the situations faced. But it is quite implausible that they do so. Either they go wrong in giving too much weight to factors that should not have so much, such as punctuality, or they give too little weight to factors that should have more, such as whether they are inflicting severe pain; or at least they give the wrong relative weight to such infliction by comparison with following the experimenter's instructions.

The switch from crude to sophisticated virtue theory is nevertheless important, if only for the sake of understanding properly the subject matter of our controversy. However, it is implausible to suppose that this proper understanding immediately provides a satisfactory response to the situationist attack. It is incumbent on virtue theory to grant that the experiments do raise legitimate doubt as to how global and robust human practical wisdom is and how global and robust are its more specific component virtues such as kindness, human decency, honesty, courage, and the rest.

That being granted, it would hardly follow that humans have no practical wisdom, none of the structure of virtues that, when properly integrated, constitute such wisdom. This sort of invalid inference is the crucial weakness in situationism to be probed here. Indeed, probing this sort of fallacy, once spotted, deflates the situationist attack even when aimed against the crude version of virtue theory. If the attack fails even against the crude version, it will be an even worse failure against the more sophisticated variant.

## 5.2 Defense of Virtue Ethics

1. The following defense of virtue ethics is based on an analogy between moral competence and driving competence. Let's define driving competence as a disposition

required of one by the situation faced. And this is a *normatively constituted* competence: it requires systematic foregrounding of the normatively relevant as such. Here the phenomenon of inattentional blindness is highly relevant. Take the seminarians in a hurry. It is not implausible that their disregard of the fellow human in need bespeaks not so much callousness as inattention, and indeed inattentional blindness. However, what seems still an open question is the extent to which such blindness is to be classed with culpable neglect as a moral failure. And this is, again, of a piece with the question whether failure to foreground the morally relevant is itself thus morally relevant.

to produce driving that is *safe*, when one is at the wheel, and *efficient* in routing to one's destination upon getting directions. Recall the contrasts applied earlier to traits, between the robust and the frail (or between the broad and the narrow, or the global and the local) and between the distinctive and the vulgar. These can be seen to apply with similar plausibility to driving competence. Someone's driving competence may be limited to quiet neighborhoods, for example, and may not extend to busy highways, nor to city driving. Such competence is then not as robust as it might be. The minimal driving competence required for safe and efficient driving in a sleepy village is very widely shared, moreover, and not as rare or distinctive as the physical abilities demanded by Formula One car races, or the navigational adroitness required by a reticulated old city.

Evaluatively relevant behavioral differences in instances of driving (one bad, one good) will not be explained by the shared vulgar competence of the two drivers. Any such behavioral difference explained by appeal to competence levels must of course appeal to some *difference* in such levels. Absenting any such difference in competence, one must appeal to some difference in situation.

2. Consider now the factors that have been found to affect the safety or efficiency of driving:

- (i) Brightness of light, even when the road is visible;
- (ii) Whether you are on a bridge when it is cold and wet (roadways on bridges being colder and potentially more slippery);
- (iii) Whether you're using a cell phone;
- (iv) Your blood level of alcohol;
- (v) Whether you received directions orally or through a map.

And so on. For some or all such factors, it must at some point have been surprising how good driving does depend on them. How should one's folk theory of driving respond to such discoveries? One possible response would be *driving situationism*, as follows:

- (a) Situations are dominant in the explanation of evaluatively relevant differences in driving behavior;
- (b) The robustness of our driving competence is rendered problematic: it is found to vary surprisingly with respect to previously unsuspected factors;
- (c) Personal integration in driving competence is not as widespread as one might have thought: thus, operational competence at the wheel does not necessarily go together with navigational competence.

These three theses are respectively analogous to the three listed by Doris in presenting the essentials of situationism (and quoted above). However, there has been no tendency to adopt driving situationism, nor is it plausible to conclude that driving competence is just an illusion or that we make a fundamental attribution error in taking people to be competent drivers when we see them display some good long

stretch of such driving, and when we explain the safety of their driving through their relevant competence.

3. Situations *will* indeed bear on our explanations of drivers' performances. But then any competence, indeed any disposition, will issue in a certain behavior only given certain triggering conditions and given a certain shape and situation of the host of that disposition. Behavior will in general have a two-ply explanation, one strand being the disposition or competence, and the other strand being the relevant triggering conditions that elicit the manifestation of the competence from the host in a given shape and situation. This is obviously true of dispositions in general. A sugar cube dissolves not just due to its solubility but also due to its insertion, while in normal shape in a normal situation.

Of course, any factor that to our surprise affects the quality of one's driving shows our driving competence not to be all we had assumed it to be. But this argues not so much for the abandonment of competence psychology as for its correction: with every such discovery we need to change our view of the shape or situation that underlies a competence.

Sometimes we have a choice between reasonable alternatives. Take the different ways one could respond to four factors that bear on safe driving: the brightness of the light, the amount of traffic, one's blood alcohol level, and one's visual acuity. Plausibly, the alcohol level belongs with the acuity as an inner basis for safe driving. After all, one's competence can change as can even one's acuity. By contrast, the amount of traffic is part of the situation, along with the quality of the light. How competent a driver one is at a time pertains to both sorts of factors. The degree of robustness of one's competence, for example, will be directly proportional to the breadth of the span of situations wherein one would produce good driving.

4. It might be replied that the analogy between driving and moral competence is very limited and potentially quite misleading. "You do not improve your moral competence by avoiding situations where it will be severely tested in strict analogy to how you improve your driving competence by avoiding bridges when it's wet and wintry." Even if this point has some force, the analogy remains effective. For one thing, you need not avoid the bridges so long as you heighten your awareness of the risk and adjust your behavior accordingly. Through such heightening and adjusting you *thereby* become more competent as a driver. Similarly, one way to improve your moral competence, surely, is to heighten your sensitivity to moral danger and to proceed with corresponding care.<sup>19</sup>

<sup>19</sup> Three studies in the April 2007 issue of the *Journal of Experimental Psychology* report on experiments conducted in Germany that examined 198 men and 92 women aged 16 to 45 while they played various games on a Sony PlayStation. Based on the results of these experiments, the psychologists think you may need to be extra cautious next time you get behind a real-life wheel. They say that people who play car-racing games drive more aggressively and have a greater risk of car accidents than people who play racing games less often, or who play "neutral" games.

5. Here are some lessons from our exercise.

The discovery of factors bearing surprisingly on our moral competence might more reasonably lead us to improve it than to reject its existence. So much for the normative lesson important for virtue *ethics*. (When we discover the bearing of sleep deprivation on good driving we tend to avoid driving when sleep deprived, thus improving our competence to drive safely.)

As for the bearing of the experimental results on virtue *psychology*, what they call for is, again, correction, not rejection. The sort of practical wisdom that explains a normal human being's ordinary behavior varies in degree from agent to agent, and also somewhat in structure, given (a) how variably humans can fall short and (b) how implausible it is to postulate a single acceptable motivational structure with sharp outlines.

We have learned of factors with a previously unsuspected bearing on our morally relevant conduct, factors that dull our discernment of the moral or other practical demands in certain situations, leaving nearly all of us with less practical wisdom than we had commonsensically claimed. Similarly, driving competence is affected by certain factors in ways that once proved surprising. A driver on his cell phone while crossing a bridge in wintry twilight will still likely reach his destination without incident, but we have long known that accidents under such circumstances are significantly more probable, so that smooth driving (prior to the discovery of the factors that are surprisingly relevant) is then *less* explicable through driving competence than we earlier thought, and more a matter of situational luck.

There are two sorts of relevant discoveries. Some factors might be thought to reduce a driver's competence temporarily, such as sleep deprivation and blood alcohol. Others might be viewed as conditions not covered by a driver's competence. That is to say, the competence might be viewed as more local and less robust than we earlier thought. Thus, driving competence is now known to be affected by ambient light, by cell-phone use, and by whether you are on a wet bridge in winter, to cite just three relevant factors.

On this view, the quality of driving competence is determined by a certain inner state of the agent's and by a set of distinguished situations. Thus, the inner shape of the driver can change so as to lower the degree of situation-relative robustness of the competence.

On an alternative conception, a driver's fundamental competence does not necessarily change with the noted inner changes. On this alternative, our driving competence is viewed as stable throughout, although the difference in expectable outcomes shows it to be less robust than previously thought, since the good outcomes are restricted relative to alcohol level and other inner states.

Whichever of the two views we take, we have discovered through the years that we possess neither the robustness of practical wisdom nor the robustness of driving competence that we had once optimistically self-attributed. We overreact, however, if we leap from that fact to the belief that driving competence and practical wisdom are just illusions.

### 5.3 The Virtue/Situation Clash: Genuine or Illusory?

Finally, the supposed clash between situationism and virtue theory turns out to be just an illusion owed to confusion and misconception.

1. Virtue theory makes no claim that normatively relevant behavioral variation is to be explained exclusively or even mainly through difference in virtues. Virtue theory is quite compatible with the view that humans are pretty much alike in the degree of virtue that they normally attain. *Compatibly with this, human conduct might still be universally explicable through the attribution of virtue to the agent.* Variability of marble rolls is not explicable through the universally shared disposition of marbles to roll, but each marble roll might still be explained in essential part through that disposition and its underlying basis, the rigid sphericity of a marble.

2. Virtue theory should accept that the experiments have shown humans to be less practically wise than folk virtue theory had imagined. Our practical wisdom now seems less robust or global than we had believed. While concluding that there are “no virtues of the sort that virtue theory had imagined,” situationists do not leap all the way to a Skinnerian nihilism of behavioral dispositions. But virtue theorists for their part should *accept* that there are no virtues *of the sort virtue theory had imagined*, since, after all, we are *less* robustly, globally virtuous than we had believed commonsensically. Situationists conclude: Forget virtues, explain by situations! But wait. Remember, behavioral explanation is two-ply, requiring when laid out fully *both* the relevant particulars of the situation, *and* the relevant non-Skinnerian dispositions. These latter look for all the world like traits, competences, virtues, however robust they may or may not turn out to be. Neither extreme position seems acceptable: neither that such traits explain with no situational help at all, nor that such situational particulars explain with no dispositional help at all. Behavioral explanation is two-ply when laid out fully.<sup>20</sup>

3. In light of the above, it would seem that any remaining substantive disagreement will be over degrees: Just how robust are the relevant human virtues?<sup>21</sup>

4. As for the situationist recommendation that we should assess the situations we enter for the relevant risks, this is fully endorsed by virtue theory, which will take the

<sup>20</sup> What is more, that *concedes* to situationism a restriction of the relevant virtues to the innermost seats of our relevant behavioral dispositions. However, common sense and virtue theory can and do recognize abilities that *encompass* the shape and situation of the agent. Thus, driving ability can be conceived to vary with changes in the shape or situation of the driver. *Competence* to drive safely on a certain road can vary with the driver's alcohol level and with the ice on the road. Corresponding varieties of competence will be distinguished in our theory of competence to be developed in what follows.

<sup>21</sup> And there is indeed a surprising array of factors that can influence our behavior subliminally and unexpectedly. For example, noise levels seem to affect helping behavior [Mathews, K. and L. Cannon. (1975). Environmental Noise Level as a Determinant of Helping Behavior. *Journal of Personality and Social Psychology* 32: 571–7]; as can fragrances in shopping malls [Baron, R. (1997). The Sweet Smell of... Helping: Effects of Pleasant Ambient Fragrance on Prosocial Behavior in Shopping Malls. *Personality and Social Psychology Bulletin* 23: 498–503]. But the like is, again, true of our driving performance.

discernment and proper weighing of such risks to be among the most important virtues in any given domain of human performance.

5. And we have been given no more reason to doubt that instilling virtue is worth the effort (on the part of parents and teachers) than to doubt that instilling driving competence and requiring its certification is worth the effort (on the part of relevant government authorities).<sup>22, 23</sup>

## 5.4 The Virtue/Situation Clash in Epistemology

So much for the defense of virtue theory on the side of ethics. Social psychology has provided an extensive basis for a similar critique of virtue theory on the side of epistemology. The critique of virtue ethics laid out in section 1 of this chapter is precisely replicable in the analogous critique of virtue epistemology. And against *this* critique of virtue epistemology there is of course a defense analogous to our foregoing defense of virtue ethics.

Because the analogy is so obvious, our remaining space is better devoted to a fuller development of a virtue epistemology defensible in the way suggested by our defense of virtue ethics.

1. We begin with a theory of competence for a virtue epistemology, one that will be safe against situationist attack. A competence is a disposition (ability) to succeed when one tries. How are such dispositions in general constituted? When complete, they have a “triple-S” constitution. Accordingly, we can distinguish three sorts of dispositions: the innermost (seat), the inner (seat + shape), and the complete (seat + shape + situation). With regard to driving competence (or ability), for example, we can distinguish between (a) the innermost driving competence: that is, the structural seat in one’s brain, nervous system, and body which the driver retains even while asleep or drunk, (b) the fuller inner competence, which requires also that one be in proper shape, i.e., awake, sober, alert, etc., and (c) the complete competence or ability to drive well and safely, which

<sup>22</sup> The third most important experiment bearing on our issues is the Stanford Prison Experiment due to Philip Zimbardo, who concludes as follows in a retrospective: “The critical message then is to be sensitive about our vulnerability to subtle but powerful situational forces and, by such awareness, be more able to overcome those forces. Group pressures, authority symbols, dehumanization of others, imposed anonymity, dominant ideologies that enable spurious ends to justify immoral means, lack of surveillance, and other situational forces can work to transform even some of the best of us into Mr. Hyde monsters, without the benefit of Dr. Jekyll’s chemical elixir. We must be more aware of how situational variables can influence our behavior.” (Zimbardo, P. G. (2007). Revisiting the Stanford Prison Experiment: A Lesson in the Power of Situation. *The Chronicle of Higher Education* 53 (30): B6. His related book is *The Lucifer Effect: Understanding How Good People Turn Evil*, published by Random House in 2007.) This critical message is one that virtue theorists can applaud. Yes, just as alcohol can deprive us of our driving competence, so group pressure can apparently deprive us of our moral competence. This would not show such competence to be an illusion, however, nor would it tend to show that belief in it is based on a fundamental attribution error.

<sup>23</sup> My thanks to Steven Lukes for illuminating formal comments at the Herdecke conference, and to Jason Kawall for helpful philosophical and bibliographical comments (reflected in footnotes 17 and 19).

requires also that one be situated with control of a vehicle, along with appropriate road conditions pertaining to the surface, the lighting, etc. The complete competence is thus an SSS (or an SeShSi) competence.<sup>24</sup>

2. What is required for possession of a competence?
  - a. A competence is a disposition to succeed when you try. So, exercise of a competence involves aiming at a certain outcome. It is a competence because it is a disposition to succeed reliably enough when one makes such attempts. A competence is hence necessarily a competence to  $\emptyset$  successfully, for some  $\emptyset$ , with one tied to a conditional of the form: if one tried to  $\emptyset$ , one would (likely enough) succeed.
  - b. Competences come in degrees of reliability, along with a threshold. However, in order to possess a competence to  $\emptyset$  it is not enough that the following conditional be true: that one would  $\emptyset$  reliably enough if one tried to  $\emptyset$ , that one would not too easily try to  $\emptyset$  without actually  $\emptyset$ -ing. This could after all be true simply because, knowing one's limits, one  $\emptyset$ 's rarely and only when in a narrow range where one would indeed succeed. Thus, one might restrict one's shots as an archer to situations where one is within two feet of the target. This might show good shot selection, assuming one is too unreliable at any longer distance, but it would not show archery competence.
  - c. What then is required for possession of a competence? Required for archery competence, for example, is a "sufficient spread" of possible shots (covering enough of the relevant shapes and situations one might be in as an agent) where one would succeed if one tried and an extensive-enough range. What constitutes this range? There must be a close-enough sphere of possible worlds where one takes shots, varied enough across the relevant range, and these shots must easily enough succeed, extensively enough across the relevant range.

It might be objected that even if one is tied down, so that, knowing one's condition, one *would not* try to shoot, this does not take away one's competence. This may be thought to refute the requirement that there must be a close-enough sphere of possible worlds wherein one does take shots. However, if one is tied down, that *does* surely bear on one's complete SSS competence, even when it does not affect the inner, SS competence, nor the innermost, S competence. One is then too poorly situated to retain the

<sup>24</sup> Once we understand competences and dispositions in this threefold way, we can make the familiar distinction of finks from masks for each component. A "fink" would be prompted to remove the competence by the trigger of the conditional corresponding to that competence. A "mask" leaves the disposition in place but prevents it from manifesting. Thus a mask might be an agent bent on stopping the manifestations, who would be provoked by the disposition's trigger to intervene in the process that normally leads to the manifestation, but without removing any of the relevant S's (the seat, shape, or situation). By contrast, a fink would block the manifestations by removing one or another of the S's. This might be the structural seat (innermost competence), shape, or situation.

A situation-fink masks both the inner competence and the innermost competence, finally, whereas a shape-fink masks the innermost competence.

complete competence needed to hit the target with one's arrows. Moreover, if indeed one is so disabled psychologically by a phobia that one cannot so much as *try* to shoot, then this does take away even one's innermost skill *S*. One is now no longer so structured psychologically as to be competent to  $\emptyset$ .<sup>25</sup>

## 5.5 Competence: A Special Case of Disposition to Succeed

1. Every competence is, again, a disposition to succeed when one aims in certain (favorable-enough) conditions while in (good-enough) shape. But not every disposition to succeed is a competence.

Thus, take a soccer goalie who faces an opponent shooter in a tie-break attempt. Suppose the goalie anticipates with a jump to the left. In such circumstances, if the shooter attempts to score by kicking to the left, the goalie is almost sure to make the save. But if the goalie chooses arbitrarily between jumping to the left, jumping to the right, and staying put, then her save is not nearly as creditable to her competence as it would be if prior experience with this shooter enabled her to anticipate the direction of the kick. If the goalie chooses her direction arbitrarily, then her anticipation of the shooter's choice is insufficiently reliable to qualify as a real competence.

2. Consider also a rank beginner who receives service for the first time in a game of tennis. When the streaking serve comes at him, he swings blindly and wildly. If the coinciding trajectories of the ball and the racket produce a wonderful return, that does not reveal an unsuspected competence seated in that receiver. Yet if he were ever to repeat that swing in those conditions, a successful return would result. Take the vigorous-enough swing with a certain arc, when the ball approaches with a certain speed and direction, so that it and the racket will happen to coincide properly. Such a swing will reliably produce such a wonderful return. But this does not reveal a competence, despite the fact that there is in that beginner a disposition to succeed reliably in such a situation if one swings thus.

3. Flukey success will generally admit a description similar to that of our lucky return of service. The performer will occupy a situation wherein he is disposed to succeed if he tries in a certain way in that situation, and he does try that way in that situation in the case at hand. There is a way of swinging available to our tennis receiver such that, relative to the specified situation of a ball traveling towards him at a certain angle and speed, if he were to swing that way, success would ensue.

<sup>25</sup> But here we have a choice. My own preferred choice is to say that one can have a "quasi-competence" to hit the target provided one is so propertied and situated that *if* one were to try to hit the target, one would be likely enough to succeed. But a *full, proper* competence would require also that one not be so disabled that one could never even try, no matter how desirable it might be to try or to succeed in hitting (a given suitable target).

4. In defining tennis-return competence we cannot relativize the way the player happens to proceed in the situation of the fluke. Rather than his competence holding only *relative* to his proceeding thus in that sort of situation, his swing manifests competence precisely because of his disposition to swing selectively that way, along perhaps with his competently helping to bring about the right sort of situation, for example by positioning himself well in preparation for receiving the serve.

Nevertheless, we cannot require of a performer that he bring about every aspect of a situation that makes it an appropriate situation to constitute a competence to succeed. Thus, the tennis player's competence is relative to certain lighting, wind, and precipitation conditions that he is not required to bring about.

Moreover, if a player tries to improve by practicing in challenging situations with a high risk of failure that cannot be held against him, if those situations are not included among the situations within which his competence is supposed to guide his performance to success. Thus, a tennis player might practice with a machine that feeds him balls much faster than any he would ever encounter in an actual game, or he might practice in lighting conditions that are significantly subpar. His risk of failure would of course rise dramatically in such circumstances, where he is still willing to test himself. But such a failure rate would not affect his degree of competence.

Here's another way to see the point: Suppose player A prefers to play with far better players, whereas player B likes to crush weaker opponents. Player B thus has a *much* higher success rate for the various performances in a tennis player's repertoire, both actually and dispositionally (given his penchant for weak opponents). But player A may still be the more competent player nonetheless.

5. Not every disposition to succeed when one tries constitutes a competence, then, although every competence will be constituted by a disposition to succeed when the agent is within certain ranges of shape and situation. A disposition to succeed is thus plausibly made into a competence by some prior selection of shape and situation, such that one seats a competence only if one is disposed to succeed upon trying when in that shape, in that situation. What shapes and situations are appropriate will of course vary from domain to domain of performance.<sup>26</sup>

6. A competence is a disposition to succeed, but it must be such a disposition properly restricted with respect to the three S's—Seat, Shape, and Situation. At least the shape and situation must be restricted to the appropriate. And these restrictions are imposed somehow within the domain of the relevant performances.

<sup>26</sup> For simplicity, I leave aside restrictions on *how you acquire* the relevant elements of competence such as the seat and the shape—restrictions that have come to the fore with the cyclist Lance Armstrong (re drug-induced shape) and with the baseball player Alex Rodriguez (re drug-derived seat). Each of these athletes enhanced his performances by enhancing his complete SSS disposition to succeed, but these dispositions did not remain competences once drug-enhanced. They did not remain ways to attain fully apt performance, performance creditable to (proper) athletic competence and thus creditable to the athlete's relevant competence.

Again, all competences are dispositions to succeed, but not all dispositions to succeed are competences. Only those dispositions are competences whose three S's fall within the proper range, a range established in the domain of the relevant performances.<sup>27</sup>

7. The terminology of competence is flexible. Consider the competence to drive safely on a certain road. Take the complete competence to drive safely at a time *t* on a certain stretch of that road. This requires that the subject satisfy all three of the S requirements. She must have the right seat/basis of the ability to drive safely (the requisite driving skill), she must be in the right shape (thus, awake and sober), and she must be properly situated with respect to that stretch of road (so that, for example, the road is not covered with oil).

However, we also allow a competence to drive safely that the driver can keep even when asleep or drunk. This inner seat would combine with appropriate shape and situation to the effect of a complete competence. Clearly one can be a good, safe driver, in possession of a competence to drive well and safely even if the nearby roads are all covered with oil. One's status as a good driver is not beholden to the condition of nearby roads. The fact that one is not competent to produce good driving on bad, even impassable, roads does not take away one's competence as a good driver. Nor would one's skill be impaired by inebriation or sleep. None of this affects the fact that one has the competence of a good driver. What is required for this is only that one seat that competence, that one host the seat/basis that, together with *appropriate* shape and situation, would dispose one to drive well if one tried. (But we could possibly distinguish between being a good driver and having the competence of a good driver. The latter would require only that one have the skills required for producing good driving if one tried, while the former would require also that one at least normally try when one drives.) Such an innermost competence, which abstracts from appropriate shape and situation, often constitutes a "skill."

8. We must not confuse ability with competence. The ability to  $\emptyset$  requires only that one have in one's repertoire of basic actions at least one action B whereby one would successfully  $\emptyset$ . ("Ability" stretches or contracts depending on whether we highlight the innermost, the inner, or the complete, but we put aside this complication for present purposes.) Competence is stronger in that it requires not only that one have the ability to  $\emptyset$  through a basic action in one's repertoire, but also that there be a set of such basic actions whereby one would succeed in  $\emptyset$ -ing and such that when one tried to  $\emptyset$  one would do so by means of at least one of them. On this account having a competence to  $\emptyset$  entails having the ability to  $\emptyset$ , but the converse is false. If one has a competence, one must have a corresponding ability, but if one has an ability one may lack a corresponding competence altogether.

<sup>27</sup> Many domains are set largely or wholly by convention, as are games, sports, and artistic domains. Other domains of human endeavor are set by our nature and needs, and by evolutionary teleology. Much is set by the approval of the group or of the species. But there is surely room for the group to fall short: moral leaders, for example, can lead the way to recognition of competences previously overlooked.

9. Finally, one's SSS profile on a certain occasion for a certain attempt to  $\emptyset$  may be such that any attempt constituting an exercise of that specific competence on that occasion would be likely enough to succeed so that the performance would be adroit or skillful, deriving as it would from a reliable-enough specific competence. And this is so even if the agent is not competent, period, in that domain of performance. After all, someone who has a very low level of competence in archery may still be competent enough to hit near and large-enough targets, so that these specific hits would derive from skillful-enough shots and might then be apt, since they would then succeed in virtue of that degree of competence. This is despite the fact that in more difficult situations, with the target somewhat farther away, the level of that archer's reliability would fall below the level required for the archer to remain competent enough.

Such an archer would thus remain competent enough to attain skillful and even apt performance when their relevant shape/situation profile was favorable enough. And this might happen even in the case of an agent who was not competent, period, in that domain, given how unusually favorable the shape/situation profile would need to be in order for them to enjoy a level of reliable-enough competence.

10. Thus, what are preselected in a domain of performance are the shape/situation dimensions within which an agent's reliability is to be assessed in determining their level of competence for performances in that domain. Thus, when an agent has a relevant specific shape/situation profile relative to a certain possible performance, their level of competence is determined relative to that specific profile. What determines their level of competence is just how reliably they would succeed when they tried while so shaped and situated. Also set by the domain is the spread of shape/situation pairs within which one must remain reliable enough in order to count as competent, period, in that domain.

However, note a perhaps surprising fact: namely, that even a quite incompetent agent might perform aptly and with a very high level of competence in a particular instance within that domain. All that would be required for this is that the agent be at that moment favorably enough shaped and situated so that he may thereby enjoy a high-enough level of reliable success with his attempts.

The fact remains that the relevant level of reliability is set by the preselected shape and situation dimensions, which will vary in how favorable they are for success by the agent. So, for an archer the relevant dimensions will involve alertness, blood alcohol level, visual acuity, degrees of light and wind, distance from the target, and the like. Notably absent is the presence of a wizard and degree of helpfulness of that wizard. Two things are then determined relative to those preselected shape and situation dimensions. And let's assume for simplicity that there is just one shape dimension and one situation dimension. With this simplifying assumption, we can say that once these two dimensions are preselected then the agent must succeed reliably enough with his attempts in broad-enough ranges of shapes and situations, in

order to count as competent, period, in a given domain. Secondly, an agent might be reliable enough to count as competent within a particular shape/situation circumstance, even if their prowess is highly restricted within the relevant shape and situation dimensions. They might be thus sufficiently restricted to count as incompetent in the situation involved, while still performing with high-enough competence and reliability within the very specific shape/situation profile that they in fact enjoy in that situation.

## 5.6 Competences, Epistemic and Otherwise

1. We can distinguish judging that *p* from merely thinking that *p*. The latter requires that one affirm in the endeavor to get it right on the question of whether *p*. But that is compatible with just psychologically guessing. Full judgment requires more: One must endeavor not only to get it right on that question, but also to do so aptly enough.

2. And here we find a reason why *judgment is automatically on a reflective higher order*. For, what does it mean that one endeavors to get it right “aptly enough”? Succeeding in this endeavor requires that one would affirm in such a way, in such circumstances and while in possession of sufficient relevant competence that, given one’s complete relevant competence, one would be likely enough to affirm correctly. That is to say, one’s success rate for such attempts must be high enough. And what would go into such an endeavor? What must one consider as one aims to affirm aptly and therefore reliably enough? Must one not consider one’s relevant SSS situation and how likely it is that a relevant attempt manifesting such a complete competence would succeed?

In that case one takes a higher-order attitude towards one’s three options: affirming, denying, suspending. One considers which of these would be best from the point of view of taking a risk if, but only if, it is proper to do so. What one considers then is the reliability of affirmation or denial of the question at hand. One thus considers whether double-omission would not be preferable given the risks involved in either affirming or denying. And this is all on the second order, since one must consider one’s relevant, first-order, complete competences, and the first-order options of affirming, denying, and double-omitting.

3. For performances more generally something similar is true across the board. Consider, for example, three-point shots in basketball. Even a player who overconfidently takes low-percentage shots too frequently may retain an excellent ability to sink three-point shots close enough to the three-point line. So, success in such an attempt does seem properly creditable to his competence, even if he tries too often when he should not do so, when he is so far from the basket that his reliability plummets. This is *especially* plausible when he is *fully aware of his limits* but tries to succeed even beyond them, as a deliberate risk taker. But it remains plausible enough even for

someone who lacks competent and full-enough knowledge of his limits and shoots beyond them only because he incorrectly takes himself to be reliable enough even when he no longer qualifies.

The difference between those two cases is as follows. In the former, the agent can still perform with reflective and full aptness when he knowingly performs within his limits. In the latter, the agent no longer performs with reflective and full aptness, at least not when too near his threshold of reliable-enough performance. Nevertheless, he can of course still perform with animal aptness even so—indeed superlatively so.

4. But what is the aim of a basketball shooter? Does she aim merely to get the ball in the hoop? If that is her sole relevant aim, then she does succeed in this, creditably so, even when she is far out, especially if her success even that far out is due to a level of competence well above the average. Important considerations about her performance as a team player are relevant to this assessment. For example, should she pass rather than shoot? However, normally performers do not seek just to succeed in the sense of attaining their basic aim no matter how unreliably. Normally they aim to succeed aptly enough, while avoiding too much relevant risk.

5. Consider now the three-point shooter as she approaches the distance to the basket where her reliability is indiscernibly near to the relevant threshold. And suppose that in fact her reliability *is* above the relevant threshold, but *indiscernibly so to her*. A statistician/coach/observer might know perfectly well that the player's reliability is still barely above the threshold. Suppose he has studied her success rate extensively, aided by a device that measures with exactitude her distance from the basket; that way he can tell that she is still reliable enough at that distance. But the player is very far from knowing any such thing. Consider now her aims as she shoots from that distance. She may attain her basic aim: namely, to sink that basket. And in that respect her shot may be apt. It may succeed in a way that manifests the competence that the statistician knows her to possess even at that distance. So, she has animal competence and her shot is not just successful, not just competent, but also apt. The success of the shot, and not just its existence, does manifest the sufficient competence resident in the player even at that distance.

6. What then is she missing? Anything? Well, although she attains her first-order objective, and does so aptly, she does not aptly attain her reflective aim of succeeding aptly enough. She does succeed reliably enough, but her attaining *this* objective is a matter of luck. Unlike the statistician, she is unable to tell that her shot is still reliable enough even at that distance. If she shoots anyhow from that distance and it turns out that her shot is reliable enough, then she may attain her objective of succeeding reliably enough. However, her attainment of this objective will not be apt. It will be relevantly lucky enough that the full success of her shot (its going through the hoop reliably, aptly enough) is not *fully* creditable to her, given the important element of luck involved. So, her first-order success will be apt but it will not be metacompetent and hence not meta-apt, and so it will not be fully apt.

## 5.7 More on the Nature and Epistemology of Competences

1. Competences are very special cases of dispositions. Suppose a solid-iron dumbbell would shatter upon hitting a certain surface only because a hovering fiend is determined to zap it if and only if it hits the surface, and just as it does so. This would make the relevant conditional true: the dumbbell *would* shatter on that occasion, but this does not make it fragile. For the latter to be the case, one would need at a minimum that the dumbbell would shatter whenever relevantly impacted in a relevantly “interesting” combination of condition and situation.

2. Just what makes a triple-S combination of skill, shape, and situation “interesting” is an interesting, and neglected, question. Such combinations constitute the innumerable competences of interest in the many performance domains recognized by human communities. Such domains—whether athletic, artistic, intellectual, medical, scientific, legal—contain performances aimed at certain distinctive aims, along with the competences whereby performers might succeed more or less competently. When a success manifests such competence then, and only then, it is apt. The SSS profile that underlies any such competence is not always subject to full and detailed linguistic formulation. But in that respect it joins much else of substantial human interest and importance, such as when conduct counts as polite.

What conduct does count as polite? This is much easier to discern by perception in a particular situation, than it is to formulate in full detail. This is so for politeness in general, as well as for particular matters of politeness, such as the distance that is appropriate for normal conversation in person, and the tone and volume of voice that counts as rude, to take just two out of many instances.

Moreover, that is not to say that rules of polite conduct are just a myth. Not all rules need to be formulated linguistically. If what is polite is defined by human convention, however, then there must be rules, in some broad sense, that communities agree upon ahead of time; there need not be a constitutional convention to institute those rules. Their institution may be more organic, less artificial than that. Moreover, the agreement within the community will not require linguistic communication. It may be instituted more implicitly than that, through persistent profiles of approval/disapproval, praise/blame, and systematic and implicit manifestation in individual and social conduct.

3. Returning to our own issues, this pattern of normativity is applicable also to competences. Thus, the SSS profiles of the competences that acquire salience for communities may be determined not by antecedent linguistically formulable convention, but rather by persistent implicit profiles of the sort that determine the content of politeness for that community.

What is more, the relation is not just a matter of coincidence or mere analogy. There is after all such a thing as competence in etiquette, and even in morality. It seems

possible to conceive of these domains either in terms of rules implicitly operative, or in terms of competences determined by proper community interest.<sup>28</sup>

## 5.8 Competence, Safety, and Reliability

1. In order to possess a performance skill (or the seat/basis of a competence), one need not satisfy any general reliability requirement. We have seen how one can exercise one's skill too often in inappropriate shapes or situations, so that one would fail with extremely high frequency while of course remaining in place (to be thus frequently exercised).

2. Skill might be present and exercised, and its exercise might even lead to success without its being manifest in that success. Thus a skilled archer might shoot skillfully, and the arrow might be deflected by an unexpected gust. However, a hovering guardian angel might happen by, determined to correct the trajectory of any well-shot arrow so that it would hit the target when (but only when) it would have done so but for the unlucky intervention of a gust. In this case, the archer's skill is exercised, and in the situation as pictured the skill does partly underlie a disposition to succeed. But this disposition does not amount to a true competence, since it is situationally based essentially on the angel. And the angel's presence is a lucky exception and not a stable component of the relevant background conditions. What would make that innermost archery seat a true skill is its ability to combine with appropriate shape and situation so as to dispose the archer to succeed reliably enough if she tried to hit a target. However, in the actual situation the archer's seat/basis fails to combine with an appropriate situation. The angel is an exceptional and ad hoc situational factor. No such angel figures in the situations of interest in the domain of archery, the situations relative to which we assess true archery competence and achievement.

Consider the archer who shoots with the unexpected gust about to cross the field and with the guardian angel ready to intervene, unbeknownst to the archer. This archer does not earn proper credit for her success, which does not really manifest competence. And the reason for this, I suggest, is that her success does not manifest a sufficiently reliable competence based simply on her level of skill and the appropriate shape and situation. The appropriate shape and situation for archery competence cannot depend essentially on interventions by an ad hoc guardian angel. The appropriate shape and situation must be taken for granted as background for archery shots. Relevant credit will accrue to the archer's skill only if the success of the shot does manifest that skill. Such credit will normally accrue only if the agent's skill produces success deriving just from reliable-enough skill in appropriate shape and situation.

<sup>28</sup> Of course, that is not to say that all such domains have their normative source or basis in human convention. Nor need we suppose convention-based domains to lie beyond objective assessment in global or specific respects. They may still be subject to such assessment on how well they further any value that they may be designed to secure.

3. Skills do come in associated clusters. An archer who knows of the guardian angel, with her power and intentions, might have an enhanced archery skill that takes that into account. Such an archer might indeed manifest her enhanced archery skill when the gust comes along and the angel puts the arrow back on track. The success of her shot would thus be properly credited to that archer's enhanced skill, and thus properly credited to her. (Compare the orchestra conductor who makes skillful use of his players in pursuit of his proper ends.)

4. Similarly, an archer could deliberately frustrate the attempt of a hovering demon to take his arrow off course. Here again an enhanced competence would be in play, to which the shot's success could be credited properly.

5. The bottom line is that the success of a performance in attaining its aim is properly credited to a skill—to an innermost seat of competence—only if that skill counts as a skill relative to pre-selected shape/situation combinations of interest in the domain of that performance and the performance's success manifests the skill: that is, only if the skill produces the success reliably in combination with the agent's relevantly appropriate shape and situation.

## 5.9 Conclusions

Situationists<sup>29</sup> have argued that virtue epistemology, whether responsibilist or reliabilist, faces serious problems posed by experimental results accumulated over years. What exactly are these problems?

One problem is supposed to be that we are surprisingly likely to be less reliable than we had thought, and that we are made less reliable by influences astoundingly *irrelevant* to the truth of the beliefs that they nevertheless do influence. So, the bottom line here is that we are *less reliable* epistemically than we had implicitly supposed all along.

There are two serious gaps in that critique, the first (though not the second) of which is independent of any particular theory of competence. We next consider how damaging these are.

In the first place, in order to show that we are lacking in epistemic competence it would need to be shown that the ways in which we form beliefs are *unreliable*, or *insufficiently reliable*. All that has been made plausible, however, and even this only for *some* domains of belief formation, is that we are *less reliable than had been thought*, and *less reliable in ways and for reasons that are surprising, and sometimes even astounding*. But from none of this does it follow that there is *any* recognized domain where we form beliefs in some distinctive way(s), wherein we are *unreliable*, or *insufficiently reliable* to

<sup>29</sup> I have found helpful the following excellent, thought-provoking works: Alfano, M. (2013). *Character as Moral Fiction*. Cambridge: Cambridge University Press. Olin, L. & Doris, J. M. (2013). Vicious Minds: Virtue Epistemology, Cognition, and Skepticism. *Philosophical Studies*, published online 19 June 2013 (DOI 10.1007/s11098-013-0153-3). Fairweather, A. & Montemayor, C. (2014). Inferential Abilities and Common Epistemic Goods. In A. Fairweather (ed.) *Virtue Epistemology Naturalized*. Synthese Library.

*count as competent* (and as competent enough to regularly attain knowledge through the exercise of our competence in that domain).

Suppose, however, that we *have* been shown to be *insufficiently* reliable in *some* domain where we had taken ourselves to be reliable enough to attain knowledge competently. Even if that is so for some specific domain, even if we are indeed unreliable believers in a particular set of circumstances distinctive of that domain, it would be fallacious to infer from this that we are more generally unreliable in some way that puts our common-sense attributions of knowledge seriously and generally in doubt.

Compare the many perceptual illusions that psychologists have detailed, which do conclusively establish how seriously unreliable our perceptual systems are in certain situations. That has given no good basis for any general rejection of our perceptual competences.

So much for the first supposed problem. A second problem is laid out as follows. First, it is pointed out that in a certain situation *X* our supposed competence to  $\emptyset$  regularly fails us. And it is inferred that we do not really have any such competence. What we have is rather something like this: We have a “competence-to- $\emptyset$ -when-not-in-*X*.”

The fallacy in that form of reasoning is exposed if we apply it to almost any athletic competence. Take a basketball player who has a highly prized competence to sink three-point shots. Placed far enough back from the basket, at center court, his reliability will plummet. So, what should we conclude from that? Must we say that strictly speaking he lacks competence as a three-point shooter? Has he at most a competence to sink three-point baskets “provided he is not at center court”? That is not how we normally think of a competence. A competence does come with implicit reference to preselected shape and situation within which it might be exercised properly. So, yes, when we invoke a competence or a virtue there is an implicit relativization to shape and situation. But, as we saw in section 8 above, that is not so much a problem as a feature.

As we have seen, a competence is a disposition to succeed if one tries through a basic action in one’s repertoire. But not every such disposition is a competence, as we have seen. Recall the tennis beginner’s disposition to succeed if he swings a certain way when the ball approaches in a corresponding way, such that he has the *ability* to swing that way in such circumstances. So, in a particular situation, the player may be disposed to return the serve well if he swings in way “*W*,” given that the ball approaches in way “*W*,” as in fact it does. This is a disposition grounded in a certain combination of seated skill, shape, and situation, true enough. And it is an ability that our agent does have. What our beginner importantly lacks, nevertheless, is a corresponding competence. What is required for possession of a given competence is that one be disposed to succeed if one tries, with a disposition based on a certain preselected range of shape/situation combinations. And this requires not just that one have the ability to do so, through any of a set of basic actions in one’s repertoire. It requires also that one *would* reliably enough try *through one or another of those success-promoting basic actions in one’s repertoire*. One needs to have such a repertoire

of basic actions that would reliably enough lead to success within the shape/situation combinations preselected as of interest by the relevant community (or even by the relevant species, for species-wide competences and aims). The fact that one does not succeed when one tries in a certain specific sort of situation may thus have no bearing at all on whether one enjoys a certain competence to succeed by  $\phi$ -ing when one tries to do so. That competence may after all involve a preselected shape/situation range that excludes that particular sort of situation. Earlier sections above aim to elucidate such competences.

We may thus conclude that neither situationist objection withstands a closer look. An interesting conclusion *is* implied by this approach: namely, that knowledge and other human accomplishments or achievements are normative in ways previously unremarked. They are normative in that they are successes that manifest competences, where competences are dispositions involving preselected shape/situation combinations wherein specifically human accomplishment is prized (or otherwise of special interest).<sup>30</sup>

<sup>30</sup> This (along with the foregoing) sketches a theory of competence meant to help explain the nature of human accomplishment, including the epistemic accomplishment of the knower. A fuller account is found in my (2015) *Judgment and Agency* (Oxford: Oxford University Press).