



# Brittle virtue or bust: a new challenge to virtue-as-skill theories

Daniel Kilov<sup>1</sup>

Received: 8 February 2023 / Accepted: 10 August 2023 / Published online: 6 September 2023  
© The Author(s) 2023

## Abstract

Buoyed by research in philosophy and moral psychology, virtue ethics has become increasingly influential in the literature. This renewed attention has also led to the development of the *situationist challenge*: empirical studies undermine the idea that we possess character traits that allow us to act virtuously across contexts. A promising reply to the situationist challenge is that we should not conceive of virtues as traits. Instead, we should conceive of them as expert skills. Here, I raise a new challenge for those who go in for the (s)kill against situationism. Expert skills are *brittle*, rather than flexible or robust. They typically develop in narrow ways and can falter in novel situations. Experts also sometimes perform worse than novices in novel situations. Such findings present a problem for the virtue-as-skill thesis. Someone who will help others only in narrow contexts or who causes more harm than good does not exemplify virtue. The brittleness of expert skills means that virtue theorists must either (a) abandon the idea that virtues explain behaviour across contexts or (b) abandon the idea that virtues are expert skills. It's *brittle* virtue or bust – you can't teach an old dogma new skills.

**Keywords** Virtue ethics · Situationism · Character scepticism · Expertise · Skill

## 1 Introduction

In recent years, virtue theorists have gravitated towards the idea that virtues are properly understood as expert skills. This move is, in part, motivated by situationist critiques originating from Harman (1999, 2000) and Doris (2002). Research from social psychology casts doubt on the idea that we possess character traits that manifest robustly across contexts. By contrast, Stichter makes the point that:

---

✉ Daniel Kilov  
Daniel.Kilov@gmail.com

<sup>1</sup> Australian National University, Canberra, ACT, Australia

one of the advantages of the “virtue as skill” thesis is that given the vast research done on skills and expertise by psychologists, there is no psychological skepticism about the ability of people to acquire skills (in contrast to being able to acquire global character traits) (2018, p. 3).

However, there are also independent reasons for thinking of virtues as skills. One reason is that both are acquired through practice. Aristotle draws the comparison explicitly:

What we need to learn to do, we learn by doing; for example, we become builders by building, and lyre players by playing the lyre. So too we become just by doing just actions, temperate by doing temperate actions and courageous by courageous actions (NE 2009 1033 a 32-b2).

One cannot learn to swim from a book; one has to get wet. In the same way, one cannot cultivate virtue without repeated opportunities to practice.

If we take the virtues-as-skills thesis seriously (and take seriously Aristotle’s prescription to ‘use as evidence what is apparent for the sake of what is obscure’ (NE 1984 1104a, 14–15)) then we must attend to the rich psychological literature on the development and shape of expert skill. This can shed light on the nature of virtue, but it also raises new challenges for virtue theory.

Some of the most surprising results in the literature on expertise speak to the *brittleness* of expert skill (see Kilov, 2020). As Feltovich and colleagues point out, expertise typically develops in very narrow and highly specific ways, such that there is ‘little transfer from high-level proficiency in one domain to proficiency in other domains – even when the domains seem, intuitively, very similar’ (2006, p.65). Further, expertise is brittle in the sense that the very factors that mediate expertise will sometimes cause experts to perform worse than novices in novel situations.

These findings present a problem for virtue-as-skill theories. Someone who will help others only in narrow contexts or who causes more harm than good does not exemplify virtue. Virtues are supposed to manifest *robustly*, that is, across contexts and in novel situations. Alfano (2013) refers to this kind of cross-situational consistency as part of the ‘hard core’ of virtue ethics. The problem is that the empirical evidence speaks against this kind of consistency. As such, virtue theorists must either (a) abandon the idea that virtues explain behaviour across contexts or (b) abandon the idea that virtues are expert skills.

In § 2 of this paper, I spell out what it would take for skills to be cross-situationally robust in the relevant sense. Specifically, I claim that it would require that the skills be flexibly deployed in novel situations.

In § 3, I argue that skills generally do not meet this condition.

In § 4, I spell out a novel argument against the virtue-as-skill thesis before anticipating replies. I focus on Stichter’s work (notably, 2018). He puts forward what appears to be the most empirically informed and well-developed virtue-as-skills account. Stichter’s account is quite different from previous accounts, such as Annas (2003) and addresses or otherwise evades objections to older views (e.g., those sur-

veyed in Dougherty, 2020). My objection is new in that it targets the most up-to-date and otherwise defensible version of the virtue-as-skill thesis.

## 2 Cross-situational robustness and the virtues of virtue

### 2.1 Cross-situational robustness as flexibility

As mentioned, the assumption that virtues exhibit a certain kind of cross-situational robustness is central to traditional accounts of virtue. On these accounts, virtues are understood as global character traits. The scientist who is conscientious in their research, for instance, is supposed to exhibit similar conscientiousness with respect to their social commitments and hobbies. However, there is considerable evidence that character traits such as conscientiousness do not function in this way (for a thorough accounting of the past and future of the situationist challenge to character traits see Doris, 2022).

However, some virtue-as-skill advocates argue that identifying virtues with character traits was a mistake all along (Stichter, 2018). If so, then it may also have been a mistake to think that virtues must exhibit the cross-situational robustness required of traits. Could the insistence on robustness be a vestigial limb we lop off as we move to embrace the virtue-as-skill thesis? I do not think so for the following reason. Much of the appeal of virtue theory derives from the idea that virtues can be invoked to explain behaviour. From the standpoint of virtue theory, to ascribe virtues (or vices) to someone is not just to evaluate their actions. It is also to explain why those actions were taken. According to MacIntyre, ‘very little will be genuinely explicable’ if we do not appeal to the role of virtues in explaining behaviour (1984, p. 199). Furthermore, virtues must manifest consistently across time and context (wherever the relevant reasons present themselves) if they are to be explanatorily useful. As Dent puts it, virtue is supposed to cause appropriate behaviour in ‘ever-various and novel situations’ (1975, p. 328). Thus, if virtues are supposed to play the explanatory role assigned to them, then they must exhibit robustness across contexts<sup>1</sup>. This is the case whether virtues are understood as traits, skills, or something else. What good is an ascription of honesty if it does not explain why someone told the truth or give some assurance that they will do so in the future or in other circumstances?

What does cross-situational robustness amount to in the case of expert skill? Invoking the Aristotelian injunction to understand the new in terms of the familiar, the relevant form of robustness is best understood in terms of the flexibility of learnable skills. As Narvaez puts it, a central motivation for adopting the virtue-as-skills approach is that:

moral behavior is pruned from the rigidity of personality temperament and put into the realm of learnable behavior. It appears more like behavior in other domains

<sup>1</sup> This view is widespread among defenders of virtue theory. In their defence of virtue theory, Winter and Tauer (2006) identify the robustness of traits as being the key issue. Similarly, Miller (2014, p. 191) also notes that virtue ethicists “are standardly committed to global traits”.

like football or chess, as a set of skills that can be learned [and deployed across contexts] (2005, p. 141).

Following Stuart Dreyfus and Hubert Dreyfus (1986), philosophers working on skill have largely assumed that expert performance is characterized by a kind of flexible coping. They have thus sought to explain this flexibility or appealed to it as evidence for particular theories of the cognitive underpinnings of expert skill. For example, Christensen and colleagues (2019) appeal to the putative flexibility of experts to argue that skilled performance cannot be explained in terms of automated processes. They argue that expert skills must be efficient, but they must also be flexible. Since automatic processes are relatively inflexible, they conclude that skills must not automate fully (p. 694). For Christensen and colleagues, experts exhibit flexibility with respect to their performances because they can dynamically respond to rapidly shifting situational factors or respond adaptively to previously unencountered aspects of complex domains of performance. As they point out (2019) the state space of a system can be large even when it has a small number of elements. A chess board has sixty-four squares and thirty-two pieces. However, the number of possible chess combinations has been estimated to be between  $10^{40}$  and  $10^{50}$ . Further, chess experts stand to gain considerably from drawing their opponents into unfamiliar situations and so are strongly incentivised to drive towards novel situations/parts of the domain.

So, expertise involves cross-situational robustness insofar as experts exhibit flexibility with respect to their domains of expertise. However, this just passes the explanatory buck to the notion of flexibility. What exactly is flexibility and do experts have it?

## 2.2 Flexibility versus sphexibility

One way to get to grips with what it means for a behaviour to be flexible is to consider an instance of its opposite: the behaviour of the Sphex wasp.

The Sphex wasp is typically understood to have a largely mechanical kind of intelligence. Before entering the burrow, its routine is to leave the prey at the entrance, inspect the inside of the burrow, and then re-emerge to drag the prey inside. However, if the prey is moved (even a few inches) while the wasp is inside the burrow, then it will bring the prey back to the entrance and re-initiate the entire inspection procedure. If the prey is moved again, then the Sphex will repeat the process again, and so on *ad infinitum*.

The Sphex' behaviour can be modelled as a kind of finite state machine. Each step in its behaviour is triggered by the combination of its internal state and readily detectable features of the environment. To offer a further example, there is evidence that the Sphex wasp will drag a cricket by its antennae. If the antennae are removed, then the wasp will go in search of new prey. It will never think to drag the cricket in some other way (by its leg, for example) (Fabre, 1915). Clearly, the behaviour of the Sphex wasp is *inflexible*.

Suppose we set out to create a kind of Super Sphex by enumerating some large number of further behavioural sequences to account for various contingencies<sup>2</sup>. For example, “if the cricket has been moved but the burrow has been inspected, then proceed to drag the cricket into the burrow” or “if the antennae are removed from the cricket, then drag it by its leg”, and so on.

Does our Super Sphex engage in flexible behaviour? In one sense, the answer seems to be “yes”. The Super Sphex engages in adaptive behaviours across a large part of the domain of its activities (those for which behavioural sequences were enumerated). In another sense, the answer is clearly “no”. The Super Sphex cannot respond to novel contingencies and will be tripped up by changes in its domain. Some changes will merely interfere with its adaptive behaviour, while others may lead to it to behave in ways that are maladaptive. Let us call the kind of flexibility that the Super Sphex exhibits *sphexibility*. We can reserve the term ‘flexibility’ for cases where behaviour in response to novel contingencies is genuinely novel and adaptive. That is, a genuinely flexible agent is one who can transfer their skills to situations outside the scope of their prior experiences.

Having drawn the distinction between flexibility and sphexibility, we can ask the crucial question: Are experts flexible or merely sphexible? It is bad news for proponents of the virtue-as-skills account if experts are merely sphexible since it seems unlikely that expert skill would exhibit the kind of robustness required to animate the virtue-as-skills account. This is because virtue theorists typically maintain that a virtuous agent is an agent whose virtue allows them to act appropriately in new situations. As Olin and Doris put it, ‘[t]alk of virtues compels at least in part because virtue – be it epistemic or moral – promises progress on unusual problems, in difficult conditions’ (2014, p. 679).

Answering the crucial question requires that we engage with the rich empirical literature on skills and expertise. Building on the survey of the literature in Kilow (2020), I provide additional examples of the brittleness of expert skill. Further, I broaden the scope to examine additional landmarks along the path to expert skill, examining evidence for brittleness in advanced students. Such evidence forms an important part of the story, since, as Stichter notes, ‘the possession of virtue is typically understood as a matter of degree... That puts virtue back on a par with expertise, since expertise also admits of degree’ (2018, p. 99). I now turn to the empirical literature on skills and expertise.

---

<sup>2</sup> Indeed, this may be a more accurate description of *actual* Sphex wasp behaviour. As Keijzer (2013) argues that, though the anecdote about the Sphex wasp is now baked into the mythology of contemporary cognitive science (repeated, for instance, by Brooks 2002; Dennett, 1978; Hofstadter, 1985; Millikan, 2004; Sterelny, 1990), the actual evidence about Sphex behaviour is equivocal: endless repetition is not standard. Perhaps the Sphex wasp is not stuck in a behavioural loop, but it seems that we philosophers sometimes are, at least with respect to promulgating the above anecdote!

### 3 Evidence for brittleness

#### 3.1 On the brittleness of skill

It is natural to think that the flexible application of learned materials is at least partly constitutive of knowledge and understanding. If a student is unable to apply the lessons of the classroom to the world beyond it, in what sense can they really be said to have *learned* anything at all? However, research suggests that the skills of even advanced students are often disappointingly brittle. The psychologist Gardner (1991) cites research carried out at Johns Hopkins, MIT and other well-regarded universities. The research documents how students who receive honour grades in college-level physics courses are often unable to solve problems that differ slightly from those they encounter in class and in textbooks. Caplan provides additional evidence:

If you throw a coin straight up, how many forces act on it midair? The textbook answer is ‘one’: after it leaves your hand, the only force on the coin is gravity. The popular answer, however, is ‘two’: the force of the throw keeps sending it up, and the force of gravity keeps dragging it down. Popular with whom? Virtually everyone – physics students included. At the beginning of the semester, only 12% of college students in introductory mechanics get the coin problem right. At the end of the semester, 72% still get it wrong (Clement 1982). After students learn how to handle complex homework and exam problems, few apply their lessons to simple real-world cases (2018, p. 57).

This problem does not plague only physics students. In a study by Voss and colleagues (1986), college graduates were asked questions about economic issues. The authors found no difference in performance between those who had taken economics classes and those who had not. The authors concluded that ‘the results suggest that classroom instruction in economics does not necessarily lead to superior performance on “everyday” economics tasks’ (1986, p 269).

Should these results trouble champions of the virtue-as-skill thesis? If one squints in just the right way, these results can almost be made to look like good news for virtue theorists. After all, it’s common knowledge that true virtue is rare, as is true expertise. Evidence that graduate students stumble when they encounter novel circumstances seems to confirm this platitude. However, taking this line will not help virtue-as-skill theorists. Recall that one of the primary appeals of the virtue-as-skills thesis is the insight that virtue is something that admits of degrees. Raising the bar for what counts as ‘genuine’ possession of virtue means that virtue will be unrealisable for almost everyone<sup>3</sup>.

---

<sup>3</sup> An additional problem with the “virtue is rare” reply in the context of skills is that skills, even great skills, are not rare. There are plenty of great surgeons, carpenters etc. My thanks to a reviewer for pointing this out.

More problematic, still, is the implicit assumption that experts differ from novices with respect to the flexibility of their skills. As we shall now see, the evidence does not support this idea.<sup>4</sup>

### 3.2 On the brittleness of expert skill

The spflexibility theory of expert performance (the theory that expertise is *merely* spflexible) predicts that expert skill is narrowly confined to specific domains. Expert skill is very sensitive to situational factors and changes to the domain of performance often result in a precipitous drop in performance. Furthermore, the inflexibility of experts' responses means that experts perform worse than novices in some situations. That is to say, the spflexibility theory predicts that expertise is often *brittle*. One finds numerous examples of these kinds of brittleness in the literature.

Experts' superior performance is domain-specific, arising only in the narrow context of their expertise. For example, Nodine and Krupinski (1998) compared search performances of expert radiologists and non-experts on non-radiographic images, such as *Where's Waldo* images. In these tasks, radiologists performed no better than the non-experts; they did not experience a generic boost to search performance.

Further, research indicates that surgical expertise is task-specific and does not transfer to even closely related tasks (Wanzel et al., 2002). Proficiency in one surgical procedure doesn't translate into competence in others, nor does it enhance performance on tasks closely mimicking their speciality. Van Sickel et al. (2007) established that neither years of practice nor the number of completed laparoscopic procedures predict performance on a laparoscopic training simulator; rather, the key factor was specific experience with the simulator. Similarly, Park et al. (2007) found that success in training simulators didn't accurately forecast performance in clinical environments. These findings shouldn't lead us to doubt surgical expertise. Instead, they underline how sensitive expertise is to the contours of a domain.

In some cases, introducing changes to a domain will interfere with expertise. For example, professional baseball hitters were unable to hit throws by Olympic softball champion Jenny Finch (Horovitz, 2004). Although Major League Baseball players routinely hit baseballs travelling at over 153 kph while Finch pitches at 110kph, differences in the field and pitching style meant that they were unable to connect with a larger, slower target.

Changes need not be dramatic to elicit a drop in performance. For instance, Saari-luoma (1991) found that experts at blindfold chess were unable to track the positions of pieces if random, rather than conventional, moves were performed. This result is consistent with the now famous research by Chase and Simon (1973), which demonstrated that the superior recall of chess experts largely disappears when random board configurations, rather than configurations of actual boards, are used in the recall task.

In a much-discussed study, Morrot and colleagues (2001) asked fifty-four expert wine tasters to taste two glasses of wine, one red and one white. Unbeknownst to the

<sup>4</sup> The reasons for brittleness in these cases will probably differ. In the early stages of expertise, brittleness is likely to result from exiguous representations or schemata. In the case of a *bona fide* expert, brittleness is a consequence of rich but entrenched representations or schemata. See Dane (2010).

tasters, both glasses were from the same bottle. The glass of 'red' contained an odourless red dye. Lexical analysis of the experts' judgements revealed that 'because of the visual information, the tasters discounted the olfactory information' (2001, p. 309). Changing the colour of the wine is a fairly superficial change to the domain. After all, visual information is not obviously relevant to judging the *taste* of wine in the way that olfactory information is. Still, it was enough to prevent the experts from making even the grossest distinction between red and white wine.

In other cases, accumulated expertise can interfere with performance in such a way that experts perform worse than novices. For example, Shiffrin and Schneider (1977) report that people trained for several thousand trials to detect visual targets among distractions found struggled when the target and distracting objects were reversed. Performance was worse than at the very beginning of training. Indeed, people needed almost two and a half thousand additional trials of practice under the reverse conditions to reach the level of performance originally obtained after one thousand five hundred trials.

These results seem to hold for more complex domains of skill. Sternberg and Frensch (1992) compared expert and novice bridge players and examined the effects of various arbitrary rule changes on their performance. Sternberg and Frensch found that the more expert a player is, the more they struggle to adapt to rule changes. Research has also shown that expert accountants struggle more than novices at applying a new tax law (one that supersedes an old tax law) (Marchant et al., 1991). Kilov (2020) discusses studies of chess problem-solving (studies by Bilalić et al., 2008; Saariluoma, 1990). Chess players of various skill levels were presented with a series of chess puzzles. When the first four puzzles were solvable by the same motif, the expert players, recognising the pattern, applied the same approach to a fifth problem, overlooking the fact that it could be solved by objectively better means. Less experienced chess players, unaffected by the *Einstellung* stimulus<sup>5</sup>, were able to discover the better solution.

In these cases, and others, experts are thwarted by the very things that under normal circumstances are central to their success. Experienced firefighters will sometimes die fleeing a fire because they did not drop cumbersome equipment. Some naval officers fail to remove their steel-capped boots when abandoning ship and drown because they puncture holes in their life rafts. Members of both professions are explicitly instructed to avoid these mistakes. When interviewed, survivors describe the experience of dropping their tools as unnatural, uncomfortable, and wrong: their biggest liabilities in these scenarios are precisely those that normally facilitate their success (Weick, 1996).

Na (2006) discovered that Go expertise hampers creative problem solving in similar domains. Expert Go players, when faced with remote association tasks adapted from Mednick (1962), underperformed compared to novices due to their inability to dismiss irrelevant Go-related knowledge. Na concludes, this confirms the "negative effect of expertise on creative problem solving" (2006, p 65).

---

<sup>5</sup> *Einstellung* effects are negative effects brought about prior knowledge when solving novel tasks. The term was coined in Luchins (1942).



It is worth noting that the empirical record may understate the extent of expert brittleness. In most of the studies discussed above, the researchers went looking for flexibility and failed to find it. Brittleness may be so pervasive a feature of skill that we cannot help but stumble across it, even when we are not actively looking for it.

### 3.3 Why are expert skills brittle?

The precise reasons for the brittleness of expertise cannot be divined from the armchair. I will, nonetheless, offer two suggestions for why we should expect brittleness to be the norm, that is, why we should expect expert performance to fail so dramatically in novel situations. These suggestions are commensurate with existing research on the nature of skill. The first concerns the embodied and embedded nature of expert performance. The second concerns the nature of the internal, representational structures that mediate expert performance.

First, our abilities are embodied. Only a creature with a body like ours can learn to play the piano, for example. As Noë points out, ‘[d]ogs couldn’t manage the feat, and chimps, who might have the hands for the task, lack the brains’ (2005, p 284). Learning new skills changes our bodies to make them better suited to certain tasks, whether by building muscles or callouses or cortical reorganization.

Furthermore, our embodied performances depend on our being able to successfully and reliably couple with embedded contexts. One cannot ski without skis or snow-covered slopes. Thus, a natural explanation for the brittleness of skills is that situational factors are, in some sense, constitutive of skilful performance. Changing the situation robs performers of situational affordances and resources necessary for successful performance.

Noë is right to emphasise the dramatic way in which learning changes our bodies. However, another significant factor in the development of expertise is the development of more effective and efficient mental representations. Ericsson and Pool state as follows:

[D]eliberate practice will also lead to physical changes in the body itself – in divers, the development of the legs, abdominal muscles, back, and shoulders, among other body parts – but without the mental representations necessary to produce and control the body’s movements correctly, the physical changes would be of no use (2016, p 84).

Furthermore, the range of acceptable responses will generally shrink as task or situational demands increase. Consequently, the representational structures that mediate expert performance must become increasingly specialized. Here is an analogy. Pocket-knives are a useful general-purpose tool. They are useful in a wide range of situations. One can quite easily use a pocketknife to uncork a bottle of wine, but a pocketknife will never be as good as a corkscrew (that is why so many pocketknives have corkscrew attachments). However, the specialized design of the corkscrew

means that it is useless for almost any other task. The same features that make it good at one task explain why it is bad at many other tasks.<sup>6</sup>

As noted at the beginning of this section, additional research is needed to fully uncover the sources of expert brittleness. Kilov (2020) suggests that skills might share trade something like the trade-offs between generality and accuracy in biological and other scientific models. Lately, I have wondered whether expert brittleness might be like overfitting in machine learning<sup>7</sup>. Still, this is all just speculation and analogy. Future work could aim to provide a genuine theory of expert brittleness and its cognitive and ecological underpinnings.

I have argued that the spectre of brittleness accompanies would-be-experts like a shadow. It is present in the performances of novices and experts alike. Experts can respond to a dazzling array of situations within their domains, but this is because they have painstakingly acquired responses to ever more *recherché* parts of their domains of performance. However, this is evidence only of an extraordinary sphexibility – impressive, to be sure, but as we shall see in the next section, inadequate to ground an account of virtue-as-skill. This is the topic of the next section.

## 4 A new challenge to virtue-as-skill

### 4.1 Parallels to the situationist challenge

My argument bears a family resemblance to situationist challenges to virtue ethics. Situationism began with the observation that situational factors can result in significant changes in behaviour. Importantly, these situational factors often lead to people behaving badly. Previous versions of the situationist challenge emphasized what Miller (2015) calls ‘Surprising Dispositions’:

Research in social psychology seems to have uncovered a range of dispositions that could not be constituents of the moral virtues. In other words, Surprising Dispositions can result in unvirtuous behaviour. They can lead people to refrain from helping others or to generally engage in behaviour that is morally reprehensible. In Milgram’s (1963) famous obedience study, 65% of participants were willing to administer a lethal, 450-volt shock to someone when prompted to do so by an experimenter in a lab-coat with a clipboard (See Blass (1999) for a historical survey of 35 years’ worth of replications). Unlike some others<sup>8</sup>, these results have survived the replication crisis, having been recently replicated by Burger (2009). In addition, Bègue and colleagues (2015) conducted a Milgram-like experiment in the context of a fake television show called ‘La Zone Xtrême’, with Milgram-like results. These

---

<sup>6</sup> This highlights the important point that brittleness sometimes confirms expertise rather than calling it into question. In fact, if brittleness undercuts claims of expertise, then it would be impossible to discover that experts are brittle. All we could discover is that our putative experts are not experts. Nonetheless, brittleness is a *shortcoming* of experts.

<sup>7</sup> Overfitting produces agents that perform perfectly on the training data but falter in novel circumstances.

<sup>8</sup> See Alfano (2018). A plague on both your houses: Virtue theory after situationism and repligate, though as a reviewer notes, replication is not the only issue. Effect sizes are also key.

experiments suggest that much of our behaviour is driven by situational factors that activate certain mental dispositions.

Being surprising, Surprising Dispositions are causally operative outside our conscious awareness. Thus,

because they are (i) widely held, (ii) causally influential in many morally relevant situations, (iii) non-virtuous in their motivational and/or behavioral effects, and (iv) unconscious in their functioning, the Surprising Dispositions are a significant impediment to virtue cultivation (Miller 2015, p. 111).

Situationism also poses a challenge to virtue-as-skills theories. Stichter (2018) has developed the virtue-as-skills thesis in some detail, paying due attention to the empirical literature on expertise. Stichter's responds to Miller by invoking an argument from Sosa (2009). Sosa draws an analogy between moral competence and driving competence. He defines driving competence as 'a disposition to produce driving that is *safe*, when one is at the wheel, and *efficient* in routing to one's destination upon getting directions' (2009, p. 283). He then asks us to consider the range of factors that we have discovered influence the safety or efficiency of driving. These factors include weather conditions, the brightness of light, use of a cell phone, blood alcohol level, and so on. Sosa notes that it must have been surprising for us to learn how competent driving is influenced by some of these factors. He asks, '[h]ow should one's folk theory of driving respond to such discoveries?' (2009, p. 284) A possible response is *driving situationism*: situational factors have a significant influence on differences in driving behaviour. The robustness of our driving competencies then becomes problematic. Driving competencies vary surprisingly as a function of previously unsuspected factors. However, no one is tempted to conclude on the basis of these factors that driving competence does not exist. Although the discovery of these surprising factors reveals the inadequacy of our theories, this only calls for the revision of those theories and not for their wholesale abandonment.

According to Stichter, this analogy might also provide advocates of the virtue-as-skills thesis a ready path to rehabilitating virtue:

Despite the disappointing current levels of moral competence on display, the skill model gives us reasons to think that this situation can be improved... Skill acquisition is a matter of learning how to become more responsive to reasons, such as learning how to alter your driving behavior in response to certain environmental conditions (e.g. you have a good reason to drive slower in icy conditions). If we can do this in the case of skill, we should be able to do so in the case of virtue (2018, p.149).

However, as we shall now see, the problem is that the facts about the brittleness of skill militate against this approach.

## 4.2 The argument from brittleness

For advocates of the virtues-as-skills thesis, to possess a virtue is to have attained a degree of skill with respect to a particular domain. To what degree, exactly? Annas sets the standards quite high:

someone learning a practical skill like building will pick up bits of know-how and technique here and there; the expert, however, will have mastery of everything relevant to that kind of building, and will have unified that mastery so as to be able to understand his own and others' successes and mistakes, *and to be able to apply his skill in new situations without further learning being required* (2003, p. 18, emphasis added).

For Annas, possessing the kind of expertise that is coextensive with virtue requires more than that someone is able to respond intelligently to specific circumstances. It requires that they have mastered the entire domain and that they are able to apply their expertise in novel situations.

Other accounts are less demanding. Adams accepts that much learning to be virtuous will come in situational or domain-specific modules, as we learn to understand and appreciate relevant considerations only after some exposure to relevant situations (2006, p. 130).

That said, Adams does think that this domain-specificity has its limits. Calling a fine-grained skill such as "sailing-in-rough-weather-with one's-friends courage" 'courage' is 'ridiculous' if it describes the only form of courage a person possesses. If the circumstances in which the behaviour manifests are too specific, then we can doubt whether the person is sensitive to the evaluative commitments that virtue requires. This would again sever the link between possessing virtue and the explanation, prediction, and evaluation of behaviour, which is a primary motivation for advancing a virtue account in the first instance.

However, as we saw in § 3, empirical evidence suggest that skills are often acquired in a fine-grained, piecemeal way. Even experts are often unable to apply their skills in new situations without additional learning. Indeed, sometimes their expertise interferes with their performance in novel situations.

This gives rise to a new critique: the argument from brittleness. Because expertise is brittle, situational factors - here understood as changes to the domain of performance - will often cause expertise to disappear, and virtue along with it. Put differently, situationist critiques often emphasize virtuous traits' lack of intra-situational stability. In this respect, the argument from the brittleness of expertise is quite different. Experts typically *will* be reliable across iterations of the same situation. Indeed, such reliability is *the* hallmark of expertise. However, agents will lack cross-situational consistency because skill is brittle.

Furthermore, Sticher's reply to Miller in § 4.1 is a non-starter with respect to the argument from brittleness. This is so since, if the account I gave in § 3.3 is on the right track, brittleness is intrinsic to the nature of skill. It is not merely a matter of interference from hidden environmental factors. Sticher might reply that, although experts can initially underperform in novel circumstances, they adapt to novel situa-

tions faster than non-experts, and therein lies their virtue. This may sometimes be the case. However, as we've seen, expertise often *interferes with and delays* adaptation to novel circumstances. Consider the cases discussed in § 3.2 in this regard. Shiffrin and Schneider's (1977) experts took 167% longer to adapt to the reversed stimuli condition than untrained participants. Similar results have been obtained in other domains (Sternberg & Frensch, 1992; Marchant et al., 1991).

Much of expertise consists in developing automatic routines and habits. Learning to override these routines can be difficult. Further, when experts do successfully intervene on their finely honed routines, the result is not always positive. These experts sometimes fall prey to "choking" and "the yips"<sup>9</sup>.

## 5 Conclusion

The virtue-as-skill approach has many virtues, and it has been skilfully developed in recent years. But would-be-advocates of this approach find themselves between the devil and the deep blue sea. On the one hand, they could make a deal with the devil and abandon the idea that the virtuous agent can skilfully respond to varied and novel situations, thereby giving up the soul of virtue ethics. On the other hand, they could scuttle the ship and give up on the idea that virtues are expert skills, since expert skills are brittle.

**Acknowledgements** I am indebted to Mark Alfano for his mentorship with respect to this paper. Thanks also to Bronwyn Finnigan for her expert feedback on an early draft. I am grateful for insightful comments from Denise Metcalfe, James Willoughby, Maxwell Riess and Ross Pain. I also benefited greatly from the incisive feedback of John Doris and an anonymous reviewer from Synthese. Finally, my thanks to Alan Hájek, who provided invaluable feedback on every draft of this paper and who exemplifies many of the best virtues of our profession.

**Funding** Not applicable.

Open Access funding enabled and organized by CAUL and its Member Institutions

**Data Availability** Not applicable.

## Declarations

**Conflicts of interest/Competing interests** None.

**Open Access** This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

<sup>9</sup>For important studies on choking see Beilock and Carr (2001) and Baumeister (1984) for philosophical discussion of choking and the yips see Papineau (2015).

## References

- Adams, R. M. (2006). *A theory of virtue: Excellence in being for the good*. Oxford University.
- Alfano, M. (2013). *Character as moral fiction*. Cambridge University Press.
- Alfano, M. (2018). A plague on both your houses: Virtue theory after situationism and repligate. *Teoria*, 38(2), 115–122. <https://doi.org/10.4454/teoria.v38i2>.
- Annas, J. (2003). The structure of virtue. In M. DePaul, & L. Zagzebski (Eds.), *Intellectual virtue: Perspectives from ethics and epistemology* (pp. 15–33). Oxford University Press.
- Aristotle (2009). In Brown (Ed.), *Nichomachean ethics, L*. Oxford University Press.
- Baumeister, R. F. (1984). Choking under pressure: Self-consciousness and paradoxical effects of incentives on skillful performance. *Journal of personality and social psychology*, 46(3), 610.
- Bègue, L., Beauvois, J. L., Courbet, D., Oberlé, D., Lepage, J., & Duke, A. A. (2015). Personality predicts obedience in a Milgram paradigm. *Journal of Personality*, 83(3), 299–306. <https://doi.org/10.1111/jopy.12104>.
- Beilock, S. L., & Carr, T. H. (2001). On the fragility of skilled performance: What governs choking under pressure? *Journal of experimental psychology: General*, 130(4), 701.
- Bilalić, M., McLeod, P., & Gobet, F. (2008). Inflexibility of experts – reality or myth? Quantifying the Einstellung effect in chess masters. *Cognitive Psychology*, 56(2), 73–102.
- Blass, T. (1999). The Milgram paradigm after 35 years: Some things we now know about obedience to authority. *Journal of Applied Social Psychology*, 29(5), 955–978. <https://doi.org/10.1111/j.1559-1816.1999.tb00134.x>.
- Burger, J. M. (2009). Replicating Milgram: Would people still obey today? *American Psychologist*, 64(1), 1–11. <https://doi.org/10.1037/a0010932>.
- Caplan, B. (2018). *The case against education: Why the education system is a waste of time and money*. Princeton University Press.
- Chase, W. G., & Simon, H. A. (1973). Perception in chess. *Cognitive psychology*, 4(1), 55–81. [https://doi.org/10.1016/0010-0285\(73\)90004-2](https://doi.org/10.1016/0010-0285(73)90004-2).
- Christensen, W., Sutton, J., & Bicknell, K. (2019). Memory systems and the control of skilled action. *Philosophical Psychology*, 32(5), 692–718. <https://doi.org/10.1080/09515089.2019.1607279>.
- Clement, J. (1982). Students' preconceptions in introductory mechanics. *American Journal of Physics*, 50(1), 66–71. <https://doi.org/10.1119/1.12989>.
- Dane, E. (2010). Reconsidering the trade-off between expertise and flexibility: A cognitive entrenchment perspective. *Academy of Management Review*, 35(4), 579–603. <https://doi.org/10.5465/amr.35.4.zok579>.
- Dennett, D. C. (1978). *Brainstorms: Philosophical essay on mind and psychology*. Harvester Press.
- Dent, N. J. H. (1975). Virtues and actions. *The Philosophical Quarterly*, 25(101), 318–335. <https://doi.org/10.2307/2218976>.
- Doris, J. M. (2002). *Lack of character: Personality and moral behavior*. Cambridge University Press.
- Doris, J. M. (2022). *Character trouble: Undisciplined essays on moral agency and personality*. Oxford University Press.
- Dougherty, M. (2020). The importance of roles in the skill analogy. *Journal of Ethics & Social Philosophy*, 17(1), 75–102.
- Dreyfus, S. E., Dreyfus, H. L., & Athanasiou, T. (1986). *Mind over machine: The power of human intuition and expertise in the era of the computer*. The Free Press.
- Ericsson, A., & Pool, R. (2016). *Peak: Secrets from the new science of expertise*. Houghton Mifflin Harcourt.
- Fabre, J. H., & de Texeira, A. (1915). *The hunting wasps*. Dodd, Mead and Company.
- Feltovich, P. J., Prietula, M. J., & Ericsson, K. A. (2006). Studies of expertise from psychological perspectives. In K. A. Ericsson, N. Charness, R. R. Hoffman, & P. J. Feltovich (Eds.), *The Cambridge handbook of expertise and expert performance* (pp. 41–67). Cambridge University Press.
- Gardner, H. (1991). *The unschooled mind: How children learn and how schools should teach*. Basic Books.
- Harman, G. (1999). Moral philosophy meets social psychology: Virtue ethics and the fundamental attribution error. *Proceedings of the Aristotelian Society*, 99(1), 315–332. <https://doi.org/10.1111/1467-9264.00062>.
- Harman, G. (2000). The nonexistence of character traits. *Proceedings of the Aristotelian Society*, 100(1), 223–226. <https://doi.org/10.1111/j.0066-7372.2003.00013.x>.

- Hofstadter, D. R. (1985). On the seeming paradox of mechanizing creativity. In D. R. Hofstadter (Ed.), *Metamagical themas: Questing for the essence of mind and pattern* (pp. 526–546). Basic Books.
- Horowitz, B. (2004). “Softball’s Jennie Finch ready to make her pitch.” *USA Today*, September 9 2004. [https://usatoday30.usatoday.com/sports/olympics/athens/2004-08-09-finch\\_x.htm](https://usatoday30.usatoday.com/sports/olympics/athens/2004-08-09-finch_x.htm).
- Keijzer, F. (2013). The Sphinx story: How the cognitive sciences kept repeating an old and questionable anecdote. *Philosophical Psychology*, 26(4), 502–519. <https://doi.org/10.1080/09515089.2012.690177>.
- Kilov, D. (2020). The brittleness of expertise and why it matters. *Synthese*, 199(1–2), 3431–3455. <https://doi.org/10.1007/s11229-020-02940-5>.
- Liaisons: Philosophy meets the cognitive and social sciences* (pp. 155–175). MIT Press.
- Luchins, A. S. (1942). Mechanization in problem solving: The effect of Einstellung. *Psychological Monographs*, 54(6), i–95. <https://doi.org/10.1037/h0093502>.
- MacIntyre, A. (1984). *After virtue: A study in moral theory* (2nd ed.). University of Notre Dame Press.
- Marchant, G., Robinson, J., Anderson, U., & Schadewald, M. (1991). Analogical transfer and expertise in legal reasoning. *Organizational Behavior and Human Decision Processes*, 48(2), 272–290.
- Mednick, S. (1962). The associative basis of the creative process. *Psychological Review*, 69(3), 220.
- Milgram, S. (1963). Behavioral study of obedience. *The Journal of Abnormal and Social Psychology*, 67(4), 371–378.
- Miller, C. B. (2014). *Character and moral psychology*. Oxford University Press.
- Miller, C. B. (2015). Russell on acquiring virtue. In M. Alfano (Ed.), *Current controversies in virtue theory* (pp. 106–117). Routledge.
- Millikan, R. G. (2004). *Varieties of meaning: The 2002 Jean Nicod lectures*. MIT Press.
- Morrot, G., Brochet, F., & Dubourdieu, D. (2001). The color of odors. *Brain and Language*, 79(2), 309–320. <https://doi.org/10.1006/brln.2001.2493>.
- Na, S. U. (2006). The effect of expertise on creative problem solving. *The International Journal of Creativity & Problem Solving*, 16(1), 65–74.
- Narvaez, D. (2005). The neo-Kohlbergian tradition and beyond: Schemas, expertise and character. In G. Carlo & C. Pope-Edwards (Eds.), *Nebraska Symposium on Motivation*, 51, 119–163. University of Nebraska Press.
- Nodine, C. F., & Krupinski, E. A. (1998). Perceptual skill, Radiology Expertise, and visual test performance with NINA and WALDO. *Academic Radiology*, 5(9 (September)), 603–612.
- Noë, A. (2005). Against intellectualism. *Analysis*, 65(4), 278–290. <https://doi.org/10.1093/analys/65.4.278>.
- Olin, L., & Doris, J. M. (2014). Vicious minds: Virtue epistemology, cognition, and skepticism. *Philosophical Studies*, 168(3), 665–692. <https://doi.org/10.1007/s11098-013-0153-3>.
- Papineau, D. (2015). Choking and the yips. *Phenomenology and the Cognitive Sciences*, 14, 295–308.
- Park, J., MacRae, H., Musselman, L. J., Rossos, P., Hamstra, S. J., Wolman, S., & Reznick, R. K. (2007). Randomized controlled trial of virtual reality simulator training: Transfer to live patients. *The American journal of surgery*, 194(2), 205–211.
- Saariluoma, P. (1990). Apperception and restructuring in chess players’ problem solving. In K. J. Gilhooly, M. T. G. Keane, R. H. Logie, & G. Erdos (Eds.), *Lines of thinking: Reflections on the psychology of thinking* (2 vol., pp. 41–57). John Wiley & Sons.
- Saariluoma, P. (1991). Aspects of skilled imagery in blindfold chess. *Acta Psychologica*, 77(1), 65–89.
- Shiffrin, R. M., & Schneider, W. (1977). Controlled and automatic human information processing: II. Perceptual learning, automatic attending and a general theory. *Psychological Review*, 84(2), 127–190. <https://doi.org/10.1037/0033-295X.84.2.127>.
- Sosa, E. (1980). The raft and the pyramid: Coherence versus foundations in the theory of knowledge. *Midwest Studies in Philosophy*, 5(1), 3–26. <https://doi.org/10.1111/j.1475-4975.1980.tb00394.x>.
- Sosa, E. (2009). Situations against virtues: The situationist attack on virtue theory. In C. Mantzavinos (Ed.), *Philosophy of the social sciences: Philosophical theory and scientific practice* (pp. 274–290). Cambridge University Press.
- Sterelny, K. (1990). *The representational theory of mind: An introduction*. Blackwell.
- Sternberg, R. J., & Frensch, P. A. (1992). On being an expert: A cost-benefit analysis. In R. R. Hoffman (Ed.), *The psychology of expertise* (pp. 191–203). Springer.
- Stichter, M. (2018). *The skillfulness of virtue: Improving our moral and epistemic lives*. Cambridge University Press. <https://doi.org/10.1017/9781108691970>.

- Van Sickle, K. R., Ritter, E., McClusky, D., Lederman, A., Baghai, M., Gallagher, A., & Smith, C. (2007). Attempted establishment of proficiency levels for laparoscopic performance on a national scale using simulation: The results from the 2004 SAGES minimally invasive Surgical trainer—virtual reality (MIST-VR) learning center study. *Surgical endoscopy*, *21*(1), 5–10.
- Voss, J. B., Blais, J., Means, M. L., Greene, T. R., & Ahwesh, E. (1986). Informal reasoning and subject matter knowledge in the solving of economics problems by naive and novice individuals. *Cognition and Instruction*, *3*(3), 269–302. [https://doi.org/10.1207/s1532690xci0303\\_7](https://doi.org/10.1207/s1532690xci0303_7).
- Wanzel, K. R., Hamstra, S. J., Anastakis, D. J., Matsumoto, E. D., & Cusimano, M. D. (2002). Effect of visual-spatial ability on learning of spatially-complex surgical skills. *The lancet*, *359*(9302), 230–231.
- Weick, K. E. (1996). Drop your tools: An allegory for organizational studies. *Administrative science quarterly*, 301–313.
- Winter, M., & Tauer, J. (2006). Virtue theory and social psychology. *J Value Inquiry*, *40*, 73.

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.