

# Modal Virtue Epistemology

BOB BEDDOR<sup>1</sup>

*National University of Singapore*

CARLOTTA PAVESE

*Duke University*

## 1. A Tale of Two Theories

Since the publication of Gettier (1963), many accounts of knowledge have been proposed, but few have withstood the test of time. Of the various proposals, two have emerged from the fray in relatively good shape, although not entirely unscathed. One of these is a *modal account*. While there are different ways of developing a modal account, the most common proposal in the contemporary literature explains knowledge in terms of *safety*. More precisely:

**Safety:** S's belief amounts to knowledge if and only if *it could not have easily been false* — i.e., if and only if in all relevantly close worlds where S forms a sufficiently similar belief, S's belief is true.<sup>2,3</sup>

The other major contender is virtue epistemology. The guiding idea behind this approach is that knowledge is a type of skillful performance. Perhaps the most well-developed version of this idea comes from Sosa (2007, 2015), who distinguishes various normative statuses that performances can attain. For a performance to be *accurate* is for it to achieve its aim; for it to be *adroit* is for it to manifest a skill. Finally, for it to be *apt* is for it to be accurate because it is adroit — that is, for it to succeed in virtue of the manifestation of a skill. Virtue epistemologists propose that knowledge is apt belief:

**Aptness-Based Virtue Epistemology (AVE):** S's belief amounts to knowledge if and only if it is accurate in virtue of manifesting an epistemic skill.<sup>4</sup>

---

<sup>1</sup> The order of the authors is alphabetical.

<sup>2</sup> While some proponents of a modal account seem to embrace both the necessity and sufficiency of safety for knowledge (Pritchard 2005; Williamson 2009; Lasonen-Aarnio 2010), others defend only the necessity (Sosa 1999; Manley 2007; Pritchard 2012, 2016).

<sup>3</sup> Here the “sufficiently similar beliefs” are beliefs on relevantly similar subject matters, which are formed *via* sufficiently similar methods, or on the basis of sufficiently similar bodies of evidence.

<sup>4</sup> See especially Sosa (2007, 2015); see also Zagzebski (1996); J. Greco (2010); Turri (2011). For rather different ways of developing aptness-based virtue epistemology, see Kelp (2017); Wedgwood (2018).

Which approach should we prefer? This has proven difficult to adjudicate, since each approach holds certain explanatory advantages over the other. The modal approach captures intuitions about various cases — in particular, lotteries and fake barns — that virtue epistemology struggles to accommodate. At the same time, virtue epistemology avoids a number of *prima facie* counterexamples to **Safety**; it also sheds light on the normative parallels between knowledge and skilled action.

Given this trade-off, some have questioned whether we need to choose between the two approaches. Why not develop a synthesis? This ambition has led some to advocate a hybrid approach, according to which knowledge requires both safety and aptness (Pritchard 2012, Kelp 2013). While we sympathize with this synthetic impulse, we argue that hybrid approaches implicitly abandon one of the chief attractions of virtue epistemology, which is that it offers a unified account of knowledge and skillful action. In brief, the problem is that hybrid approaches do not derive the safety condition on knowledge from any property shared by knowledge and skillfulness.

In view of this difficulty, this paper develops an alternative synthesis: “Modal Virtue Epistemology.” It is a form of virtue epistemology in that it takes knowledge to be a type of skillful performance. But it is also a modal approach, since it goes on to understand both skillfulness and knowledge in purely modal terms — specifically, in terms of success across a range of counterfactual scenarios. We show that this synthesis preserves the benefits of both **Safety** and AVE, while overcoming the main difficulties for both. It also avoids the main challenge to hybrid approaches, since it takes both knowledge and skillfulness to be reducible to a shared modal quantity.<sup>5</sup>

## 2. Tradeoffs: Safety vs. AVE

### 2.1 Advantages of Safety

A primary motivation for **Safety** is that it captures the intuition that knowledge is, in some sense, incompatible with luck (Pritchard 2005, 2012, 2016). As a result, it offers a promising diagnosis of many Gettier cases that have bedeviled alternative analyses of knowledge. Take the fake barns case (Goldman 1976): Henry is driving through fake barn county and happens to glance at the only real barn in the region. He forms the belief that he is looking at a barn. **Safety** captures the intuition that his belief does not qualify as knowledge. He could have easily looked out the window a moment earlier; had he done so, he would have formed the same belief, but it would have been false. Or consider lottery cases (Williamson 2000; Hawthorne 2004; Pritchard 2005): Lottie believes, on statistical grounds, that her ticket is a loser. As a matter of fact, it is. **Safety** explains the intuition that Lottie’s belief does not qualify as knowledge. After all, there is at least one relevantly close world where she holds this belief but her ticket wins.

As a number of writers have noted, it is less clear that virtue epistemology captures these intuitions.<sup>6</sup> In the fake barns case, Henry forms a correct belief about whether he is looking at

---

<sup>5</sup> We offer Modal Virtue Epistemology as *one* way of developing a unified account of skillfulness and knowledge that fares better than extant approaches. A very different unification strategy would be to adopt an intellectualist account of skills, according to which skills are themselves defined in terms of knowledge (Pavese 2013; 2016a,b; Stanley and Williamson 2017). A detailed comparison between these two strategies will be left to future work.

<sup>6</sup> Lackey (2007); Pritchard (2012); Kelp (2013).

a barn as the result of an epistemic skill (*vision*). His belief is thus apt, and hence — by the lights of **AVE** — qualifies as knowledge. Similarly, Lottie’s belief that her ticket will lose is based on careful consideration of the statistical odds. Assuming Lottie judges the odds accurately, she arrives at the truth in virtue of exercising an epistemic skill. So her belief also qualifies as apt, and hence knowledge, by the lights of **AVE**.

Of course, some virtue epistemologists may reject intuitions about fake barns and lotteries. Recently, Sosa (2015: 119) has challenged lottery intuitions by providing a structurally analogous case where knowledge is arguably present. In Sosa’s scenario, you call into Apple’s support center and are assigned to one of their many reliable operators. However, you could easily have been randomly assigned to a lone disgruntled employee who dispenses misleading information. Sosa contends that you can come to know the answer to your question, despite the nearby possibility of being misled.

However, it would be premature to reject lottery intuitions on this basis alone. The intuition that Lottie does not know is widely attested among both philosophers and non-philosophers.<sup>7</sup> Moreover, it is revealing that Sosa’s case is formulated in terms of knowing *the answer to a question* (p.119), as opposed to knowing some proposition (e.g., *that the caller needs to download the latest operating system*). As some authors have observed, we tend to have laxer standards for the former sort of knowledge ascription (Hawthorne 2000; Stanley 2011). Suppose Fred needs to get to Boston and is looking for directions. We might be happy to tell him that John knows what the best route is, even if we are aware that John learned the route from an unreliable source. Exactly what we should conclude from this is a delicate matter. But one tempting option is to say that in such cases the agent does not, strictly speaking, know the answer to the question. However, the knowledge ascription is pragmatically acceptable, since the main purpose of the ascription is to say who can give the *correct* answer to the question.<sup>8</sup>

Barn cases and lotteries aside, there is independent theoretical reason to impose a safety requirement on knowledge (D. Greco 2016; Pavese 2018). A number of epistemologists have suggested that knowledge plays a central role in explaining successful action: the fact that the burglar *knows* that the jewels are in the house provides a better explanation of the burglar’s successful discovery of the jewels than the fact that he merely believes it, or even truly believes it (Williamson 2000; Nagel 2013). In philosophy of science it is often held that good explanations are *robust*, in the sense that they show that their explananda are not mere flukes: a good explanation of some outcome *o* will predict *o*’s occurrence across a range of relevantly close circumstances. If knowledge requires safety, then we have a clear explanation for why knowledge provides a robust — hence good — explanation of successful action. By contrast, if knowledge does not require safety, the robustness of the relevant explanations will be in jeopardy. (If a belief is unsafe, then there will be relevantly close circumstances where the belief is false. And in those circumstances, the action will be unsuccessful.) Thus views that

---

<sup>7</sup> For experimental work suggesting that ordinary speakers share the lottery intuition, see Turri and Friedman (2014).

<sup>8</sup> Whereas lottery intuitions are attested among non-philosophers, experimental work on fake barns tells another story (Colaço et al. 2014). Some may think that these results lend support to the virtue epistemological verdict that Henry’s belief amounts to knowledge (cf. Sosa 2007). However, two points are worth making. First, many epistemologists report having very robust “No-knowledge” intuitions about fake barn cases. Second, one virtue of **Safety** is that it is able to explain why there are divergent judgments in these cases. After all, there are different ways of demarcating the set of relevantly close worlds — a point which we will explore in more detail in §5.

renounce a safety requirement on knowledge have a comparatively difficult time underwriting the explanatory role of knowledge attributions.

## 2.2 Advantages of Virtue Epistemology

At the same time, virtue epistemology seems to have some important advantages over **Safety**. First, it offers an axiological framework that applies to both knowledge and actions. In doing so, it sheds light on the normative parallels between the two. To illustrate, consider the following practical analogues of epistemic performances (cf. Sosa 2007):

- *Wayward Gust*: Wendy is a competent archer who shoots at a target. However, an unforeseen gust of wind diverts her arrow off-course.
- *Double Gust*: Like Wendy, Don is a competent archer who shoots at a target. His arrow is also blown off-course by an unforeseen gust. However, a second gust of wind intervenes at the last moment, redirecting his shot towards the bullseye.
- *No Gust*: Norm is a competent archer who shoots at a target in normal conditions (no wind, clear visibility, etc.). His shot lands straight on the bullseye.

Wayward gust is the practical analogue of a justified false belief; Double Gust is the practical analogue of a Gettiered belief; No Gust is the practical analogue of knowledge. **AVE** explains these parallels by identifying a normative status that each of these practical performances shares with its epistemic analogue. Wendy's performance — much like a justified false belief — is adroit but inaccurate. Don's performance — much like a Gettiered belief — is adroit and accurate but inapt. Finally, Norm's performance, much like ordinary knowledge, is apt: it attains its aim in virtue of skill.

By contrast, the modal approach does not explain these parallels. As it stands, **Safety** is entirely silent about the normative statuses of actions. The obvious way of trying to extend **Safety** to the practical domain is to insist that all skillful performances are safe — i.e., they succeed in all of the relevantly close worlds where they are attempted. This would capture the intuition that something is amiss in Double Gust: Don's shot could easily have failed, had not the second wind intervened. However, this would fail to capture the intuition that an action can be skillful despite being inaccurate (as in Wayward Gust). The worry, then, is that we can only do justice to the normative parallels with skilled action by adopting the rich axiology of virtue epistemology.

Even if we set aside skillful action and focus on knowledge, many have objected to both the necessity and sufficiency of safety. Start with necessity: a number of philosophers have offered cases in which it seems that an agent knows  $p$ , even though there is a nearby world where they form a false belief about  $p$ .<sup>9</sup> To illustrate with an example from Kelp (2009): at 3:00 Russell forms a belief about the time by looking at a clock. The clock is in perfect working order. However, there is a demon waiting in the wings, who desperately wants Russell to believe it is 3:00. Had Russell glanced at the clock a minute earlier or a minute later, the demon would have manipulated the hands of the clock to incorrectly display, "3:00." Intuitively, Russell knows it is 3:00, even though there's a

---

<sup>9</sup> See e.g., Neta and Rohrbaugh (2004); Comesaña (2005); J. Greco (2007); Baumann (2008); Kelp (2009); Bogardus (2014), among others.

nearby world where he is misled. Here **AVE** appears to be in a better position to explain our intuitions: since the demon did not intervene, the correctness of Russell's belief is due to the exercise of an epistemic skill.

On to sufficiency: Pritchard (2012) gives the example of Temp, who forms beliefs about the temperature by consulting a broken thermometer. However, Temp has a guardian angel in the room who controls the thermostat, ensuring that the room's temperature matches the reading displayed on the thermometer. Assuming the angel manipulates the thermostat in all nearby worlds, any belief Temp forms about the temperature in such worlds will be true. But intuitively Temp's thermal beliefs do not amount to knowledge. Here too, **AVE** offers an explanation: the accuracy of Temp's belief is not due to his epistemic skill but rather to the angel's interventions.

### 2.3 Going Hybrid?

Give these trade-offs, it is natural to hope for an approach that combines the advantages of **Safety** and **AVE**. This hope motivates hybrid approaches, which take knowledge to require both an aptness condition and a safety condition (Pritchard 2012; Kelp 2013). For example, Pritchard defends:

**Anti-Luck Virtue Epistemology (ALVE):** S's belief amounts to knowledge if and only if it is safe in virtue of the exercise of an epistemic skill.

While this is an attractive idea, hybrid approaches face at least two drawbacks. First, they still face the counterexamples to the necessity direction of **Safety** (since they take safety to be necessary for knowledge). Second, recall that a chief advantage of **AVE** is that it provides a unified account of knowledge and skillful action. At first blush hybrid approaches seem to share this advantage: according to **ALVE**, the property shared by knowledge and skillful action is their aptness. However, notice that the modal dimension of knowledge (its safety) does not follow from this shared property. Rather, safety is taken to be a *special* feature of knowledge — one that distinguishes knowledge from other kinds of skillful performances. Indeed, this is precisely what makes **ALVE** a *hybrid* of the modal and virtue epistemological approaches. However, we will now argue that skillful actions have a modal dimension of their own. Hence a truly unified account of knowledge and skillful action should account for *both* of their modal dimensions. Moreover, it should do so by deriving these modal dimensions from whatever knowledge and skillful action have in common.

### 3. The Modal Dimension of Skillful Action

While the idea that knowledge has a modal dimension has been widely explored, the idea that skillfulness also has a modal dimension has received far less attention. To motivate this claim, start by considering two archers, an expert and an incompetent novice. *Ceteris paribus*, we would expect their difference in skill to result in a difference in success rate: the expert will hit the target more frequently than the novice, at least in normal circumstances.

Indeed, it is standard to rank competitors in various endeavors on the basis of their success rates. To give just a few examples, consider the Hollinger player efficiency rating (PER) in basketball, the Batter-Fielder/Pitcher Wins (BFW/PW) in baseball, and the Elo

rating system in chess. While they differ in their details, all of these systems rank competitors using some function of their successes and failures.

Success rates are also used by the cognitive sciences to measure cognitive and motor skills. For example, the ability to reposition a screen cursor from a common central origin to a series of peripheral targets is measured by the subject's rate of success across trials above baseline. And a subject is counted as having the ability to pick up a glass on a table if they can repeatedly pick it up at varying speeds (Krakauer et al. 2000).

This suggests that skillfulness has a statistical dimension: in a wide variety of domains we use an individual's rate of success to measure their skillfulness. However, we should not identify skillfulness with the *actual* frequency of success. Consider the novice archer who only attempts to hit a target on a single occasion. By luck, they manage to strike the bullseye. Their performance is not skillful, even though they have a perfect success rate.

The natural solution is to go modal: what really matters for skillfulness is success rate across a range of counterfactual circumstances. That is:

**Skillfulness as Modal Robustness:** S's performance is skillful if and only if it is *modally robust* — i.e., in a sufficiently high proportion of relevantly close worlds where S engages in a sufficiently similar performance, their performance attains its aim.<sup>10</sup>

Here what counts as a “sufficiently high proportion” will vary with the task in hand. For very demanding tasks such as *attempting to hit a home run*, success in only, say, 8% of the relevantly close worlds may qualify as sufficiently high.<sup>11</sup>

Our primary argument for **Skillfulness as Modal Robustness** is that it explains the statistical dimension of skillful performance. For example, it explains the difference in skillfulness between the expert and the novice: the expert will hit the bullseye in a higher proportion of the relevantly close worlds than the novice will. A second, and related, consideration in favor of **Skillfulness as Modal Robustness** is that it explains the *gradability* of skillfulness. Chess grandmasters are more skilled than chess masters, who are in turn more skilled than an average chess club player. **Skillfulness as Modal Robustness** captures these differences: one's degree of skillfulness at a task is reflected in the proportion of relevantly close worlds in which one succeeds at the task.

A final argument for **Skillfulness as Modal Robustness** comes from the thought that there is an important connection between skills and *dispositions*. This thought is championed by virtue epistemologists, who claim that a performance is adroit if it manifests a disposition to achieve its aim. But how should we understand dispositions in the first place? In the contemporary literature, one prominent approach is to analyze dispositions in modal terms (Manley and Wasserman 2008; Vetter 2014; Aimar 2018). On Manley

---

<sup>10</sup> Cf. J. Greco (2010: 77), who defends a modal account of ability.

<sup>11</sup> Some might worry that talk of a “sufficiently high proportion of relevantly close worlds” will not be well-defined when there are infinitely many relevantly close worlds. In response, one option is to partition an infinite set of worlds into a finite number of cells. We could then restrict our attention to the cells containing only worlds where S engages in a sufficiently similar performance and evaluate whether S's performance succeeds in a sufficiently high proportion of these cells. A related option is to go probabilistic: once we've partitioned the worlds, we could define an objective probability measure over the cells in this partition. We could then say S's performance is modally robust if it has a sufficiently high objective probability of success. For the rest of this paper, we will set this complication aside; see Manley and Wasserman (2008: 81-82) for discussion of related issues.

and Wasserman's (2008) account,  $x$  is disposed to  $\phi$  if and only if  $x \phi$ s in a sufficiently high proportion of the worlds where some stimulus condition obtains. For example, a vase is fragile (disposed to break when dropped) provided it breaks in a sufficiently high proportion of the worlds where it is dropped.

Suppose that we combine a modal account of dispositions with the virtue epistemological idea that performances are skillful if and only if they manifest a disposition to achieve their aim. This in turn delivers **Skillfulness as Modal Robustness**.

Some might object to **Skillfulness as Modal Robustness** on the grounds that it has trouble accommodating *exceptional performances*. For example, in 1979 Pietro Mennea shattered the world record for the 200 meter sprint, clocking in at an incredible 19.72 seconds. His performance was exceptional, in that for many years nobody came close to repeating it. But surely it was skillful. This may be thought to cause trouble for **Skillfulness**: in virtually all of the relevantly close worlds where Mennea attempted this performance, he failed.

In response, observe that whether an action counts as skillful depends on how the action is described. Under the description, *running the 200m sprint within 19.72 seconds*, Mennea's performance is not skillful, since he does not succeed in this activity in a suitably high proportion of relevantly close worlds. However, under other descriptions the same performance will count as skillful: for example, *running the 200m faster than most athletes alive in 1979*, or *running the 200m within 19.80 seconds*. This allows us to preserve the intuition that Mennea's performance is skillful, since there is a sufficiently competitive description of the performance under which it succeeds in a sufficiently high proportion of relevantly close worlds.

#### 4. Introducing Modal Virtue Epistemology

In the last section we argued that skillfulness, like knowledge, has a modal dimension. Thus it is a desideratum on any unified account of the two that it predicts their respective modal dimensions. Hybrid accounts such as **ALVE** fail to satisfy this desideratum.

To clarify the problem, it may help to distinguish between a superficial worry and a deeper worry. The superficial worry is that, as it stands, **ALVE** does not predict that skillful action requires modal robustness. Here proponents of **ALVE** could reply that a modal robustness requirement follows from the aptness condition: aptness requires adroitness, which could in turn be cashed out in modal terms, courtesy of a modal account of dispositions. But this brings us to a deeper worry, which is that this strategy does not derive the modal properties of skillful action and knowledge from any common element. According to **ALVE**, the property shared by knowledge and skillful action is their aptness. And while the modal dimension of skillful action may follow from this property, the modal dimension of knowledge (safety) does not.

This difficulty also points towards a more promising way of synthesizing the modal and the virtue epistemological accounts — one that avoids the additional complexity of the hybrid view. Perhaps what really unites knowledge and skillful action is not aptness but rather their shared modal dimension.

To develop this thought, suppose that we combine **Skillfulness as Modal Robustness** with the virtue epistemological idea that belief is a type of performance, which can be evaluated in terms of whether it manifests an epistemic skill. We can then propose that a

belief amounts to knowledge just in case it exhibits the highest degree of epistemic skillfulness:

**Knowledge as Maximally Skillful Performance:** S knows  $p$  if and only if S has a maximally skillful belief that  $p$ .

Let **Modal Virtue Epistemology (MVE)** be the conjunction of **Knowledge as Maximally Skillful Performance** with **Skillfulness as Modal Robustness**.<sup>12</sup>

This is the first-pass version of our view; in §5, we will introduce some complications and refinements. But before plunging into complications, it may be helpful to discuss how our view differs from more familiar forms of virtue epistemology.

The first major difference between **MVE** and **AVE** is that **MVE** entails **Safety**. After all, a belief is maximally skillful if and only if it is maximally modally robust (by **Skillfulness as Modal Robustness**). But for a performance to be *maximally* modally robust is for it to succeed in *all* of the relevantly close worlds. And a belief will satisfy this condition just in case it is safe. **MVE** thus enjoys all the explanatory benefits that come with a safety requirement.

At the same time, **MVE** also captures the modal dimension of skillful action. And unlike **ALVE**, it derives the modal dimensions of knowledge and skillful action from the common element uniting the two. According to **MVE**, this common element is a modal quantity that comes in degrees (modal robustness). What makes knowledge distinctive is that it possesses this quantity to the highest degree. For **MVE**, then, the difference between the skillfulness that makes for knowledge and the skillfulness that makes for other kinds of skillful action is one of degree rather than kind.

The final difference between **MVE** and more traditional versions of virtue epistemology is that **MVE** makes no appeal to aptness, or the *in virtue of* relation more generally. Instead, it explains knowledge and skillfulness in purely modal terms.

**MVE** is thus importantly different from traditional forms of virtue epistemology. But it gives rise to serious challenges. As we have seen, **MVE** entails **Safety**. But as we saw in §2.2, **Safety** faces significant hurdles. Can **MVE** surmount these obstacles? This will be the topic of the next section, where we explore how to develop **MVE** in a way that makes progress on the challenges to the sufficiency of safety for knowledge (§5.1), the challenge of explaining the normative parallels between the epistemic and practical domains (§5.2), and the challenges to the necessity of safety for knowledge (§5.3).

## 5. Defending Modal Virtue Epistemology

### 5.1 Is Modal Robustness Sufficient for Knowledge and Skillfulness?

As we saw in §2.2, the idea that safety is sufficient for knowledge faces a powerful counterexample in Pritchard's (2012) case of Temp, who forms beliefs about the temperature on the basis of a broken thermometer, while the thermostat is manipulated by a

---

<sup>12</sup> The idea that knowledge is maximally skillful will appeal to those who think of knowledge as the highest epistemic achievement. But some might resist this thought, on the grounds that there can be even more exalted epistemic statuses, such as epistemic certainty. For the purposes of this paper, we will set this issue aside. However, see Beddor (ms.) for one way of cashing out the relation between knowledge and epistemic certainty in modal terms.

guardian angel. Similar cases can be used to cast doubt on the sufficiency of modal robustness for skillfulness. Suppose Marie is terrible at shooting hoops. But whenever she attempts a shot, her guardian angel ensures that she makes the basket. Assuming her guardian angel assists her in all nearby worlds, her performance will have a perfect success rate. Nonetheless, it is not skillful (cf. Sosa 2007: 29).

Is there any way to defend the sufficiency of modal robustness for skillfulness and knowledge in the face of these counterexamples? Let us start with Marie. Note that her case is structurally analogous to what are known as “mimics” in the dispositions literature: it is a case in which something lacks a disposition to  $\phi$  but — due to peculiar features of the circumstances — mimics the behavior of something that is disposed to  $\phi$ . Manley and Wasserman (2008) give the example of a concrete block that has incurred a sorcerer’s wrath. Were the block ever to be dropped, the sorcerer would transform the ground into a rapidly accelerating diamond surface. The block is not fragile but it mimics the behavior of a fragile object. Marie’s case is similar. She lacks the disposition to sink baskets but, thanks to her guardian angel, her behavior is indistinguishable from that of someone who has this disposition.

This suggests that the problem is not specific to knowledge or skillfulness *per se*. Rather, it is a problem that arises for dispositions more generally. And this in turn suggests a natural strategy for making progress: consider how modal accounts of dispositions handle mimics, and see whether a similar strategy can be used to handle Marie.

On a modal account, the concrete block is fragile if and only if it breaks in a sufficiently high proportion of the worlds where it is dropped. As Manley and Wasserman observe, the block only shatters in sorcerer-inhabited worlds, which constitute a small fraction of the possible worlds where the block is dropped. As long as we take the sorcerer-free worlds into consideration, modal accounts of dispositions have no trouble handling mimics.

The crucial idea behind Manley and Wasserman’s suggestion is that dispositions quantify over worlds above and beyond the nearest worlds. This strategy extends to the case of Marie. As long as we take the *relevantly* close worlds to be a sufficiently large set of worlds that includes many fairly distant worlds, we correctly predict that her shot is not skillful. After all, many of those worlds will be angel-free. At those worlds, her shot invariably fails.

Pursuing this strategy raises a question: if the relevantly close worlds are not just the nearest worlds, what determines which worlds count as relevantly close? Here is one suggestion. When we assess a task such as shooting hoops, we implicitly associate the task with a set of conditions that we take to be *normal* for its performance and assessment. What sense of normality is at issue? We should resist identifying normal conditions with those that are *favorable*, in the sense of being conducive to the attainment of the aim of the task. Otherwise, Marie’s shot would count as skillful; after all, what is more favorable than having an angel’s assistance? We should likewise resist any purely statistical conception of normality: in a world where most aspiring athletes have guardian angels, we would still not regard Marie’s shot as skillful. A more promising option is to identify the normal conditions for a task with those which we would consider to be *fair* for performing and assessing the task. We suspect it will be difficult to give a precise, non-circular analysis of what these conditions consist in. However, our intuitions about cases reveal a tacit grasp of these conditions. For example, in the case of shooting hoops, fair conditions include freedom from external intervention, sufficient lighting, etc. Moreover, it

seems that we manifest our tacit conception of fair conditions when setting up competitions of various sorts.

This suggests a *sufficient* condition for relevance: the worlds that are relevant for assessing the skillfulness of a task include the nearest worlds where conditions are *normal* for the task at hand. To ensure the actual world is always among the relevantly close worlds,<sup>13</sup> we propose identifying the worlds that are relevantly close to  $w$  with the worlds where conditions are at least as normal for the task at hand as those that obtain at  $w$ . Applied to Marie: the worlds where conditions are *at least as normal* for shooting hoops will include all the worlds where conditions are *more normal* for shooting hoops. And so they will include worlds where Marie takes a shot without angelic assistance. At all of these worlds, Marie's performance fails.

Some may worry that this proposal rules out the possibility of agents who are only skilled at succeeding in *abnormal situations*. Consider Archie the archer: in normal circumstances, his performance is middling at best. But it drastically improves in gale-force winds. Is there not some sense in which Archie is skilled? Much depends on how we individuate Archie's task. Ordinary language allows us to explicitly relativize tasks to conditions: we can talk about whether he is *skilled at hitting the bullseye* (full-stop), or about whether he is skilled at *hitting the bullseye in such-and-such conditions*. Our account allows that Archie is skilled at *hitting the bullseye in gale-force winds*. After all, he succeeds at this task in a sufficiently high proportion of worlds where conditions are normal for this task (since these will all be worlds where gale-force winds are present). At the same time, our account explains why we would be reticent to claim that Archie is skilled at hitting the bullseye *full-stop*.<sup>14</sup>

Therefore, by adopting a standard strategy for handling mimics — namely, expanding the sphere of relevantly close worlds — we defuse the counterexample to the sufficiency of modal robustness for skillfulness. Can this strategy also defuse the Temp counterexample to the sufficiency of modal robustness for knowledge?

According to the present proposal, Temp's belief is safe (maximally modally robust) if it is true in all of the worlds where conditions are at least as normal for the task at hand as those that obtain at Temp's world.<sup>15</sup> Which worlds are these? Temp's task is *forming beliefs about the temperature on the basis of a thermometer*. Now, consider the worlds where Temp consults a broken thermometer but there is no hidden helper manipulating the thermostat. Intuitively, at these worlds conditions are *more* normal for the task at hand than they are at Temp's world. (After all, having a hidden helper controlling the thermostat is a highly abnormal circumstance.) But at these helper-free worlds, consulting a broken thermometer will typically lead to false beliefs. And so by expanding the relevantly close worlds to include worlds where conditions are at least as normal as those

---

<sup>13</sup> This seems plausible insofar as one wants a notion of relevant closeness that validates a standard centering requirement. We discuss conceptions of skillfulness that relax this requirement in §5.2.

<sup>14</sup> This reticence is reflected in our ordinary scoring practices. For example, the results of certain athletic competitions are not treated as records if they are obtained under various abnormal conditions, such as the presence of an exceptionally strong tailwind.

<sup>15</sup> The view of knowledge that emerges bears some important similarities to the account developed in Goodman and Salow (2018). Their account is motivated on largely independent grounds. However, it also cashes out safety in terms of what obtains in the worlds that are *at least as normal as* the world of evaluation.

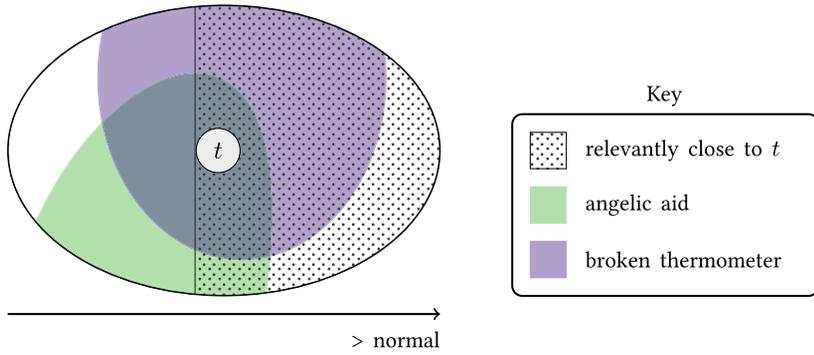


Fig. 1. Temp. (Temp’s belief is unsafe because it is false in many of the relevantly close worlds where there is a broken thermometer but no angelic aid).

which obtain at Temp’s world, we deliver the result that Temp’s belief does not qualify as safe, as desired. (See Fig. 1.)

Thus by making the sphere of relevantly close worlds sensitive to normal conditions, we overcome the challenges posed by Marie and Temp. At the same time, we retain the main advantages of a traditional safety account — in particular, its handling of lotteries and fake barns. Start with lotteries: if Lottie’s ticket had won, conditions would have been as normal for the task at hand as they are at Lottie’s world. Hence the current approach correctly predicts that she does not know. On to fake barns: since Henry is in fake barn county, he is in conditions that are relatively abnormal for the task at hand (*forming perceptual beliefs about barns*). Arguably, however, these conditions would not have become *more* abnormal if, while still in fake barn county, he simply happened to glance a facade rather than a real barn. And so at least some facade-gazing worlds count as relevantly close. Hence the current approach correctly predicts that Henry’s belief does not amount to knowledge.<sup>16</sup>

## 5.2 Unity

**MVE** fares better than a simple modal account of knowledge when it comes to unifying the practical and the epistemic domains. After all, it says that what unites knowledge and skillful action is their shared modal robustness. However, one might worry that this connection is not enough. As we saw in §2, **AVE** captures the normative differences between Norm, who skillfully hits the target in normal conditions, Wendy, who misses the target because of a gust of wind, and Don, who hits the target by a lucky fluke in the double gust case. Moreover, **AVE** provides a story about what these cases have in common with their epistemic analogues (knowledge, justified false belief, and Gettiered belief). It does so because it has three normative statuses — accuracy, adroitness, and aptness — at its disposal. On the other hand, it is doubtful whether **MVE** can account

<sup>16</sup> These comparative judgments about the normality of *looking at a barn in fake barn county* vs. *looking at a facade in fake barn county* may not be universally shared. However, the original “No Knowledge” judgment in fake barn cases is not universally shared. An advantage of the current approach is that it explains the divergences in people’s intuitions about fake barns. This variation is attributed to divergences in how people rank the various worlds for normality.

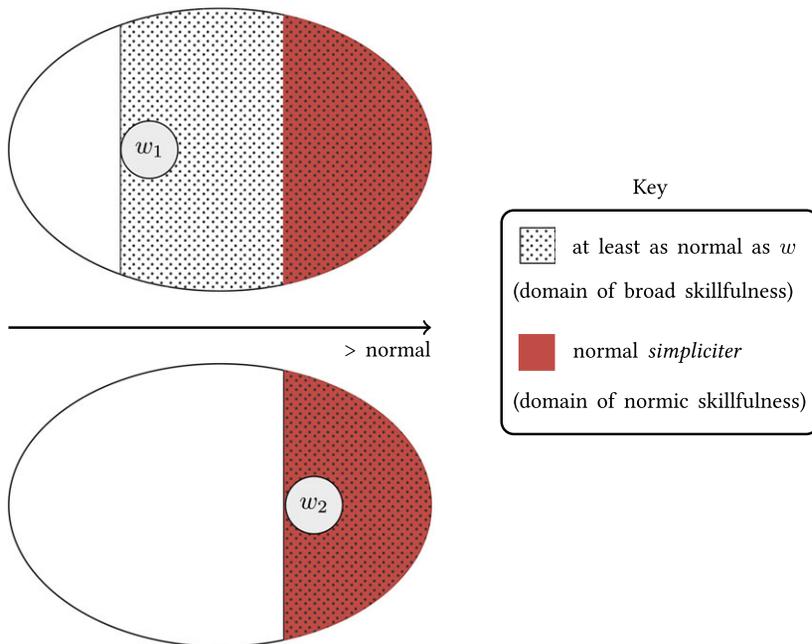


Fig. 2. Two Types of Skillfulness. (Broad skillfulness and normic skillfulness coincide when conditions are normal *simpliciter* ( $w_2$ ), but not when they are abnormal ( $w_1$ )).

for these differences. After all, **MVE** appears to have only *two* normative statuses: accuracy and (degrees of) modal robustness.<sup>17</sup>

Adequately responding to this challenge requires complicating our initial account: we need a more nuanced taxonomy of normative statuses than **MVE** currently provides. In §5.1, we stressed the intuitive connection between skillfulness and normal conditions. There are, however, at least two ways that we could use normal conditions to carve out a domain of quantification. First, we could take the route developed in §5.1, quantifying over worlds where conditions for the relevant task are *at least as* normal as those obtaining at the world of evaluation. Say that a performance is “broadly skillful” just in case it succeeds at a sufficiently high proportion of these worlds. Alternatively, we can zoom into the worlds where conditions are normal *simpliciter* for the relevant task. Adapting some terminology from Smith (2010), say that a performance is “normically skillful” just in case it succeeds at a sufficiently high proportion of these worlds. When the conditions are normal for the relevant task at a world  $w$ , a performance will be broadly skillful at  $w$  just in case it is normically skillful at  $w$ . But when conditions are abnormal at  $w$ , these statuses will come apart. (See Fig. 2.)

We now have three normative statuses, where before we had two. This enriched taxonomy captures the similarities and differences between our three archers. On the most natural way of filling out the cases, Norm, Wendy, and Don have equally high success rates across those worlds where conditions are normal for archery (worlds where there are no unforeseen gusts of wind or angelic interventions). Hence, their shots are equally normically skillful. However, they differ in whether they are broadly skillful. Norm’s

<sup>17</sup> Thanks to a referee for raising this point.

shot clearly is: since he is in normal conditions for archery, for him normic skillfulness and broad skillfulness coincide. By contrast, Don's shot is not broadly skillful: consider the worlds where the first gust blew his arrow off-course but no second gust intervened, or it intervened in a slightly different manner, failing to put his arrow back on course. At these worlds, conditions are at least as normal as those obtaining at Don's world. And at all these worlds his shot fails. What about Wendy's shot? Well, it fails at her world, and at many other worlds where a similar gust intervened. And conditions at these worlds are just as normal as those that obtain at her world. So her shot is not broadly skillful — or, at any rate, it is less broadly skillful than Norm's.

Thus by appealing to the difference between broad skillfulness and normic skillfulness, we can account for our three archers' shots. But which of these two notions corresponds to our ordinary notion of skillfulness? To answer this, consider again Marie. When she she tries to sink the basket, her shot does not have a high degree of either broad or normic skillfulness, since it fails in all angel-free worlds. Still, it is *somewhat* broadly skillful. After all, it succeeds in all of the angelic worlds, and in many of these worlds conditions are at least as normal as they are at hers. However, some may have the intuition that Marie's shot is not *at all* skillful. Arguably, this intuition is tracking normic skillfulness: she succeeds in *none* of the worlds where conditions are normal *simpliciter* for shooting hoops.

That said, intuitions about Marie's precise level of skillfulness may be rather murky. For this reason, we do not want to take a stand on which of these two notions corresponds to everyday "skillfulness" talk. Indeed, it may be that such talk is context-sensitive, sometimes picking out the one property and sometimes picking out the other. What matters is that both properties are important normative statuses, and that our normative appraisals of various performances are sensitive to both, as revealed by our responses to the three archers.<sup>18</sup>

While the distinction between broad skillfulness and normic skillfulness bears some resemblance to the distinction between adroitness and aptness, there are important differences. For one thing, we have carved out our distinction in purely modal terms, without recourse to the notion of success *in virtue of* the exercise of a skill. Moreover, broad skillfulness is not equivalent to aptness. Unlike aptness, broad skillfulness does not entail accuracy: an unsuccessful performance might be broadly skillful, so long as it succeeds at sufficiently many relevantly close worlds. And although *maximal* broad skillfulness does entail accuracy, it still comes apart from aptness, as can be seen from lotteries and fake barns: these beliefs are apt but unsafe, hence not maximally broadly skillful.

The distinction between broad and normic skillfulness also allows us to capture the difference between knowledge and justification. Translating our account of knowledge into present terms: a belief amounts to knowledge if and only if it is maximally broadly skillful — i.e., it is maximally modally robust across worlds where conditions are at least as normal for believing in that way. We could then hold that a belief is justified if and

---

<sup>18</sup> Here the situation is analogous to the traditional virtue epistemological distinction between adroitness and aptness. Both are quasi-technical terms, and it is not straightforward which of these is supposed to correspond to our everyday "skillfulness" talk. Despite this, virtue epistemologists claim that our normative appraisals are sensitive to both notions.

Table 1:  
Comparison between AVE and MVE

		AVE			MVE		
epistemic	practical	accurate	adroit	apt	accurate	normically skillful	broadly skillful
Justified false belief	Wayward gust	✗	✓	✗	✗	✓	✗
Gettiered belief	Double gust	✓	✓	✗	✓	✓	✗
Fake barns / Lotteries	Fielding	✓	✓	✓	✓	✓	✗
Knowledge	Norm	✓	✓	✓	✓	✓	✓

only if it is normically skillful — i.e., it is modally robust across worlds where conditions are normal *simpliciter* for believing in that way.<sup>19</sup>

This approach has a number of virtues. It explains why knowledge is factive but justification is not. After all, a world  $w$  will always be among the worlds where conditions are *at least as normal as* those which obtain at  $w$ , but at  $w$  conditions might not be normal *simpliciter*. It also explains why knowledge entails justification but not *vice versa*. After all, the set of worlds where conditions are normal *simpliciter* is a subset of the set of worlds where conditions are at least as normal as those which obtain at  $w$ . Finally, it allows us to capture the intuition that Temp and Henry are both justified in their beliefs. After all, if conditions were normal *simpliciter* for Temp’s task, there would be no hidden helper, but the thermometer would also be functioning. And if conditions were normal *simpliciter* for Henry’s task, then he would be in an environment free of fake barns.<sup>20</sup> (See Table 1 for a point-by-point comparison with AVE.)

### 5.3 Is Modal Robustness Necessary for Knowledge and Skillfulness?

A final challenge to MVE concerns whether modal robustness is necessary for knowledge or skillfulness. We saw in §2 that a number of authors have proposed counterexamples to the necessity of safety for knowledge. Similar counterexamples arise for the idea that modal robustness is necessary for skillfulness. For example, Pritchard (2009) describes the case of Fielding, who randomly selects one of a hundred targets and sends the arrow straight into the bullseye. Unbeknownst to him, all of the other targets were protected by a forcefield. His performance is skillful, though it would appear not to be modally robust: in the vast majority of relevantly close worlds where he shoots at a target, his shot is deflected by the forcefield.

These are important challenges. Nonetheless, we do think there are some promising defenses available — defenses that fall out of the moves made in the preceding sections.

Start with the counterexamples to the necessity of safety for knowledge. Recall Kelp’s example of Russell, who comes to know that it is 3:00 on the basis of a properly

<sup>19</sup> This way of understanding justification dovetails nicely with the account of justification in Smith (2010, 2016), which independently stresses the connection between justification and normality.

<sup>20</sup> In order for a belief to be justified, does it need to be *maximally* normically skillful? An affirmative answer will be attractive to those who think that Lottie cannot be justified in believing her ticket is a loser (cf. Smith 2010, 2016). But this answer is not mandatory; one could also hold that a belief is justified provided that it attains a sufficiently high threshold of normic skillfulness.

functioning clock. If Russell had looked at the clock a minute later he would have formed a false belief, since a demon would have manipulated the hands of the clock. But note that if the demon had interfered, this interference would have rendered conditions significantly more abnormal for the task at hand (*forming beliefs about the time based on a clock*) than they are at Russell's world, where the demon is present but does not intervene. So our view predicts that Russell's belief is safe (maximally broadly skillful), since it is true in all worlds where conditions are *at least as normal* for the task at hand.

This strategy also helps with other counterexamples. Consider Neta and Rohrbaugh's scenario (2004: 402): Beth is participating in an experiment where she is asked to recall the number of flashes displayed on a screen. She belongs to the control group. But she could have easily been assigned to another group, in which case she would have been given a drug that interferes with one's ability to count visual stimuli. Despite this, it seems that Beth is able to know the number of flashes when they are displayed. On our view, if Beth had been assigned to other group, the drug would have impaired her visual abilities, making conditions considerably less normal for the task at hand (*counting visual stimuli*). Thus Beth's belief is safe, since it is true in all worlds where conditions are at least as normal as those that obtain at her world.<sup>21,22</sup>

What about Pritchard's counterexample to the necessity of modal robustness for skillfulness (Fielding)? As we saw in §5.1, it is important to clarify which task is being evaluated for skillfulness. If the task is *shooting non-forcefield-protected targets*, then our account predicts that Fielding's shot is both broadly and normically skillful. If the task is *shooting targets in general*, then arguably his shot is not broadly skillful, since many of the worlds where conditions are at as least as normal will also contain forcefields. And at most of these worlds, Fielding's shot will be deflected. However, his shot is normically skillful, because it succeeds in worlds where conditions are normal *simpliciter*, since these worlds are forcefield-free. According to this diagnosis, Fielding's situation is a practical Gettier case. And this seems like the right result, since the case is structurally analogous to the fake barns case.

To be clear: our claim is not that a normality-based framework, on its own, handles *all* of the apparent counterexamples to the necessity of modal robustness for skillfulness and knowledge. Rather, our claim is that judgments about skillfulness are closely connected to judgments about the normal conditions for performing a task — a connection that yields new avenues for warding off at least *some* of the standard counterexamples to

---

<sup>21</sup> Neta and Rohrbaugh (2004) also give a related example: Hugo grabs a water bottle from a refrigerator, without knowing that most of the water bottles within have been polluted with a tasteless, colorless toxin. By luck, Hugo seizes the only unpolluted bottle. As he drinks it, he correctly believes he is drinking unadulterated water. Intuitively, this belief does not qualify as knowledge. Neta and Rohrbaugh suggest that any safety condition that delivers the correct verdict about Beth will deliver the wrong result about Hugo. However, our proposal has the resources to distinguish between these two cases. Both Beth and Hugo are in rather abnormal circumstances. The key difference, however, is that if Beth had been given a drug that impaired her faculties, this impairment would have made conditions even more abnormal for the relevant task. By contrast, if Hugo had simply selected one of the other bottles, this this would not have made conditions any more abnormal, since no external agent or substance would have tampered with his belief-forming process. (In this regard, Hugo resembles Henry in the fake barn case, whereas Beth resembles Russell in Kelp's case.) More generally, intuitions about these cases — together with intuitions about Temp and Marie — suggest that our judgments about normal conditions are sensitive to whether external factors (be they meddlesome or helpful) interfere with the performance of the task.

<sup>22</sup> This strategy also extends in a similar fashion to handle Bogardus' (2014) "atomic clock" counterexample.

modal accounts.<sup>23</sup> Moreover, the appeal to normality is not *ad hoc*, because it falls out of a dispositional conception of skillfulness, together with the idea that knowledge is a type of skillful performance. In this regard, the relevance of normality considerations is a consequence of the virtue epistemological dimension of the theory. This is another respect in which the modal approach benefits from an integration with virtue epistemology.

## 6. Final Remarks and Comparisons

In this essay, we have developed and defended a modal version of virtue epistemology. The view agrees with traditional virtue epistemology that knowledge is a type of skillful performance. But whereas traditional virtue epistemology assigns a central role to *aptness*, our view dispenses with this notion, explaining knowledge and skillfulness in purely modal terms.

Some may question whether aptness is so easily jettisoned. We argued (§5) that some of the standard motivations for aptness can be accommodated by reconstruing the relevant closeness relation in terms of normal conditions. However, some may worry that once relevant closeness is reconstrued in this manner, safety (and modal robustness more generally) will end up entailing aptness. After all, the standard counterexamples to the safety-aptness entailment come from cases like Temp. But, on our view, Temp's belief is not safe after all. If safety does entail aptness, then **MVE** and **AVE** are closer than they initially appeared.

However, even if Temp is not a counterexample to the safety-aptness entailment, there are principled reasons for denying this entailment holds. For a performance to be apt it must succeed *in virtue of* the manifestation of a skill. But as the literature on grounding teaches us, the *in virtue of* relation resists being understood in purely modal terms. Take the example of Socrates and his singleton set (Fine 1995: 271). They exist in all the same worlds but the singleton set exists in virtue of Socrates, as opposed to the other way around. *Mutatis mutandis*, the fact that maximal modal robustness entails success does not mean that the latter happens in virtue of the former.

A second way of trying to bridge the gap between **MVE** and **AVE** would be to identify knowledge with *maximally apt* belief, where a belief is maximally apt if it succeeds in virtue of being maximally adroit. Proponents of **AVE** could then combine this with the view that maximal adroitness amounts to maximal modal robustness. This would entail a safety condition on knowledge, since on the resulting view a belief would amount to knowledge provided *it is true in virtue of being safe*.

---

<sup>23</sup> In particular, Comesaña's (2005) "Halloween Party" case, Baumann's (2008) bank robber case, Sosa's (2007) rubbish chute example, as well as some of J. Greco's (2007) cases are not obviously susceptible to the normality-based treatment offered here. While a full discussion of these cases is outside the scope of this paper, we should note that modal accounts have other resources at their disposal. In particular, by getting clear on which circumstances of a world need to be held fixed when determining whether conditions at some other world count as at least as normal, defenders of **Safety** can try to accommodate some of these recalcitrant cases. See Bogardus (2014) and Pritchard (2016) for development of related ideas. A separate, but related, worry about the necessity direction of **Safety** is whether it allows for inductive knowledge (see e.g., J. Greco 2007). Here too, an adequate discussion would require a separate paper. However, it is worth noting that even if a set of premises does not entail some conclusion, it could still be that in all the relevantly close worlds where the premises are true, the conclusion is also true. And so there is nothing in the nature of **Safety** that precludes inductive knowledge.

However, there are two difficulties with this suggestion. First, it is unclear what it means for a belief to be true in virtue of being safe. For a belief to be safe is for it to be true at the actual world *and* at all the other relevantly close worlds. In this sense, safety is a conjunctive condition. But a conjunction is true in virtue of its conjuncts, not the other way around (e.g., Fine 2012: 50). Second, it is unclear what work the appeal to the *in virtue of* relation is doing in this account. If all of the important explanatory work — all of the work in handling Gettier cases, lotteries, Temp’s case, etc. — is done by safety, why add the part about being true *in virtue of* being safe? The appeal to aptness will be an idle wheel.<sup>24</sup>

In conclusion, modal and virtue epistemological approaches have long been viewed as competitors. And even hybrid approaches take the modal condition on knowledge and the virtue epistemological condition to be distinct. On our approach, however, the modal condition falls out of the virtue epistemological condition, since skillfulness itself is understood in modal terms. The result is a new form of virtue epistemology, which combines the benefits of both approaches while overcoming many of their difficulties.<sup>25</sup>

## References

- Aimar, S. 2018. “Disposition Ascriptions,” *Philosophical Studies*, 1–26.
- Baumann, P. 2008. “Is Knowledge Safe?” *American Philosophical Quarterly* 45 (1): 19–30.
- Beddor, B. ms. “Certainty First.”
- Bogardus, T. 2014. “Knowledge under Threat,” *Philosophy and Phenomenological Research* 88 (2): 289–313.
- Colaço, D. et al. 2014. “Epistemic Intuitions in Fake-Barn Thought Experiments,” *Episteme* 11 (2): 199–212.
- Comesaña, J. 2005. “Unsafe Knowledge,” *Synthese* 146 (3): 395–404.
- Fine, K. 1995. “Ontological Dependence,” *Proceedings of the Aristotelian Society* 95 (1): 269–290.
- 2012. “Guide to Ground,” In F. Correia and B. Schneider (Eds.), *Metaphysical Grounding: Understanding the Structure of Reality*: 37–80. Cambridge University Press.
- Gettier, E. L. 1963. “Is Justified True Belief Knowledge?,” *Analysis*, 23 (6), 121–123.
- Goldman, A. 1976. “Discrimination and Perceptual Knowledge,” *Journal of Philosophy* 73: 771–791.
- Goodman, J. & Salow, B. 2018. “Taking a chance on KK,” *Philosophical Studies*, 175 (1), 183–196.

---

<sup>24</sup> This is not to deny that there *is* a property of aptness. Our position here is analogous to a metaphysician who concedes that the notion of grounding is coherent, but denies that this notion is needed to explain important metaphysical notions such as physicalism, constitution, and the like. Similarly, we concede that the notion of aptness is coherent and even that performances can be assessed in terms of whether or not they are apt. But we deny that it plays any role in explaining knowledge or skillfulness.

<sup>25</sup> Thanks to Andy Egan, Simon Goldstein, Chris Kelp, Bernhard Salow, Mona Simion, and two anonymous referees at *Philosophy and Phenomenological Research* for helpful comments. For useful discussion, we are also grateful to participants in Ernie Sosa’s dissertation seminar at Rutgers, as well as to audiences at the University of Edinburgh, the 2017 Bled Epistemology Conference, the National University of Singapore, and 3rd Cologne-Leuven Epistemology Meeting.

- Greco, D. 2016. "Safety, Iteration, and Explanation," *Philosophical Issues* 26 (1): 187–208.
- Greco, J. 2007. "Worries about Pritchard's Safety," *Synthese* 158 (3): 299–302.
- 2010. *Achieving Knowledge: A Virtue-theoretic Account of Epistemic Normativity*. Cambridge University Press.
- Hawthorne, J. 2000. "Implicit Belief and A Priori Knowledge," *The Southern Journal of Philosophy* 38 (1): 191–210.
- 2004. *Knowledge and Lotteries*. Oxford University Press.
- Kelp, C. 2009. "Knowledge and Safety," *Journal of Philosophical Research* 34: 21–31.
- 2013. "Knowledge: The Safe-Apt View," *Australasian Journal of Philosophy* 91: 265–78.
- 2017. "Knowledge First Virtue Epistemology," Carter et al. (eds), *Knowledge First: Approaches in Epistemology and Mind*, Oxford University Press.
- Krakauer, J. et al. 2000. "Learning of Visuomotor Transformations for Vectorial Planning of Reaching Trajectories," *The Journal of Neuroscience* 20 (23), 8916–8924.
- Lackey, J. 2007. "Why We Don't Deserve Credit for Everything We Know," *Synthese*, 158 (3), 345–361.
- Lasonen-Aarnio, M. 2010. "Unreasonable Knowledge," *Philosophical Perspectives* 24: 1–21.
- Manley, D. 2007. "Safety, Content, Apriority, Self-knowledge," *The Journal of Philosophy* 104 (8): 403–423.
- and Wasserman, R. 2008. "On Linking Dispositions and Conditionals," *Mind* 117 (465): 59–84.
- Nagel, J. 2013. "Knowledge as a Mental State," In T. Szabo Gendler & J. Hawthorne (Eds.), *Oxford Studies in Epistemology* 4: 273–308. Oxford University Press.
- Neta, R. and Rohrbach, G. 2004. "Luminosity and the Safety of Knowledge," *Pacific Philosophical Quarterly* 85 (4): 396–406.
- Pavese, C. 2013. *The Unity and Scope of Knowledge*, PhD Dissertation. Rutgers University.
- 2016a. "Skill in epistemology Part I: Skill and Knowledge," *Philosophy Compass* 11: 642–649.
- 2016b. "Skill in epistemology Part II: Skill and Know-How," *Philosophy Compass* 11: 650–660.
- 2018. "Know-How, Action, and Luck," *Synthese*.
- Pritchard, D. 2005. *Epistemic Luck*. Oxford University Press.
- 2009. "Apt Performance and Epistemic Value," *Philosophical Studies* 143: 407–416.
- 2012. "Anti-Luck Virtue Epistemology," *Journal of Philosophy* 109.3: 247–279.
- 2016. "Epistemic Risk," *Journal of Philosophy* 113 (11): 550–571.
- Smith, M. 2010. "What Else Justification Could Be," *Noûs*, 44 (1), 10–31.
- 2016. *Between Probability and Certainty: What Justifies Belief*. Oxford University Press.
- Sosa, E. 1999. "How to Defeat Opposition to Moore," *Philosophical Perspectives* 13: 137–149.
- 2007. *A Virtue Epistemology*. Clarendon Press.
- 2015. *Judgment and Agency*. Oxford University Press.

- Stanley, J. 2011. *Know How*. Oxford University Press.
- & Williamson, T. 2017. “Skill,” *Noûs* 51 (4): 713–726.
- Turri, J. 2011. “Manifest Failure: the Gettier Problem Solved,” *Philosophers’ Imprint* 11 (8): 1–11.
- & Friedman, O. 2014. “Winners and Losers in the Folk Epistemology of Lotteries,” In J. R. Beebe (Ed.), *Advances in Experimental Epistemology*: 45–69. London: Bloomsbury.
- Vetter, B. 2014. “Dispositions Without Conditionals,” *Mind* 123 (489): 129–156.
- Wedgwood, R. 2018. “The Internalist Virtue Theory of Knowledge,” *Synthese*.
- Williamson, T. 2000. *Knowledge and Its Limits*. Oxford University Press.
- 2009. “Probability and Danger.” *Amherst Lecture in Philosophy*.
- Zagzebski, L. 1996. *Virtues of the Mind*. Cambridge University Press.