

Purifying impure virtue epistemology

Fernando Broncano-Berrocal¹

Published online: 3 February 2017
© Springer Science+Business Media Dordrecht 2017

Abstract A notorious objection to robust virtue epistemology—the view that an agent knows a proposition if and only if her cognitive success is because of her intellectual virtues—is that it fails to eliminate knowledge-undermining luck. Modest virtue epistemologists agree with robust virtue epistemologists that if someone knows, then her cognitive success must be because of her intellectual virtues, but they think that more is needed for knowledge. More specifically, they introduce independently motivated modal anti-luck principles in their accounts to amend the problem of eliminating luck—this makes their views instances of impure virtue epistemology. The aim of the paper is to argue, firstly, that such a move lacks adequate motivation; secondly, that the resulting impure accounts equally fail to handle knowledge-undermining luck. On a more positive note, these results bolster a more orthodox virtue-theoretic approach to knowledge that assigns a fundamental explanatory role to the notion of ability. In this sense, the paper also sketches an account of ability and a corresponding account of knowledge that explains how success from ability (of the right kind) is incompatible with success from luck.

Keywords Virtue epistemology · Aptness · Safety · Luck · Epistemic luck

1 Introduction

The *core thesis* of virtue epistemology is that knowing that *p* is a matter of succeeding cognitively—i.e., coming to believe *p* truly—*because of* an exercise of intellectual virtue.

✉ Fernando Broncano-Berrocal
f.broncano.berrocal@hum.ku.dk;
<http://www.fbroncano-berrocal.com>

¹ Department of Media, Cognition and Communication, University of Copenhagen, KUA2, Building: 16-1-20 2300, Copenhagen, Denmark

There are three main disputes within virtue epistemology. The first one concerns the notion of intellectual virtue. While *virtue responsibilists* (e.g., Baehr 2011) think that intellectual virtues are character traits such as open-mindedness or conscientiousness, *virtue reliabilists* (e.g., Greco 2009; Sosa 2010) hold that intellectual virtues should be conceived as reliable cognitive abilities such as vision or reasoning. For the purposes of the paper, I will assume the virtue reliabilist conception of intellectual virtues.

The second dispute concerns the interpretation of the technical ‘because of’ relation in the core thesis. According to some virtue reliabilists, to succeed cognitively because of cognitive ability is for that success to *manifest* the exercise of cognitive ability—this is the *manifestation reading*. Advocates of this interpretation (e.g., Kelp 2013a; Sosa 2011, 2015; Turri 2011) model knowledge in dispositional terms: in the same way as salt manifests its solubility when stirred into water, a cognitive success constituting knowledge manifests the agent’s cognitive abilities if exercised in appropriate internal and external conditions. An alternative reading has it that to get things right because of cognitive ability the relevant abilities must figure as a causal factor in the explanation of such a cognitive success. This *explanatory reading* splits into two sub-readings. According to some authors (e.g., Greco 2010a), the agent’s cognitive abilities must be the *most salient* explanatory factor of why she comes to believe something true—this is the *strong explanatory reading*—whereas others (e.g., Pritchard 2012a; Sosa 2007) hold that they must be a salient but *not necessarily the most salient* factor—this is the *weak explanatory reading*. These different interpretations are important insofar as they help put forward different views in the virtue epistemological landscape.

The third dispute within virtue epistemology, and the most relevant to the purposes of this paper, concerns the scope of its core thesis. The more orthodox stance is that of *robust virtue epistemologists*, who defend that knowing that *p* is just a matter of coming to believe *p* truly because of an exercise of cognitive ability. By contrast, *modest virtue epistemologists* argue that, although knowing *p* requires believing *p* truly because of cognitive ability, more is needed for knowledge. The reason why they retreat to such a modest position is that, in their opinion, no version of virtue epistemology is able to successfully eliminate (modally understood) knowledge-undermining luck. As we will see, they typically introduce independently motivated modal conditions on knowledge to amend the problem.

This means that different versions of virtue epistemology are motivated differently. For example, robust virtue epistemologists back up their principles with a unifying rationale that is in the spirit of the core thesis.¹ This makes their views instances of *pure virtue epistemology*.² By contrast, as just pointed out, modest virtue epistemologists typically tack on non-virtue-theoretic modal principles. Their views are considered versions of *impure virtue epistemology*, not

¹ See Sect. 3.

² The terminology is from Kelp (2013a). While the term ‘robust’ puts more emphasis on the way a virtue epistemological view defines the concept of knowledge—namely, solely using virtue conditions—the term ‘pure’ emphasizes how it is motivated. From now on, I will use the two terms indistinctly.

because they feature modal conditions *per se*—virtue conditions can be modal—,³ but because the motivation they give for them is not only different from the motivation they offer for their (modest) virtue conditions, but also intendedly independent of the core thesis.

At first blush, its simplicity and elegance makes robust (i.e., pure) virtue epistemology an appealing approach to knowledge. After all, it does all the explanatory work using a single theoretical resource: the notion of cognitive ability. But if it does not yield an adequate anti-luck epistemology—as impure virtue epistemologists insistently argue—it might be better to retreat to a modest position and resort to independently motivated non-virtue-theoretic principles which can handle the problematic cases of epistemic luck. The *aim of this paper* is to argue that such a move not only lacks adequate motivation, but also that the resulting ill-motivated impure views equally fail to eliminate knowledge-undermining luck.

The *plan* is as follows. Section 2 explains why some virtue epistemologists have retreated to modest positions and adopted impure views. Section 3 argues that, to avoid the charge of ad hocery—impure views give the impression to be designed just to deal with troublesome cases—they need to offer a solid theoretical motivation. Section 4 argues that the motivation that impure virtue epistemologists actually provide fails. Section 5 continues the attack by arguing that the two main impure views in the literature also fail to eliminate knowledge-undermining luck. In addition, it argues that the prospects for views that feature modal anti-luck principles—and in particular the safety principle—are dim. Section 6 outlines an alternative way to develop virtue epistemology in terms of the notion of cognitive ability.


The *results* of the paper might be used in two ways. Virtue epistemology skeptics might use them in conjunction with criticism of robust virtue epistemology to cast doubt on the whole project: robust or modest, pure or impure, virtue epistemology does not seem to yield an anti-luck epistemology. More optimistically, robust virtue epistemologists—who firmly believe that the notion of cognitive ability has proven to be too fruitful when it comes to long-standing questions such as the normativity, nature and value of knowledge to make that path worth taking—can use the arguments presented here to undermine the position of their modest rivals and seduce them to the pure side of virtue epistemology.

2 From pure to impure virtue epistemology

To see what has motivated the transition to an impure approach, let's begin by considering the two main robust views in the literature as well as the main objections they face. As we have seen, the core idea of pure or robust virtue epistemology is that knowing *p* is just a matter of coming to believe *p* truly because of an exercise of cognitive ability. If one shares this idea and thinks that the correct reading of the technical 'because of' is the manifestation reading (e.g., Sosa

³ See Sect. 5.3.

2011, 2015), one will endorse the view, called *dispositional robust virtue epistemology*, that *S* knows that *p* if and only if *S*'s cognitive success (concerning *p*) manifests her cognitive abilities.⁴ If one opts for the strong explanatory reading instead (e.g., Greco 2010a), one will be committed to *explanatory robust virtue epistemology*, the view that *S* knows that *p* just in case the exercise of *S*'s cognitive abilities is the most salient or important factor in the set of causal factors that explain why she comes to believe *p* truly. Schematically:⁵



Lackey (2007, 2009) and Pritchard (2012a) have argued that if one wants to analyze the concept of knowledge solely in terms of the notion of cognitive ability—as both pure accounts do—one will have to strive hard to solve the following two problems: (1) explaining knowledge that partly stems from epistemically relevant factors beyond one's cognitive agency and (2) eliminating knowledge-undermining luck.

The real problem, Lackey and Pritchard argue, is that (1) and (2) cannot be solved independently. What is worse, they pull pure views in opposite directions. Indeed, the objection takes the form of a dilemma called the *creditworthiness or attributability dilemma*: in order to solve (1), pure virtue epistemologists need to weaken their accounts, but then they cannot solve (2); in order to solve (2), they need to strengthen their accounts, but then they cannot solve (1).⁶

Lackey (2007, 2009) illustrates (1) with ordinary cases of testimonial knowledge—e.g., the kind of knowledge that one acquires when asking an adult passerby for directions in a train station. Pritchard (2010) suggests that the same point can be also made with cases of knowledge from instruments. In a nutshell, the problem is the following: if knowledge is *just* a matter of exercising individual cognitive abilities, then it is puzzling how it can be acquired from a reliable speaker or an instrument.⁷

⁴ See Kallestrup and Pritchard (2016) for the terminology.

⁵ 'VE' stands for 'virtue epistemology'.

⁶ An important difference in the way Lackey and Pritchard respectively formulate the dilemma is that, while Lackey thinks that problem (1) shows that the agent does not *deserve credit* for her cognitive success, Pritchard thinks that (1) shows that the agent's cognitive success is not completely *attributable* to (in the sense of not being explained by) her cognitive abilities.

⁷ More specifically, in the case of testimonial knowledge, the speaker's memory when recalling the relevant facts and her linguistic competences when passing the relevant accurate information *partly explain* the hearer's cognitive success, which proves explanatory robust virtue epistemology too strong. In addition, although the hearer's cognitive success may *manifest* her cognitive abilities to a rather low degree, the fact that it is to a large degree due to factors beyond her cognitive agency seems to prove

The cases that best help illustrate problem (2) are cases with the structure of the famous fake barn case (Goldman 1976), to wit, cases in which someone reliably forms a perceptual true belief about the presence of an object (e.g., a barn), but in which she could easily have formed a false belief by perceiving a fake of that object (e.g., a fake barn). In such cases, subjects lack knowledge because they get things right by sheer luck. The problem is that, although the unfortunate subjects exercise their perceptual faculties in seemingly the same manner that typically yields knowledge, the kind of dangerous epistemic luck in play does not affect the way in which they actually form their beliefs—Pritchard calls it *environmental luck*, in that the environment itself is modally inhospitable for forming knowledgeable beliefs.⁸ The point then is that, since environmental luck is incompatible with knowledge, the theoretical resources of pure views are not enough to constitute an adequate anti-luck epistemology.⁹

Enter impure virtue epistemology. Epistemologists like Kelp (2013a) or Pritchard (2012a, b) are convinced that virtue conditions are doomed to fail when it comes to environmental luck. But they also believe that the idea that knowing has to do with exercising cognitive ability—the *ability intuition*, as Pritchard calls it—is too widespread and fruitful to abandon the entire virtue epistemology project. If one agrees with them and thinks that the correct reading of the technical ‘because of’ is the manifestation reading, one will be inclined to endorse the weak view, called *dispositional modest virtue epistemology*, that if *S* knows that *p*, then *S*’s cognitive success (concerning *p*) manifests her cognitive abilities. Those who prefer the weak explanatory reading will opt for *explanatory modest virtue epistemology* instead, that is, the view that if *S* knows that *p*, the exercise of *S*’s cognitive abilities is a salient but not necessarily the most salient factor in the set of causal factors that explain her cognitive success.

Kelp and Pritchard believe that modest views are able to escape the attributability dilemma. According to Pritchard (2010, 2012a), problem (1) is not troublesome for explanatory modest virtue epistemology because, compatibly with the view, the hearer’s cognitive success amounts to knowledge only when it is to some degree explained by her informant-selection abilities—which is different from saying that it needs to be fully explained by them.¹⁰ For Kelp (2013a), it is dispositional modest virtue epistemology that is weak enough to avoid (1). His point is that competent hearers manifest their informant-selection abilities as well as their linguistic

Footnote 7 continued

dispositional robust virtue epistemology too strong too. See Sect. 6 and Broncano-Berrocal (2016) for a reply to the latter point.

⁸ Another important difference in the way Lackey and Pritchard formulate the dilemma is that, while Lackey illustrates (2) with Gettier-style cases in general, Pritchard only appeals to cases of environmental luck—according to Pritchard, standard Gettier-style cases are not cases of environmental luck.

⁹ More specifically, in cases of environmental luck such as the fake barn case the agent’s perceptual success is (a) fully explained by the exercise of her perceptual faculties and (b) manifests them—i.e., the right-hand sides of the two biconditionals that respectively define explanatory and dispositional robust virtue epistemology hold.

¹⁰ See Pritchard (2010) for a similar point concerning knowledge from instruments.

capacities if they are in appropriate internal and external conditions—e.g., while awake, sober, sufficiently attentive and before informants who are not attempting to deceive them. As a matter of fact, this is what happens in cases of testimonial knowledge, or so Kelp argues.¹¹

In order to solve (2), the problem of eliminating environmental luck, Pritchard and Kelp append the *safety principle* to their respective modest virtue epistemological views, i.e., the principle that if *S* knows that *p*, then in most close possible worlds in which *S* forms a belief in *p* in the same way as in the actual world *p* is true.^{12,13}

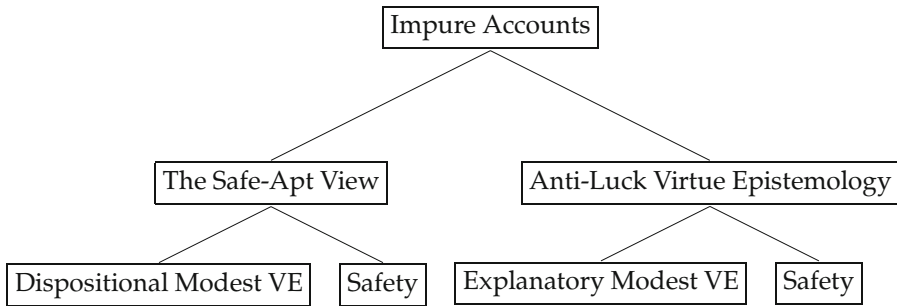
The resulting impure views are the following. Kelp's (2013a), which he dubs the *safe-apt view*, says that *S* knows that *p* if and only if *S*'s cognitive success is apt—in the sense that it manifests *S*'s cognitive abilities—and safe. Pritchard (2012a, b), by contrast, holds that *S* knows that *p* just in case the exercise of *S*'s cognitive abilities partly explains her safe cognitive success—in the sense that *S*'s cognitive abilities need to be a salient but not necessarily the most salient factor in set of causal factors that explain it. Pritchard calls this view *anti-luck virtue epistemology*. There are two ways to interpret Pritchard's account. The weaker one is that knowledge arises whenever the agent's cognitive success is partly explained by her cognitive abilities and safe. A stronger reading (the one that Pritchard plausibly has in mind) is that known beliefs are beliefs whose safety is partly explained by the exercise of cognitive ability—i.e., the exercise of cognitive ability must partly explain why those beliefs end up 'hitting' the truth not only in the actual but also in close

¹¹ Kelp (2014) gives a similar argument for knowledge from instruments.

¹² See Pritchard (2007), Sosa (1999) and Williamson (2000) for different formulations of safety. The arguments presented here apply to any formulation of safety that is introduced independently of knowledge.

¹³ It is in principle possible to defend impure versions of virtue epistemology with anti-luck conditions other than safety. The obvious and perhaps the only candidate is the *sensitivity principle* (due to Dretske 1971; Nozick 1981), which roughly says that if *S* knows that *p*, then in the closest possible worlds in which *p* is not true, *S* does not believe that *p*. However, for several reasons, sensitivity is no alternative. To begin with, unlike safety, it does not preserve knowledge closure—roughly the principle that if *S* knows that *p*, and *S* knows that *p* entails *q*, then *S* knows that *q*. For example, it follows from sensitivity that, although (*p*) having hands entails that (*q*) you are not a brain in a vat, the latter is something that is not closed under known entailment because in the closest possible worlds in which (not-*q*) you are a brain in a vat, you would still believe that (*q*) you are not one—i.e., sensitivity does not hold. Although certainly unwelcome, the violation of closure might not be such a terrible consequence if the chosen modal principle can still eliminate luck. The real reason why virtue epistemologists think that sensitivity is not a good partner for their accounts is that there is a general agreement that cases of inductive knowledge prove that sensitivity is not necessary for knowledge (see Pritchard 2012a). Suppose that you are allergic to peanuts and know that every time you have eaten peanuts you have suffered an anaphylactic shock. Suppose further that you know on this basis that the peanut you have just accidentally swallowed will trigger an anaphylactic shock. If this were not true—e.g., if the peanut were covered with a plastic layer that would prevent the usual allergic reaction—then you would surely continue to believe that it is true. So, despite you have inductive knowledge, sensitivity does not hold. When one is in the business of analyzing knowledge in reductive terms, the last thing one wants is to include a condition that knowledge does not require. For these reasons, here I will be only concerned with the most plausible versions of impure virtue epistemology, that is, Kelp's and Pritchard's respective safety-based accounts. In Sect. 5.3, we will see that several cases also speak against safety being necessary for knowledge. This makes the prospects for safety-based virtue epistemology dim, or so I will argue.

possible worlds. In this sense, the strong reading entails the weak one. The following schema represents the two impure accounts:¹⁴



One might wonder why Kelp’s safe-apt view and Pritchard’s anti-luck virtue epistemology count as versions of impure virtue epistemology. This is the reason: their respective motivations for tacking on safety are different from the motivations they offer for their modest virtue conditions, to the point that they are intendedly introduced independently of the ability intuition as well as of the main tenet of virtue epistemology—namely, the idea that knowing is a matter of succeeding cognitively because of an exercise of cognitive ability. For Pritchard and Kelp, then, the concept of knowledge has a two-part structure: on the one hand, knowledge is due to a proper exercise of one’s cognitive agency, but competent beliefs can be lucky; on the other, knowledge arises in the absence of dangerous epistemic luck, but non-lucky beliefs might fail to be competent. Therefore, the incompatibility between luck and knowledge does not exclusively stem from the fact that known beliefs are competently formed. Independent anti-luck principles are needed.

3 The need for theoretical motivation

If modal principles like safety help handling troublesome cases of luck, why do impure virtue epistemologists need to justify their inclusion? To put it briefly, if their presence in a virtue epistemological account is not justified, the account faces a *charge of ad hocery*.¹⁵ To begin with, Pritchard (2012a: 274) admits that, *qua* epistemologists, we cannot be home and dry unless some explanation is given for why knowledge has a two-part structure in the first place. Kelp (2013a) also acknowledges that the consequence of retreating to an impure position is that the resulting view will be “more complex and less elegant than its pure cousins” (Kelp 2013a: 12), and admits that coping with the cases (e.g., eliminating environmental luck) cannot be the only motivation for introducing an independent modal condition, on pain of ad hocery.

¹⁴ As before, ‘VE’ stands for ‘virtue epistemology’.

¹⁵ Greco (2010b: 227) also expresses this concern.

To see this objection in more detail, let's start by considering the following possible reply.¹⁶ Suppose that a plausible counterexample speaks against one's preferred account of knowledge. Why is the introduction of new conditions on knowledge ad hoc? After all, introducing new conditions to deal with counter-evidence is just the typical way in which theories of knowledge have been developed. Imagine, for the sake of the argument, that at the very beginning of epistemology the only available view was that knowledge is belief. In the light of new evidence (namely, cases of false belief without knowledge), early epistemologists might have been pulled to introduce a condition on truth. But there's nothing ad hoc with this move, just theoretical development. In addition, the generic justification condition might have been introduced likewise, without any trace of ad hocery yet. Moreover, the reply continues, something similar might have happened with virtue epistemology. At some point in the development of the discipline, virtue epistemologists turned out to be convinced (with good reason) that ability conditions are needed, but only after the discovery that competent beliefs can fall short of knowledge due to harmful environmental luck, they realized that extra (modal) anti-luck conditions are also needed. Again, this is not ad hoc, just theoretical development.

This strand of argument might sound appealing at first blush, but it is ultimately unsatisfactory. As it is well-known, the development of contemporary epistemology has been anything but a smooth process. The post-Gettier industry, as Williamson (2013: 1–2) puts it, not only has produced increasingly gerrymandered definitions, but they have also succumbed, one after another, to counterexamples. Indeed, current dissatisfaction with the analytical project stems from the fact that many conditions on knowledge have been introduced simply with the aim of handling counterexamples.¹⁷ Accordingly, what the previous reply obviates is the fact that the more controversial a condition or a thesis is (in epistemology or elsewhere), the more motivation is needed to append it to one's account (of knowledge or of some other concept). More specifically, there is a significant difference between introducing the initial and uncontroversial belief and truth conditions on knowledge (and even the generic and unexplained justification condition of the JTB account) and the technical and controversial principles that have flooded the literature since Gettier's challenge to the classical view (Gettier 1963). Therefore, some rationale needs to be given for why the latter should guide our thinking about knowledge, a rationale that must be independent of whether such principles get the cases right, in other words, of whether there is good reason to think that the resulting analyses are extensionally adequate.¹⁸

¹⁶ Thanks to an anonymous reviewer for pressing this point.

¹⁷ See Shope (1983) for an exhaustive collection of counterexamples and gerrymandered definitions.

¹⁸ A proof that the JTB conditions are uncontroversial is that it generates more controversy to argue against their inclusion in an account of knowledge than assuming without argument that the concept of knowledge features them. For controversial arguments against the truth condition, see Tienson (1974). For well-known yet controversial arguments against the belief condition, see Williamson (2000). The list of epistemologists who implicitly or explicitly hold that knowledge does not entail (some form of) epistemic justification is quite extensive. For example, most modal analyses of knowledge (e.g., Nozick

If anything, pure virtue epistemology counts with an excellent motivation. After many failed analyses of knowledge, the notion of epistemic virtue was introduced by Sosa (1980, 1991) with the aim of providing a unifying theoretical framework under which long-standing epistemological problems could be solved and rival theories could coexist.¹⁹ With the years, pure virtue epistemology—again by the hand of Sosa (e.g., 2007, 2010, 2011)—has gone to the effort of further developing such a framework with the introduction of *performance normativity*. More specifically, current pure virtue epistemology finds its motivation in a straightforward application of the normativity of performances in general to cognitive performances. The first step is to assume that beliefs are ways to perform cognitively (i.e., cognitive performances). The second step consists in applying normative properties used to assess performances in general to evaluate beliefs. For example, in the same way as we can deem an agent's performance (e.g., a shot) as 'apt' when it is successful because of ability, we can attribute 'aptness' to a cognitive performance (i.e., a belief) when it is successful (i.e., true) because of cognitive ability. The third step consists in accounting for the nature, value and normativity of knowledge in terms of these simple theoretical resources.²⁰

There is no doubt that failing to eliminate dangerous epistemic luck is an important shortcoming of pure virtue epistemology, but it is a shortcoming that it can to a large extent offset with its capacity to tackle long-standing epistemological problems within an elegant and unifying theoretical framework: performance normativity. In this sense, appending modal principles to virtue epistemology independently of this framework just for the sake of handling counterexamples (i.e., without any further motivation) can be fairly regarded as yet one more instance of the kind of ad hoc theorizing that has already turned many epistemologists into skeptics about the analytical project. In sum, if impure theorists want to steer clear of the ad hocery objection, they need to provide a solid rationale for justifying that the concept of knowledge features an independent anti-luck layer on top of

Footnote 18 continued

1981) can be construed along these lines. For relevant discussion on the relation between knowledge and justification more generally, see Foley (2004).

¹⁹ One of the main motivations that Sosa (1980) had for introducing the notion of epistemic virtue (reliabilistically construed) was that it is a way of overcoming the incompatibility between coherentist and foundationalist conceptions of the structure of knowledge: grounding beliefs on perception or memory—being grounded on reliable cognitive faculties is a foundationalist base property—is as intellectually virtuous as believing propositions that cohere with the rest of believed propositions, insofar as both contribute towards getting one to the truth.

²⁰ More specifically, (a) Sosa and other pure virtue epistemologists typically define knowledge as apt belief; (b) as just pointed out, they explain the normativity of knowledge as an instance of performance normativity; finally, (c) they account for the value of knowledge over mere true belief in terms of the thesis that achievements, which have final value, are apt successes: since knowledge is an apt success of a cognitive sort, it is also an achievement (of a cognitive sort), so it has final value—unlike mere true belief, which only has instrumental value. Finally, in a new admirable effort to unify philosophical theories and debates, Sosa (2015) has shown that the idea of aptness of performance can be also used to account for perception and action.

cognitive agency. As we will see next, they find it in different answers to the question of what the point of the concept of knowledge is.

4 The point of the concept of knowledge and the modal account of luck: a failed motivation

4.1 Anti-luck virtue epistemology

In order to motivate anti-luck virtue epistemology, Pritchard (2012a) starts by assuming Edward Craig's genealogical account of the concept of knowledge, which is based on a narrative about how our present concept of knowledge evolved out of a more primitive concept whose function was to *flag good informants* in a society of humans with informational needs (Craig 1990).

On the basis of this assumption, Pritchard argues that flagging a good informant might either mean identifying an informant that is reliable (i.e., an informant possessing the relevant reliable cognitive abilities) or identifying information that one can rely on (i.e., information that would not lead one astray), and claims that the ambiguity imposes two independent demands on the concept of knowledge: an *ability constraint* (when one knows, one's cognitive success must be the product of an exercise of cognitive ability) and an *anti-luck constraint* (when one knows, one's cognitive success must not be a matter of luck).

Pritchard motivates the adoption of explanatory modest virtue epistemology by a straightforward appeal to the ability constraint. But to motivate the further adoption of safety, he needs to give independent reasons for cashing out the anti-luck constraint in modal terms. This is why he appeals to the *modal account of luck* as the correct view of luck. The modal account roughly says that an event (e.g., winning a fair lottery) is lucky if and only if it occurs in the actual world but fails to occur in most nearby possible worlds in which the relevant initial conditions for the event are the same as in the actual world.²¹ Accordingly, the *modal account of epistemic luck* says that an agent's belief that *p* is true by luck if and only if in most close possible worlds in which she forms a belief in *p* in the same way as in the actual world *p* is false.²² According to safety, knowing that *p* entails that in most close possible worlds in which the agent forms a belief in *p* in the same way as in the actual world *p* is true. So safety is the perfect anti-luck condition when epistemic luck is so understood.

The assumption of the modal account of luck is the only way Pritchard can prevent an open season of unwelcome uses of his anti-luck constraint. For if he did not interpret it in modal terms, the constraint could be used to motivate the adoption of non-modal anti-luck principles, such as Unger's anti-accidentality principle that if *S* knows that *p*, it is not accidental that she is right that *p*.²³ But if the anti-luck

²¹ See Pritchard (2014) for a recent defense of the modal account of luck.

²² See Pritchard (2005) for a seminal defense of the modal account of epistemic luck.

²³ See Unger (1968) for an early anti-luck analysis of the concept of knowledge.

constraint could be used to motivate non-modal anti-luck principles, nothing would prevent pure virtue epistemologists from using it to motivate the adoption of their own views. Or is it not the elimination of knowledge-undermining luck (e.g., the kind of luck involved in Gettier-style cases) one of the main goals of virtue epistemology? Therefore, Pritchard does need to assume that the modal account of luck is the correct view of luck, but as it turns out, it is not. Let's see why.

A *coincidence* is an inexplicable event whose constituents are produced by independent causal factors in such a way that we cannot explain why they come together. More specifically, there is no common nomological antecedent of the components of the coincidence or a nomological connection between them to which we can appeal to explain why the coincidence in question occurs.²⁴ Events that arise out of coincidence are lucky, and the modal account of luck fails to account for an important subset of them: lucky events arising out of a coincidence where at least one of the two unrelated components is bound to occur. Consider an example:

Takeover

For the past three months, hundreds of corporations have been secretly trying to take over Sansa's firm, an event that would have very unwelcome consequences for her workers: they would be fired. Sansa knows how to stop these hostile takeovers. For any attempt, she just needs to file a legal complaint via an online submission system. However, unbeknownst to Sansa, there is a problem with the document targeted against the takeover attempt of company number 978, the Cersei Group. An unusual interference in the data stream has modified the contents of the submitted file in such a way that the competent authority has received a document with so many arguments justifying the acquisition that it has decided to give green light. When Cersei, the CEO of the Cersei Group, is about to seal the takeover effectively, the Stannis Group (its long-standing competitor) discloses a scandal that makes Cersei's company's shares drop 99%—no one at the Stannis Group, not even Stannis, the CEO, knows that Cersei was trying to take over Sansa's firm. As a consequence, the corporation goes bankrupt and the takeover does not succeed. In fact, it is the Stannis Group that takes over Cersei's company. The disclosure of the scandal was part of Stannis's meticulous and independent plan to bring down and take over the competition. As luck would have it, it was scheduled one year ago at coincidentally the same time Cersei was about to close the takeover.

Sansa's workers are lucky that the Cersei Group fails to bring the hostile takeover to an end and that in this way they can keep their jobs. The reason why this is lucky for them is that it arises out of a coincidence composed of two unrelated events: (a) the unusual interference in the data stream modifying the contents of the decisive document that would have instantaneously stopped the takeover and (b) the intervention of the Stannis Group. While (a) is a modally fragile event in that it would fail to occur in most close possible worlds, (b) would occur in all of them

²⁴ This account is due to Owens (1992), which is built on some ideas by Hart and Honoré (1959: 74).

because of Stannis's meticulousness and determination to carry out his own takeover plan. This means that in *all* close possible worlds Sansa's workers keep their jobs due to a coincidence, which proves the modal account of luck wrong—the account predicts that it is not lucky that they keep their jobs.

The inability to account for the many structurally equivalent cases of luck by coincidence that can be generated is a major flaw of the modal view.²⁵ Why is this negative for anti-luck virtue epistemology? Because, as we saw before, the modal account of luck is the crucial step between the anti-luck constraint and Pritchard's adoption of the safety principle. If the nature of luck is not modal—at least as Pritchard understands it—why should the epistemic principle imposed by the anti-luck constraint be a modal principle? In conclusion, Pritchard fails to motivate the transition to an impure virtue epistemological account. The problem is even worse: as we will see, the assumption of the modal account of luck and of a corresponding safety principle prevents anti-luck virtue epistemology from eliminating knowledge-undermining luck out of coincidence.²⁶

4.2 The safe-apt view

Kelp (2011, 2013a, b) gives further arguments against Pritchard's motivation for anti-luck virtue epistemology. For instance, he argues that since Pritchard thinks that the concept of knowledge features a condition that rules against agents we cannot rely on, he cannot explain cases of knowers who cannot be relied on. Kelp also attacks Craig's genealogical account and thinks that anti-luck virtue epistemology is better off without it. According to Kelp, the best motivation for Pritchard's view as well as for his own account comes from an alternative explanation of the point of the concept of knowledge, namely that the concept of knowledge serves to mark when a given agent is entitled to inquire no further into a given question, i.e., as an *inquiry stopper*. Let's see whether this serves to motivate his own account of knowledge.

Kelp starts with the observation that we are sometimes entitled to inquire no further into a question when we lack the answer to that question. For example, you might look at someone who is looking at an object that you cannot see and you might be interested in knowing whether the object has a certain property. Since you know that the person you are looking at counts with a reliable visual system, you can point out that she is entitled to inquire no further into the issue. Attributing knowledge to that person in that situation is appropriate for that reason, Kelp argues,

²⁵ Its main rival, the *lack of control account of luck*—roughly, the view that an event is lucky for an agent if and only if the agent fails to exercise control over the event—naturally explains *Takeover* by pointing out that Sansa's workers are lucky to keep their jobs because they lack control over the two components of the coincidence (the data interference and the Stannis Group's intervention). See Broncano-Berrocal (2015a) for a recent version of this view.

²⁶ More recently, Pritchard (2015) also appeals to the modal accounts of risk and epistemic risk to motivate the adoption of safety—roughly, the view that risky events are events that could very easily occur and the view that risky beliefs are beliefs that could be very easily false. Since Pritchard's modal accounts of luck/epistemic luck and risk/epistemic risk are meant to go hand-in-hand, the arguments presented here against the former apply *mutatis mutandis* to the latter.

and he wonders what sorts of conditions would govern a concept that serves that function and is applicable when we ignore the answer to a given question. Here is what he says:

To begin with, we would expect it to respect a *modal condition*: were we to find out that the attributee might so easily have been mistaken, we could not attribute to him an entitlement to inquire no further into the issue. (...) At the same time, we would also expect a concept with the envisaged function to respect a *competence condition*: were we to find out that the attributee does not have the relevant cognitive competence, we could not attribute to him an entitlement to inquire no further (Kelp 2013a: 277; emphasis added).

However, why should we prefer a modal or a competence condition over other kind of epistemic conditions? For example, if an internalist found out that the attributee is not internally justified, she would not attribute an entitlement to inquire no further into the issue to that person. The problem is that Kelp's reasoning is an epistemological wild card, in the sense that it can be used by any leading internalist or externalist view available on the market for claiming its right to use its own necessary conditions for knowledge so as to argue that those are precisely the ones that govern the concept of knowledge when serving the function he thinks the concept has. Given the recent history of epistemology, that permissiveness is a formidable drawback to an impure account of knowledge that struggles to be chosen over its less complex and more elegant pure rivals.²⁷

This argument undermines Kelp's motivation for the safe-apt view even if he is right about the point of the concept of knowledge. For his account shows at best that knowledge functions as an inquiry stopper when governed by safety and competence criteria, but does not show that it possesses such a function *only* when governed by such criteria. That being so, other theorists may still use Kelp's account for postulating their own conditions on knowledge as criteria for the function of stopping inquiry.²⁸

²⁷ Kelp could still appeal—as Pritchard's does—to Craig's narrative about how our present concept of knowledge evolved out of a more primitive concept without thereby committing himself to Craig's thesis that the function of the concept of knowledge is to flag good informants. In fact, in his 2011 paper he gives an alternative to that thesis adopting the essentials of Craig's narrative. Accordingly, in reply to my argument, Kelp could argue that the primitive concept of knowledge featured his two epistemic conditions because the relevant notions of competence and easy mistake were plausibly mastered even by agents at that primitive stage, but not so easily would they have mastered internalist notions, such as the notion of defeater, which is too technical for agents with a primitive concept of knowledge in mind. As a rejoinder, my argument does not rely on any crucial link between Craig's primitive concept of knowledge and internalist conditions. It suffices that epistemologists with very different sympathies can in principle use Kelp's account of the point of the concept of knowledge and/or Craig's narrative to motivate their own accounts of knowledge.

²⁸ Surprisingly, the argument may still work even under the assumption of the stronger claim. Suppose that the competence condition is a fundamental criterion governing the function of the concept of knowledge. Compatibly with it as well as with the essentials of Craig's narrative that Kelp (2011) upholds, it is plausible to suppose that agents mastering the primitive concept of knowledge could have applied it solely to open-minded, epistemically humble or conscientious inquirers, or more generally, to inquirers exhibiting a series of character traits. As virtue responsibilists sometimes point out, agents may possess such intellectual virtues independently of how reliably they are concerning a range of questions or

In conclusion, while Pritchard's motivation for anti-luck virtue epistemology rests on a false assumption (the modal account of luck), Kelp's motivation for the safe-apt view is vacuous in the sense that epistemologists with very different sympathies could use it to motivate their own accounts of knowledge—which is fully compatible with Kelp's account of the point of the concept of knowledge being correct. Absent the motivation for the two leading impure views, there is one less major reason why we should prefer them over their more appealing pure cousins. The only asset they have left is that they seem to be able to eliminate knowledge-undermining luck.

5 Knowledge-undermining luck: a failed elimination

Kelp and Pritchard think that the main advantage of combining the safety principle with their modest virtue conditions is that, while the latter can satisfy the ability intuition, the former can do all the anti-luck work. As we will see next, however, their views fail in their purpose of eliminating malignant epistemic luck. Let's start with anti-luck virtue epistemology.²⁹

5.1 Anti-luck virtue epistemology

Recall that the modal view of luck fails to account for many lucky events that result from coincidences. Anti-luck virtue epistemology inherits this problem via the assumption of the modal view of epistemic luck, which fails to account for many lucky true beliefs out of coincidence.³⁰

By way of illustration, if Sansa forms the belief that her company will not be taken over on the basis of her excellent inductive justification to believe that filing a legal complaint will make the competent authority stop any hostile takeover, her belief does not amount to knowledge because it is by coincidence that she gets things right.

To see this, notice that the structure of this *epistemic version of Takeover* quite unsurprisingly mirrors the structure of standard Gettier-style cases. Standard Gettier-style cases typically involve beliefs that are true by coincidence, or as they

Footnote 28 continued

how easily could they have been mistaken concerning a particular question. This means that a competence condition could have been a fundamental criterion for the primitive concept of knowledge without that serving Kelp's purposes of motivating reliability-based conditions such as aptness or safety.

²⁹ See also Broncano-Berrocal (2014a) for an argument to the conclusion that the virtue condition of anti-luck virtue epistemology renders the account insufficient for knowledge.

³⁰ Someone might object that, so far, it has been assumed without argument that the modal account of epistemic luck must be a closeness account. However, that might be conceptually mistaken insofar as it is seems in principle possible to conceptualize epistemic luck in terms of distant possible worlds. Freitag (2014), for instance, thinks that there is such a thing as epistemic luck that discloses itself in distant possible worlds. But as I argue in Broncano-Berrocal (2015b), Freitag provides no convincing reason for the existence of 'distant' epistemic luck. Absent any similar positive proposal in the epistemological literature, the identification of modal accounts with closeness accounts seems reasonable.

are sometimes described, they have a ‘double stroke of luck’ structure (Zagzebski 1994). In a first stroke of bad luck, the competent authority allows the takeover because the file that Sansa submits is modified by an unusual data interference. In a second stroke of good luck, the takeover fails because the Cersei Group goes bankrupt due to the disclosure of the scandal by the Stannis Group. In this way, since the epistemic version of Takeover is structurally equivalent to standard Gettier-style cases, claiming that Sansa knows entails accepting that paradigmatic cases of malignant epistemic luck are cases of knowledge, which obviously goes against the widespread intuition that knowing by luck is impossible. Against this intuition, however, anti-luck virtue epistemology credits Sansa with knowledge. Let’s see why.

Its virtue condition (explanatory modest virtue epistemology) says that the exercise of the agent’s cognitive abilities must be a salient but not necessarily the most salient factor in the set of causal factors that explain her cognitive success—in the present case, in the explanation of why Sansa gets it right about her company not being taken over. Sansa’s legal skills, and in particular her excellent inductive basis to believe that the kind of legal document submitted would stop any takeover, definitely explain why her company is not purchased by hundreds of corporations. Consequently, it is hard to see how Sansa’s cognitive abilities are not a salient factor in the explanation of her cognitive success.³¹

The bad news for Pritchard is that, since safety is also satisfied, he cannot explain away the intuition that Sansa does not know. Nearby possible worlds are worlds in which (a) the unusual data interference does not occur and hence Sansa’s document continues to be as effective as usual and in which (b) the release of the scandal still dooms the Cersei Group to bankruptcy. In those worlds, Sansa’s company is not purchased. So her true belief is safe. In addition, her belief is luckily true because, like Gettiered beliefs, it is true by coincidence. Therefore, anti-luck virtue epistemology makes knowledge compatible with coincidentally true belief. In other words, it fails to eliminate knowledge-undermining luck.³²

³¹ Moreover, if for any given reason it were incorrect to say that the condition on knowledge stated by explanatory modest virtue epistemology obtains, then it would be wrong to say that it held in previous cases in which Sansa came to know that her company was not going to be taken over by other companies after submitting the same kind of legal document and after inferring it *from the same inductive grounds*. By parity of reasoning, if the condition stated by explanatory modest virtue epistemology was satisfied in past cases of inductive knowledge, then it must be also satisfied when the Cersei Group tries to take over Sansa’s firm. Either one has to deny that there was knowledge in those past cases, or else to admit that the right-hand side of explanatory modest virtue epistemology obtains in the latter case (a case of ignorance). For obvious reasons, none of the options is appealing to anti-luck virtue epistemology.

³² Let me dispel some potential worries. (a) Even if it turns out that *Takeover* does not prove Pritchard’s modal account of luck wrong—as defended in the previous section—the epistemic version of *Takeover* does show that Pritchard’s modal account of epistemic luck is incorrect—the case shows that knowledge can be undermined by luck when safety and the ability condition of anti-luck virtue epistemology hold. That is all is needed to establish the conclusion that anti-luck virtue epistemology lacks adequate motivation and fails to exclude knowledge-undermining luck. (b) One could reply to this that, in order to motivate his account, Pritchard does not need that safety is sufficient for the absence of knowledge-undermining luck; instead, he only needs it to be necessary. Although this is already problematic insofar as there are cases of knowledge without safety (see Sect. 5.3), Pritchard could argue on that basis that explanatory-salience-style virtue accounts of knowledge are not adequate unless aided with a safety

Pritchard could complain that this kind of counterexample only proves the weak interpretation of anti-luck virtue epistemology wrong, since the *safety of Sansa's belief* is down to the coincidence and has nothing to do with the exercise of her cognitive abilities, and in particular nothing to do with her excellent inductive reasoning. The strong interpretation of anti-luck virtue epistemology, the reply continues, is not in trouble insofar as it says that it is the safety of Sansa's cognitive success that must be partly explained by the exercise of her cognitive abilities.³³

Unfortunately, the strong version does not seem to be in a better position to eliminate luck if all it requires is that safety is *partly explained* by cognitive ability. Consider the following *modified version of the case*. Suppose that Sansa only realizes that her company will not be taken over after having written an anonymous comment in a financial blog in which she reflects at length on how good the inductive grounds are for believing that filing a legal complaint is the best method that anyone could use to stop a hostile takeover. Suppose in addition that Sansa does not disclose that the Cersei Group is trying to take over her firm, so the whole operation remains secret. Now suppose that the reason why the Stannis Group decides to take action against Cersei's company is that Stannis, its CEO and regular reader of the blog, is impressed by the excellent reasoning presented by the anonymous writer—in this case Sansa. Finally, suppose that there is no chance that Sansa does not write that comment and that Stannis does not read it, so these two things continue to occur in all nearby possible worlds.³⁴

In the modified case, the exercise of Sansa's reasoning abilities is an important albeit deviant factor in the explanation of why her belief is safe—i.e., of why she gets things right in the actual and nearby possible worlds. First, Sansa submits the online document and, as before, its contents change due to the interference in the data stream. Right after that, she writes the comment in the blog and only then forms the belief that her company will not be taken over on the basis of her excellent inductive reasoning. *Because of that very reasoning*, Stannis decides to intervene. Finally, Stannis's intervention cancels out the unwelcome consequences that the modification of the file's contents would have (i.e., the takeover).

The fact that Stannis intervenes because of being impressed by Sansa's reasoning does not alter the coincidence in question: (a) the unusual interference in the data stream still distorts the decisive document; (b) the Stannis Group still intervenes stopping the takeover. In this way, Sansa's true belief is no less by coincidence than before. But since all the conditions of the strong reading of anti-luck virtue

Footnote 32 continued

condition even if safety does not suffice for meeting the relevant anti-luck constraint either. In this sense, he could argue that, even if the counterexample is successful against anti-luck virtue epistemology as well as against his modal account of epistemic luck, it is not exclusively problematic for anti-luck virtue epistemology but for any version of virtue epistemology (pure or impure) that features an ability condition understood in terms of explanatory salience. This is untrue, however. The epistemic version of *Takeover* is not problematic for explanatory robust virtue epistemology. After all, the fact that Sansa comes to believe truly that her company will not be taken over is not *fully* explained by her excellent inductive grounds, but partly by the lucky coincidence that stops the takeover.

³³ See Sect. 2 for the two interpretations of anti-luck virtue epistemology.

³⁴ A similar but relevantly different Gettier-style case is proposed by Greco (2012).

epistemology are in place, the account fails to eliminate knowledge-undermining luck.³⁵

5.2 The safe-apt view

Let's move now to Kelp's safe-apt view. We have seen that safety is of no help in cases of coincidentally true belief with the structure of the previous cases. Dispositional modest virtue epistemology, however, might be of help. Remember that a cognitive success manifests cognitive ability only if the conditions are appropriate. In the previous cases, Sansa's actual circumstances are not appropriate, since the online submission system works reliably only in the absence of data interferences, which means that Sansa's safe belief does not manifest that reliability under her present circumstances. In other words, Sansa's belief is safe but not apt—dispositional robust virtue epistemology can handle the cases for the same reason.

Since the safe-apt view says that knowledge is safe, apt belief, the view is able to rule out the previous cases as cases of knowledge. But that comes at a cost. Recall that Kelp and Pritchard's motivation for introducing an independent safety condition is that it can eliminate knowledge-undermining luck. In the previous cases, however, it is the virtue condition which does all the anti-luck work. Add to that point the fact that explanatory robust virtue epistemology can handle the cases as well—Sansa's competent inductive reasoning to the conclusion that her company will not be taken over is not the most salient factor in the explanation of her cognitive success: the intervention of the Stannis Group is also a salient factor. The result is a tie between safety, on the one hand, and the two leading versions of pure virtue epistemology on the other. For while safety cannot eliminate epistemic luck out of modally robust coincidence (i.e., the kind of knowledge-undermining luck involved in the previous cases), the two pure views can, and while they cannot eliminate environmental luck (i.e., the kind of knowledge-undermining luck involved in the fake barn case), safety can.

Of course, the safe-apt view can handle both the fake barn case and the previous cases. But at that point of the dialectic, there is an impasse between two theoretical options. Either one can adopt one of the two well-motivated pure virtue-theoretic accounts of knowledge (either dispositional or explanatory robust virtue epistemology), which nevertheless fail to handle some cases of knowledge-undermining luck, or else one can opt for the view that gets the cases right but whose motivation is either vacuous or that it gets the cases right (the safe-apt view). More generally, the dilemma that virtue epistemologists face is to decide which problem they want to live with and eventually solve: the problem of eliminating knowledge-undermining luck or the ad hocery worry raised in Sect. 3.

³⁵ An alternative conception of safety that tolerates no error in *very close* possible worlds—which Pritchard (2007) uses to explain lottery cases—would not save anti-luck virtue epistemology. If we fill in the details of the case to the effect that, in order to prevent the scandal from coming out significant changes to actuality would be needed, then very close possible worlds are worlds in which the scandal is disclosed and Cersei's company goes bankrupt. Since in *all* such worlds Sansa's company is not taken over, the modified safety condition holds as well.

Fortunately, there is no need to be torn in two. The safe-apt view also fails to rule out cases of environmental luck as cases of knowledge, which means that the second option is automatically discarded and the problem of eliminating knowledge-undermining luck becomes a problem for all views discussed thus far, pure and impure ones. Consider an inferential version of the fake barn case:

Inference

While traveling by train, Barnaby averts his eyes from the screen of his cellphone to admire the landscape of the barn county. Unbeknownst to Barnaby, a construction company has recently replaced all barns in the area with barn façades, except for the only one listed as national heritage, which has been left intact. As luck would have it, the first object with the appearance of a barn he looks at is that barn. He consequently forms the true belief that the object in front of him is a barn. From that belief, he competently infers (via existential generalization) that some object in the field is a barn—alternatively, that there is a barn in the field.

Barnaby's inferred belief is apt (in the sense that it manifests Barnaby's reasoning abilities), insofar as he competently infers it from the apt belief that the object in front of him is a barn—as in a good case of perception, the premise belief is also apt because there is a non-deviant causal chain from the object perceived to the belief formed. In sum, the safe-apt view's virtue condition is met.

However, Barnaby does not know that there is a barn in the field (or that some object in the field is a barn) because, like in the non-inferential version of the fake barn case, it is a matter of environmental luck that he gets things right. The question that immediately comes to mind is whether the luck in play (and hence the absence of knowledge) can be explained in terms of the violation of the safety condition. Unfortunately for Kelp, the answer is negative. To see this, note that Barnaby's belief in the proposition that some object in the field is a barn is true in the actual world and continues to be true in nearly all (if not all) nearby possible worlds in which Barnaby holds it on the same deductive basis. The reason is simple: there is always a barn in the field. This means that, despite being lucky, a belief can be apt *and* safe.

In reply, Kelp could argue that Barnaby's belief is not knowledge because the premise true belief is by luck and hence not known and, for some reason, inference transmits lack of knowledge in *Inference*. More specifically, Kelp could argue—using the safe-apt view's theoretical resources—that the reason why Barnaby's safe apt belief in the conclusion does not amount to knowledge is that the premise apt belief is unsafe and, in general, competent inference from an unsafe (albeit apt) belief fails to transmit knowledge.

However, such an explanation is untenable—or at least must be regarded as ad hoc—unless it is grounded on a general principle of lack of knowledge transmission. Unfortunately for Kelp, there does not seem to be any such principle to which he can appeal to support the said explanation with some degree of universality. In particular, the following principle is not necessarily true: if S does not know that p —e.g., due to the fact that S 's belief that p is unsafe—and S competently deduces q from p and believes that q as a result, then S does not come to know that q on the

basis of that deduction. Consider a counterexample to this principle by Baumann (2014). Suppose that Peter forms the true belief that Jack is singing in the other room. However, it is usually Jill who does the singing around there and Peter cannot distinguish Jack's from Jill's voice. In this way, Peter's belief is luckily true (it is unsafe). However, if, from that unsafe belief, Peter infers that someone is singing in the other room, then his belief in the conclusion amounts to knowledge.

The challenge for Kelp would consist, then, in justifying why the lack of knowledge transmission principle above applies to *Inference* while explaining away Baumann's case as well as similar cases of knowledge.³⁶ It is unclear, however, why the lack of safety of Barnaby's apt premise belief is supposed to prevent the inferred safe apt belief from being knowledge when in the parallel case Peter's safe apt belief amounts to knowledge despite the fact that the belief from which he infers it is unsafe and true by luck. In the absence of a satisfactory explanation, there is no other option but to consider Kelp's safe-apt view insufficient for knowledge.

In conclusion, not only the two main impure views are inadequately motivated, but they are in no better position than pure accounts when it comes to eliminating knowledge-undermining luck. Even if none of the discussed theories seems to yield an adequate anti-luck epistemology, simplicity and the fact that the latter do count with a solid motivation are two weighty reasons that tip the scales in favor of the pure side of virtue epistemology.

5.3 The safety-based approach and the safety dilemma

Those who assume that epistemic luck is to be understood in modal terms, and specifically in terms of close error possibilities, will probably attempt to find alternative ways (other than impure virtue epistemology) to develop a safety-based approach. After all, safety is the perfect anti-luck principle when knowledge-undermining luck is modally understood—I have argued that, as a matter of fact, this is not the correct way to account for it. As we have seen, introducing safety independently of virtue epistemology is problematic, but one might still try to have safety on board by deriving it from one's preferred virtue condition or by reinterpreting it as a normative property of performances in general so as to use it, in turn, to account for knowledge in terms of the satisfaction of the ability intuition.³⁷ The respective views of Sosa (2015), Turri (2011) and to some extent Carter (2014) and Jarvis (2013) can be construed along these lines.

Without getting into the details of these views, there are reasons to think that the prospects for a safety-based approach in general are dim. The crux of the matter is that, as several commentators have insistently argued, safety is not necessary for knowledge. To prove this, they use cases in which there is the strong intuition that an agent knows that p despite the fact that she could easily have believed p falsely. Consider the following case by Bogardus (2014). Imagine that you consult the most

³⁶ For example, cases of inferential knowledge from falsehood (see Warfield 2005).

³⁷ The resulting views, even if safety-based, are pure versions of virtue epistemology insofar as safety is not motivated independently of the core thesis.

accurate functioning clock on Earth. Intuitively, you come to know the time. But suppose that there is an isotope such that, if it were to decay, it would disrupt the clock's internal mechanism. If we stipulate that the isotope could easily have decayed at any moment but fortunately not at the precise time at which you have looked at the clock, then it seems that you have acquired knowledge of the time without your belief being safe—your belief is true, but the clock could easily have stopped making you get things wrong in nearby worlds.³⁸

Abstracting from the details of this and other examples, the general problem is that safety does not allow to distinguish between cases in which close error possibilities are knowledge-undermining, such as the fake barn case, and cases in which they are epistemically harmless, such as the clock case. So, unless more is said, all indicates that safety is not necessary for knowledge.³⁹

This result is devastating for anyone who aims to use safety as an anti-luck condition—pure virtue epistemologists included. For suppose, for the sake of the argument, that safety were able to eliminate all forms of knowledge-undermining luck—this paper has shown that this is actually not the case—then the following dilemma arises: if a virtue epistemological view does not include safety as a necessary condition for knowledge, then it is too weak to eliminate knowledge-undermining luck (*first horn*), but if it does, then it is too strong to account for knowledge involving epistemically harmless close error possibilities (*second horn*). In this way, the *safety dilemma* pulls safety-based virtue epistemological accounts in opposite directions—obviously, this does not speak in favor of (current versions of) impure virtue epistemology.⁴⁰

6 The ability approach: towards a robust enough virtue epistemology

Not all hope is lost on the virtue epistemology front. There is one further promising strategy that virtue epistemologists can follow to get around the environmental luck problem and retain the solid motivation of pure virtue epistemology (namely, the performance normativity framework). In fact, the results of this paper help bolster

³⁸ Other counterexamples may be found in Baumann (2008), Comesaña (2005), Kelp (2009), Neta and Rohrbaugh (2004), and Sainsbury (1997).

³⁹ One way out would be to invoke a suitably formulated principle for individuating belief-forming methods. The idea is to achieve the following result: cases in which close error possibilities are epistemically harmless are cases in which the method used in close possible worlds is different from the one used in the actual world, whereas safety requires true belief across close possible worlds in which the agent uses the *same* belief-forming method as in the actual world—one also has to ensure that cases of knowledge-undermining error possibilities are such that the agent uses the same method in all close possible worlds, so that safety fails. See Broncano-Berrocal (2014b) for this kind of strategy to save safety from counterexamples, but see Bogardus and Marxen (2014) for some objections. An alternative way out would be the outright denial of the intuition that the counterexamples are cases of knowledge. This is Pritchard's strategy (2012b), but see Bogardus (2014) and Kelp (2016) for compelling arguments against it.

⁴⁰ The same kind of dilemma can be generated for combinations of virtue epistemology with the sensitivity principle, insofar as there are plausible cases showing that sensitivity is not necessary for knowledge. See footnote 13.

this approach insofar as it no longer aims to supplement virtue epistemology with safety.

The *ability approach* starts with the observation that the first horn of the safety dilemma assumes that an account which does not feature safety as a necessary condition for knowledge is too weak to handle dangerous epistemic luck—for the theories at which the dilemma is targeted no alternative (non-modal) anti-luck condition is acceptable. But this assumption can be certainly rejected. Indeed, as soon as one grants that the modal account of epistemic luck is incorrect, resorting to safety becomes much less appealing. The alternative then consists in giving a more elaborated account of the notion of ability in order to explain how success from ability (of the right kind) is incompatible with success from luck.⁴¹

What follows is a sketch of such an account.⁴² The guiding idea is that types of abilities can be distinguished by the functional role they play. On the one hand, there are abilities that play a *task-completion role*, such an archer's disposition to hit a certain kind of target or an agent's disposition to form true beliefs about propositions pertaining to a specific domain. They are aimed at delivering outputs (e.g., successful hits, true beliefs) which serve to complete tasks (hitting targets, forming true beliefs) under appropriate conditions. On the other hand, there are abilities that are not aimed at completing tasks in this way, but play a *precautionary* or *protective role* instead. They are dispositions to stop the completion of a task by a task-completion ability when the conditions are not suitable for it. For example, a world-class archer might be disposed to refrain from delivering arrows whenever the wind speed is too high to hit the mark. Professional poker players are typically disposed to avoid believing a great number of bluffs, especially when they come from amateurs. In general, we tend not to form visual beliefs when the light conditions are bad.

All precautionary or protective abilities are sensitive to the inappropriateness of actual circumstances, but their protective role also typically involves *sensitivity to the inadequacy of circumstances in nearby possible worlds*, in the sense that they would stop the completion of tasks in circumstances that are similar to the actual ones. In addition, since the kind of appropriate circumstances in which precautionary abilities operate reliably (i.e., their range of reliability) are actual or possible circumstances which are inappropriate for the corresponding task-completion abilities, the two types of abilities operate with *independent levels of reliability*, that is, although they often go hand-in-hand, it is perfectly possible for an agent to be competent at completing a task but fail to be sensitive to the inappropriateness of the situation, and the other way around.⁴³ Finally, if a precautionary ability stops the fulfillment of a task when the circumstances in the actual and nearby worlds are appropriate for it, its operation is unreliable. For example, a distrustful person in a place where it is openly known that everybody tells the truth does not count as

⁴¹ Alternatively, how success can depend on ability (of the right kind) more so than it depends on luck. See Carter (2014) for this idea.

⁴² See Broncano-Berrocal (2016) for a full statement of this account.

⁴³ See Sosa (2015) for a rich discussion of risk assessment and the role played by monitoring abilities.

competent. So in order to function reliably in circumstances that are appropriate for completing tasks, precautionary abilities must remain inactive. Otherwise they fail to fulfill their precautionary role by being *overprotective*.

There is the question of whether precautionary abilities are *innate*, *acquired* or *temporarily possessed*. The answer is: all of them, as long as they play their distinct functional role. Agents might have the ability to refrain from completing certain tasks in virtue of innate or acquired dispositions (e.g., an archer might be endowed with an innate ability to correctly gauge the wind speed, or she might simply learn how to do it), but they might well be as contingent as a disposition to use relevant information about the inadequacy of the circumstances (e.g., archers might simply consult anemometers or trust the reliable testimony of experts). The same applies to cognitive tasks.

Finally, cognitive precautionary abilities are *second-order but not necessarily reflective* in that, while they are dispositions to stop the activation of other cognitive dispositions, agents need not be aware of their operation. This means that the kind of knowledge they help produce is not necessarily reflective. In this sense, they are not tied to any distinction between different kinds of knowledge—e.g., first-order, reflective or any superior type of knowledge that comprises both.⁴⁴ What they help produce is knowledge, *simpliciter*.

If one grants that we are endowed with belief-forming dispositions as well as precautionary cognitive abilities, in other words, if one grants that our *epistemic resources* when it comes to a cognitive task (such as forming a true belief) need to comprise both kind of abilities vis-à-vis that task, and one also sympathizes with the core thesis of virtue epistemology—namely, that knowing is just a matter of succeeding cognitively because of cognitive ability—one will be happy to embrace the idea that knowledge arises just in case the agent gets it right because of both kinds of abilities, i.e., just in case all epistemic resources are in place when succeeding cognitively.

In this way, the first step towards a *robust enough virtue epistemology* consists in understanding *apt belief* as belief that is true because of all epistemic resources, whereas *partly apt belief* is belief that is true because of part of them. Key to this version of virtue epistemology is the idea that a belief that is partly apt falls short of knowledge, whereas only (completely) apt belief amounts to it. Like other pure views, the view is located in a performance normativity framework, as it evaluates belief with normative properties of performances in general—namely, with aptness, the property of being successful because of ability. Finally, by further assuming the manifestation reading of the ‘because of’ relation, the account understands knowledge as a cognitive success (an apt belief) that *manifests* all epistemic resources, i.e., not only first-order belief-forming dispositions but also corresponding precautionary cognitive abilities.⁴⁵

⁴⁴ See especially Sosa’s distinction between animal, reflective knowledge and knowing full well (Sosa 2015).

⁴⁵ See Broncano-Berrocal (2016) for a full statement of this account.

A theoretical benefit of adopting the manifestation reading is that the idea that knowledge is a cognitive success because of ability need no longer be at odds with the fact that factors beyond one's cognitive agency often contribute to that success (such as a reliable speaker). After all, dispositions need *reciprocal partners* to be manifested—e.g., when salt dissolves in water, this manifests both the disposition of salt to dissolve in water and the disposition of water to dissolve salt. So on what grounds should then one concede that testimonial knowledge does not manifest (or slightly manifests) the hearer's disposition to select good informants (and avoid blatantly insincere ones) when trusting the testimony of a carefully selected speaker who has in turn manifested her disposition to pass accurate information? That the hearer's cognitive success manifests her cognitive abilities when competently trusting a speaker is not only compatible with but also needs that the speaker manifests hers when speaking truly: they are, after all, reciprocal dispositional partners. In this way, contrary to what the *first horn of the attributability dilemma* says, testimonial knowledge does not pose a threat to manifestation accounts when understood in this manner.⁴⁶

In addition, accounting for knowledge in terms of all epistemic resources—rather than just in terms of first-order belief-forming dispositions—helps give a simple diagnosis of cases such as the fake barn case, *Inference* or the epistemic version of *Takeover*, which in turn solves the problem of luck thus avoiding the *second horn of the attributability dilemma*. In a nutshell, what all these cases have in common is that the agents' relevant precautionary abilities (if any) fail to stop belief formation when they should, given that actual or possible circumstances are inappropriate for (first-order) belief formation.

Consider the widely discussed *fake barn case* and its variant *Inference*. In both cases the agents competently form their respective beliefs by means of their first-order belief-forming abilities: they look at real barns, make competent inferences, and so on. But they fail to be sensitive to the fact that circumstances in nearby possible worlds are inappropriate. Those circumstances are such that they would encounter fakes. Fakes in general are designed to deceive, so circumstances involving them cannot be of the sort that are appropriate for standard belief formation. In the fake barn case and *Inference*, the relevant precautionary abilities are reliable only relative to typical inappropriate circumstances for belief formation, such as bad light conditions. Therefore, preventing belief formation in the presence of fake barns is beyond their standard range of reliability. So while their beliefs are partly apt (because they are competently formed), they are not completely apt (because they fail to manifest all epistemic resources).

The epistemic version of *Takeover* can be understood along similar lines: abnormal interferences in the data stream are not part of sort of the circumstances that are typically appropriate for forming beliefs on the assumption that the information sent has arrived unaltered.⁴⁷ Given how the case is described, there is

⁴⁶ See Broncano-Berrolcal (2016) for a more detailed statement of this argument.

⁴⁷ Being sensitive to such interferences is obviously beyond the range of reliability of the precautionary abilities of standard users. While we might distrust an obviously malfunctioning device, we tend to trust

no way in which Sansa could have monitored the data interference and hence come to know—rather than merely believed truly—that her company was not going to be taken over. Therefore, her getting it right does not manifest her precautionary abilities and hence the relevant epistemic resources she would have needed to acquire knowledge—note that the only relevant difference with the fake barn case and *Inference* is that *Takeover* involves *actual* inappropriate circumstances.

This version of pure virtue epistemology can therefore deal with problematic cases of epistemic luck—including the cases that prove impure views wrong—and can thus steer clear of the second horn of the attributability dilemma. However, Kelp (2016) has recently argued that pure views fall prey to a dilemma akin to the safety dilemma. His point is that if pure accounts can handle dangerous epistemic luck, then they must regard cases of knowledge with the structure of the clock case—or *epistemic Frankfurt cases* as he calls them—as cases of ignorance, and *vice versa*.⁴⁸

In reply, I elsewhere argue that precautionary abilities need not track or be sensitive to every actual or modally close condition that would prevent cognitive success, but only to *determining conditions* (Broncano-Berrocal 2016). Success from ability in general can be prevented in at least two ways. One is by removing the ability from the agent before the output is produced; another is by masking its exercise.⁴⁹ In this sense, there are conditions that *enable* the exercise of ability and conditions that *determine* its success or failure. Precautionary abilities need to be sensitive only to the latter.

In the *clock case*, the absence of a decayed isotope simply enables the normal functioning of the reliable clock. By contrast, the presence or absence of fakes of objects (e.g., fake barns) are obviously decisive (i.e., not merely enabling) factors for the success or failure of a perceptual performance which arises out of (already) reliable cognitive ability. In other words, while in the clock case the decayed isotope removes the relevant ability from the agent—insofar as the agent's reliability is conditional on the clock's reliability—in the fake barn case the presence of barn façades masks its exercise. In the former case, there is knowledge because the relevant precautionary abilities need not track factors that disrupt the enabling conditions of the relevant first-order cognitive abilities. In the latter, there is no knowledge because the relevant precautionary abilities fail to track factors that determine the failure of their corresponding first-order abilities. In sum, Kelp's dilemma does not put a period to the ability approach to virtue epistemology.

Those who seek answers to long-standing epistemological problems will find virtue epistemology especially appealing. But the infamous problem of eliminating dangerous epistemic luck threatens to overshadow its main achievements. What this paper has shown is that one of the apparently more promising anti-luck approaches,

Footnote 47 continued

(and competently so) that the information we send will arrive safely. One only has to consider how excellent our inductive basis is to believe that sent e-mails effectively arrive to their addressees.

⁴⁸ See Kelp (2016) for his specific arguments against Greco's and Sosa's respective accounts.

⁴⁹ See Fara (2008, pp. 846–847), for discussion on masking.

impure virtue epistemology, is not only ill-motivated, but also fails. One might quickly infer from this that virtue epistemology is yet one more failed project in the reductive agenda. However, the problem of luck can be more optimistically seen as an opportunity to expand and further improve the theory. Indeed, a better understanding of the notion of ability helps explain how success from ability is incompatible with success from luck. This is the plausible first step towards developing a robust enough virtue epistemology: a view that can steer clear of the attributability dilemma.

Acknowledgements I would like to thank Chris Kelp, Klemens Kappel, Giacomo Melis, Josefine Pallavicini, Bjørn Hallsson and Frederik J. Andersen for helpful discussion and feedback as well as two anonymous reviewers for *Philosophical Studies* for valuable suggestions. Funding was provided by H2020 Marie Skłodowska-Curie Actions (Individual Fellowship).

References

- Baehr, J. S. (2011). *The inquiring mind: On intellectual virtues and virtue epistemology*. Oxford: Oxford University Press.
- Baumann, P. (2008). Is knowledge safe? *American Philosophical Quarterly*, 45, 19–30.
- Baumann, P. (2014). No luck with knowledge? On a dogma of epistemology. *Philosophy and Phenomenological Research*, 89, 523–551.
- Bogardus, T. (2014). Knowledge under threat. *Philosophy and Phenomenological Research*, 88, 289–313.
- Bogardus, T., & Marxen, C. (2014). Yes, safety is in danger. *Philosophia*, 42, 321–334.
- Broncano-Berrocal, F. (2014a). Anti-luck (too weak) virtue epistemology. *Erkenntnis*, 79, 733–754.
- Broncano-Berrocal, F. (2014b). Is safety in danger? *Philosophia*, 42, 63–81.
- Broncano-Berrocal, F. (2015). Luck as risk and the lack of control account of luck. In D. Pritchard & L. Whittington (Eds.), *The philosophy of luck* (pp. 3–27). London: Wiley-Blackwell.
- Broncano-Berrocal, F. (2015b). No luck in the distance: A reply to Freitag. *Theoria*, 82, 89–100.
- Broncano-Berrocal, F. (2016). A robust enough virtue epistemology. *Synthese*,. doi:10.1007/s11229-016-1043-y.
- Carter, J. A. (2014). Robust virtue epistemology as anti-luck epistemology: A new solution. *Pacific Philosophical Quarterly*,. doi:10.1111/papq.12040.
- Comesaña, J. (2005). Unsafe knowledge. *Synthese*, 146, 395–404.
- Craig, E. (1990). *Knowledge and the state of nature. An Essay in conceptual synthesis*. Wotton-under-Edge: Clarendon Press.
- Dretske, F. (1971). Conclusive reasons. *Australasian Journal of Philosophy*, 49, 1–22.
- Fara, M. (2008). Masked abilities and compatibilism. *Mind*, 117, 843–865.
- Foley, R. (2004). A trial separation between the theory of knowledge and the theory of justified belief. In J. Greco (Ed.), *Ernest Sosa and his critics* (pp. 59–71). Hoboken: Blackwell.
- Freitag, W. (2014). Safety, sensitivity and “distant” epistemic luck. *Theoria*, 80, 44–61.
- Gettier, E. (1963). Is justified true belief knowledge? *Analysis*, 23, 121–123.
- Goldman, A. I. (1976). Discrimination and perceptual knowledge. *Journal of Philosophy*, 73, 771–791.
- Greco, J. (2009). Knowledge and success from ability. *Philosophical Studies*, 142, 17–26.
- Greco, J. (2010a). *Achieving knowledge*. Cambridge: Cambridge University Press.
- Greco, J. (2010b). The value problem. In D. Pritchard & S. Bernecker (Eds.), *The Routledge companion to epistemology* (pp. 219–231). London: Routledge.
- Greco, J. (2012). A (different) virtue epistemology. *Philosophy and Phenomenological Research*, 85, 1–26.
- Hart, H., & Honoré, A. (1959). *Causation in the law*. Oxford: Oxford University Press.
- Jarvis, B. (2013). Knowledge, cognitive achievement, and environmental luck. *Pacific Philosophical Quarterly*, 94, 529–551.
- Kallestrup, J., & Pritchard, D. (2016). Dispositional robust virtue epistemology versus anti-luck virtue epistemology. In M. A. Fernández (Ed.), *Performance epistemology*. Oxford: Oxford University Press.

- Kelp, C. (2009). Knowledge and safety. *Journal of Philosophical Research*, 34, 21–31. doi:10.5840/jpr_2009_1.
- Kelp, C. (2011). What's the point of "knowledge" anyway? *Episteme*, 8, 53–66.
- Kelp, C. (2013a). Knowledge: The safe-apt view. *Australasian Journal of Philosophy*, 91, 265–278.
- Kelp, C. (2013b). How to motivate anti-luck virtue epistemology. *Grazer Philosophische Studien*, 88, 211–225.
- Kelp, C. (2014). Epistemology extended. *Philosophical Issues*, 24, 230–252.
- Kelp, C. (2016). Epistemic Frankfurt cases revisited. *American Philosophical Quarterly*, 53, 27–37.
- Lackey, J. (2007). Why we don't deserve credit for everything we know. *Synthese*, 158, 345–361.
- Lackey, J. (2009). Knowledge and credit. *Philosophical Studies*, 142, 27–42.
- Neta, R., & Rohrbaugh, G. (2004). Luminosity and safety of knowledge. *Pacific Philosophical Quarterly*, 85, 396–406.
- Nozick, R. (1981). *Philosophical explanations*. Cambridge: Harvard University Press.
- Owens, D. (1992). *Causes and coincidences*. Cambridge: Cambridge University Press.
- Pritchard, D. (2005). *Epistemic luck*. Oxford: Oxford University Press.
- Pritchard, D. (2007). Anti-luck epistemology. *Synthese*, 15, 277–297.
- Pritchard, D. (2010). Cognitive ability and the extended cognition thesis. *Synthese*, 175, 133–151.
- Pritchard, D. (2012a). Anti-luck virtue epistemology. *Journal of Philosophy*, 109, 247–279.
- Pritchard, D. (2012b). In defence of modest anti-luck epistemology. In K. Becker & T. Black (Eds.), *The sensitivity principle in epistemology*. Cambridge: Cambridge University Press.
- Pritchard, D. (2014). The modal account of luck. *Metaphilosophy*, 45, 594–619.
- Pritchard, D. (2015). Anti-luck epistemology and the Gettier problem. *Philosophical Studies*, 172, 93–111.
- Sainsbury, R. M. (1997). Easy possibilities. *Philosophy and Phenomenological Research*, 57, 907–919.
- Shope, R. K. (1983). *The analysis of knowing: A decade of research*. Princeton: Princeton University Press.
- Sosa, E. (1980). The raft and the pyramid: Coherence versus foundations in the theory of knowledge. *Midwest Studies in Philosophy*, 5, 3–26.
- Sosa, E. (1991). *Knowledge in perspective: Selected essays in epistemology*. Cambridge: Cambridge University Press.
- Sosa, E. (1999). How to defeat opposition to Moore. *Philosophical Perspectives*, 13, 141–154.
- Sosa, E. (2007). *A virtue epistemology*. Oxford: Oxford University Press.
- Sosa, E. (2010). How competence matters in epistemology. *Philosophical Perspectives*, 24, 465–475.
- Sosa, E. (2011). *Knowing full well*. Princeton: Princeton University Press.
- Sosa, E. (2015). *Judgment and agency*. Oxford: Oxford University Press.
- Tienson, J. (1974). On analysing knowledge. *Philosophical Studies*, 25(4), 289–293.
- Turri, J. (2011). Manifest failure: The Gettier problem solved. *Philosophers' Imprint*, 11, 1–11.
- Unger, P. (1968). An analysis of factual knowledge. *Journal of Philosophy*, 65, 157–170.
- Warfield, T. (2005). Knowledge from falsehood. *Philosophical Perspectives*, 19, 405–416.
- Williamson, T. (2000). *Knowledge and its limits*. Oxford: Oxford University Press.
- Williamson, T. (2013). Knowledge first. In M. Steup & J. Turri (Eds.), *Contemporary debates in epistemology*. Hoboken: Blackwell.
- Zagzebski, L. (1994). The inescapability of Gettier problems. *The Philosophical Quarterly*, 44, 65–73.