



Making a difference in virtue epistemology

Fabian Hundertmark¹ · Steven Kindley¹

Received: 3 March 2020 / Accepted: 22 June 2021 / Published online: 3 July 2021
© The Author(s) 2021

Abstract

Virtue Reliabilism holds that knowledge is a cognitive achievement—an epistemic success that is creditable to the cognitive abilities of the knowing subject. Beyond this consensus, there is much disagreement amongst proponents of virtue reliabilism about the conditions under which the credit-relation between an epistemic success and a person’s cognitive abilities holds. This paper aims to establish a new and attractive view of this crucial relation in terms of difference-making. We will argue that the resulting theory, Difference-Making Virtue Epistemology, can deal with cases of epistemic luck and testimonial knowledge while revealing the common core of knowledge and other achievements.

Keywords Knowledge · Virtue epistemology · Achievements · Greco · Explanation · Safety

1 Introduction

Whenever a person’s goal-directed activity succeeds, we may ask why this activity was successful. Accordingly, we might wish to inquire into whether the activity in question was successful due to the person’s abilities or due to sheer luck, and we will appraise said activity accordingly. Imagine a soccer player who scores what seems to the observer to be an awe-inspiring goal, dribbling Pelé-style in the box, avoiding the keeper in order to score. When it later turns out that the player tripped, lost her orientation, twirled around, lost sight of the ball, and dragged the ball with her by sheer luck so that it ended up in the net, this will still be a successful performance. Nevertheless, it is not a case of a success that we will credit to the agent or her abilities. If, on the other hand, the success was indeed due to the player’s abilities, we will be entitled to appraise the person and her successful activity positively

✉ Fabian Hundertmark
fhundertmark@uni-bielefeld.de

Steven Kindley
stevenkindley@gmail.com

¹ Abteilung Philosophie, Universität Bielefeld, Postfach 10 01 31, 33501 Bielefeld, Germany

by saying that “she brought something off” or that she “achieved something.” That is, some cases of success are mere flukes, whereas others are real achievements.

Virtue reliabilism is a theory of knowledge. Its core thesis is that knowledge is an instance of the overarching phenomenon of personal achievement. That is, knowledge is taken to be a particular kind of success, namely an epistemic success, that is due to a particular kind of ability, namely cognitive ability—or, as it is sometimes put, an *intellectual virtue* (Sosa, 1991, 2007, 23f.; 2015, p. 15; Greco, 2010a; Riggs, 2002, 103ff.). In this paper, we are concerned with *strong virtue epistemology* (Pritchard, 2012, p. 265), a version of virtue reliabilism which analyzes knowledge in terms of two related ideas. The first idea is that knowledge is a true belief that has been produced by the exercise of a cognitive ability of the person holding this belief. The second idea is that in cases of knowledge a person has a true belief, rather than a false belief or no belief at all, due to the fact that her belief was produced by her cognitive ability. If these two conditions are met, the person’s epistemic success is creditable to her abilities, and therefore knowledge.

Let us quickly elucidate the notions of epistemic success and cognitive ability. A person is understood to have been epistemically successful if and only if she has a true belief, rather than a relevant false belief.¹ Prime examples of cognitive abilities are, e.g., perception, memory, or various forms of (truth-preserving) inferential reasoning. We take cognitive abilities to be dispositions of a person to form beliefs in ways that are cognitively successful (cf. Greco, 2010a, p. 77; Sosa, 1991, p. 140; 2015, pp. 24–96; see also Jaster, 2020, chap. 5.4 for a similar view).² Importantly, cognitive abilities are understood to be *personal* abilities (Greco, 2010a, p. 77; Sosa, 2015, 26f.). That is why knowledge is understood to be an achievement by virtue reliabilists: The idea is that because knowledge is a kind of epistemic success for the existence of which the abilities *of the subject* are responsible, it is a kind of success for which, by extension, the subject *herself* is responsible (cf. Greco, 2010a, p. 10).

Schematically, the main idea of strong virtue reliabilism can be spelled out as follows, whereby we will call (i) the “causal clause” and (ii) the “because clause”:

S knows that *p* if and only if

- (i) *S*’s true belief that *p* was causally produced by an exercise of a cognitive ability of *S* and

¹ In order to deal with necessary true beliefs and also Kripke’s red barn case (2011, p. 186), it is vital not to restrict epistemic success in a specific situation to the truth of the proposition that a person *actually* believes. Instead, we will understand the notion of epistemic success to be tied up with a person’s having an answer to a relevant question. For arguments to this effect see Grundmann (2020, sec. 2), Hirvelä (2019) and Zhao (2019, secs. 2, 4 and 5).

² Sosa (1991, p. 225) distinguishes generation faculties, like perception, from transmission faculties, like inferential reasoning. While the first generate beliefs from mental states that are not beliefs, the latter generate beliefs from other beliefs. Plausibly, cognitive success for generation faculties consists in epistemic success (the fact that a true rather than a false belief is generated). In contrast, the operation of transmission faculties is cognitively successful if a belief that bears a specific relation to initial beliefs is produced.

- (ii) *S* has formed a true belief, rather than a false belief or no belief at all (concerning some question to which *p* is one possible answer – see fn. 1) *because** *S*'s belief (concerning some question to which *p* is one possible answer) was causally produced by an exercise of a cognitive ability of *S*.

According to the causal clause, a subject's true belief amounts to knowledge only if it was causally produced by the subject's cognitive abilities, and not by some other processes or dispositions that do not count as such. This requirement arguably enables virtue reliabilists to deny that true beliefs resulting from so-called strange and fleeting processes, such as in Plantinga's brain lesion case (1988, 30f.) and Keith Lehrer's case of Mr. TrueTemp (1990, p. 131), are cases of knowledge (Greco, 2010a, chap. 9). Furthermore, this requirement helps virtue reliabilism to deal with cases in which the agent clearly lacks the relevant ability (Greco, 2007, 301f.). There are two good reasons, though, why this requirement by itself is not sufficient for our purposes.

The first reason for the additional requirement on knowledge expressed by the because clause is that it promises to deal with cases of epistemic luck (Nimtz, 2013, 192f.).³ In these cases, a belief is undoubtedly produced by cognitive ability, but it still seems to be a matter of luck that the subject ends up with a true belief. Cases of epistemic luck include everyday cases of luckily true belief, classic Gettier-cases (Greco, 2010a, sec. 5.2; Zagzebski, 1996, sec. 3; Turri, 2011) and also cases of environmental epistemic luck, such as Alvin Goldman's barn-façades case (Goldman, 1976, 772f.).

The second reason for incorporating the because clause is that it promises to deliver an explanation for why knowledge is more valuable than mere true belief, even if this belief was produced by ability (e.g., Riggs, 2002, pp. 87–93). According to virtue reliabilists, this is because knowledge by its very nature is an achievement of the person, whereas mere true belief is not. Insofar as we value personal achievements over mere lucky successes, knowledge is more valuable than mere true belief. This paper is concerned with the question of how exactly the because clause is to be understood in order to deal with what are conceived to be the most problematic cases for virtue reliabilism. We will, for the most part, bracket issues concerning epistemic value. After all, virtue reliabilism's solution to the value problem stems from the general thesis that epistemic normativity is a kind of performance normativity, not from the particular details of the view.

In Sect. 1, we will engage with John Greco's way of spelling out the because clause in terms of explanatory salience. While Greco's approach is subject to a dilemma concerning testimonial knowledge and cases of epistemic luck (Lackey, 2007; Pritchard, 2008, 2012), we will argue that this problem arises due to

³ Nimtz (2013) is right to credit most of the advantages strong virtue reliabilism provides to the because-relation mentioned in the because clause. Appeal to cognitive abilities, however, has at least two connected roles to play. First, it excludes knowledge through dispositions that do not count as abilities (e.g., strange and fleeting processes). Secondly, together with the because-relation, it shows that knowledge is just a special case of a more general class of personal achievements, which in turn is meant to explain the value of knowledge.

Greco's appeal to *explanatory salience* and not to *explanation* per se. In Sect. 2, we will introduce our theory, called "Difference-Making Virtue Epistemology", or "DMVE," for short. According to DMVE the "because" in the because clause indicates an explanatory relation that can be interpreted with the help of the notion of a "difference maker". In Sect. 3, we will show that this interpretation helps DMVE to avoid Lackey's dilemma. In Sect. 4, we will contrast DMVE with Pritchard's anti-luck virtue epistemology. We will argue that DMVE offers a superior solution to cases of epistemic luck while staying true to the idea that knowledge can be analyzed in virtue theoretic (and closely allied) terms alone. Section 5 offers a short conclusion.

2 Greco's explanatory salience requirement

Greco (2010a) takes the "because"-relation to be causal-explanatory. In order to explain Greco's view, let us first take a look at how Greco deals with Gettier cases. Here is a simple one:

(The Stopped Clock) Maria takes a look at a clock the hands of which show that it is 3:30. So, Maria comes to believe that it is 3:30 and her belief is actually true. But unbeknownst to Maria, the clock stopped working precisely 24 hours before.

According to our intuition, Maria does not know that it is 3:30, even though she has the corresponding true belief. Greco argues that in cases like this "*S* believes from an ability and *S* has a true belief, but the fact that *S* believes from an ability does not explain why *S* has a true belief" (Greco, 2010, p. 74, emphasis added). Greco thinks that Maria's belief that it is 3:30 does not amount to knowledge since her perceptual and clock-reading-abilities do not explain why she has acquired this true belief rather than an alternative false one.

What is the rationale for excluding some *causally relevant* factors (e.g., Maria's abilities) from the *explanatory relevant* ones? In order to answer this question, Greco appeals to the idea that "our causal explanations intend to pick out an *important* or *salient* part of the causal story" (Greco, 2010a, 74). A good explanation does not pick out just *any* part of the causal history of an event. Otherwise, the election of Donald Trump in November 2016 could be explained by citing events in the distant past like the big bang or by citing background conditions like the fact that the election took place on a specific day. Instead, a good causal explanation for some event, Greco rightly insists, should pick out the most *salient* factor in its causal history. One explanation for why Trump was elected could, for example, appeal to his strong media presence.

This appeal to explanatory salience gives us the following version of virtue reliabilism (Greco, 2010a, 71ff.):

Knowledge is Success Through Ability (KSA)

S knows that *p* if and only if

- (i) S 's true belief that p was causally produced by an exercise of a cognitive ability of S and
- (ii) the fact that S 's belief (concerning some question to which p is one possible answer) was causally produced by an exercise of a cognitive ability of S is *the most salient factor in a causal explanation* for why S formed a true belief, rather than a false belief or no belief at all (concerning some question to which p is one possible answer).

Greco admits that it is difficult to give a general recipe for how explanatory salience gets determined. To deal with Gettier-cases, however, he appeals to the plausible idea that salience is, amongst other things, affected by *our interests and purposes* on the one hand and by what is *ordinary or usual* on the other. Consequently, Greco asserts: “[O]ur intellectual abilities have a default salience in explanations of true belief. In Gettier cases, this default salience is trumped by something abnormal in the way that S gets a true belief” (Greco, 2010a, 75). Had Maria looked at a well-functioning clock, her cognitive abilities would have been the salient factor for her epistemic success. In the stopped clock case, however, salience shifts to the unusual fact that Maria coincidentally looked at the broken clock at just the right time. Her abilities lose their status as the most salient factor in the explanation of epistemic success. Consequently, KSA yields the result that Maria lacks knowledge in this case.

Unfortunately, KSA is subject to a dilemma worked out by Jennifer Lackey (2007, 2009) that we hold to be detrimental to the account. She argues that KSA, as it stands, is unable to deal with cases of testimonial knowledge and that any weakened version of the salience requirement to the effect that KSA classifies cases of testimonial knowledge correctly, implies that Greco’s account loses its ability to deal with Gettier-cases. Let us start with the first horn of Lackey’s dilemma by considering the following case of testimonial knowledge:

(Chicago Visitor) “Having just arrived at the train station in Chicago, Morris wishes to obtain directions to the Sears Tower. He looks around, approaches the first adult passer-by that he sees, and asks how to get to his desired destination. The passer-by, who happens to be a Chicago resident who knows the city extraordinarily well, provides Morris with impeccable directions to the Sears Tower by telling him that it is located two blocks east of the train station. Morris unhesitatingly forms the corresponding true belief.” (Lackey, 2007, p. 352; 2009, p. 29)

Plausibly, Morris acquires knowledge about the location of the Sears Tower. So, Greco has to say that Morris’ cognitive abilities are the most salient factor in a causal explanation for why Morris acquires a true rather than an alternative false belief. The most salient causal factor in an explanation of why Morris ends up with the belief in question, however, seems to be the presence of a helpful passer-by and *her* cognitive abilities, not the fact that Morris’ belief was produced by his cognitive abilities.

As a response, one may modify the strength of the salience requirement in KSA. This modification brings us to the second horn of the dilemma. Plausibly, any

weakening on the salience-requirement keeps Greco's proposal from being able to deal with Gettier-cases. Let us, for example, suppose that Greco maintains that the fact that a belief was produced by *S*'s cognitive ability needs not to be *the most salient*, but only *one amongst other salient* factors in an explanation of epistemic success. Following this proposal, he could maintain that Morris's cognitive abilities are a salient factor in an explanation of his epistemic success. On the upside, this liberality concerning salience allows KSA to ascribe knowledge to the Chicago visitor. On the downside, however, this move prevents Greco from denying knowledge in cases of epistemic luck. After all, in the stopped clock case Maria's ability to read the clock indeed is *one* salient factor in the explanation of how she acquired a true, rather than a false belief.

Here is our diagnosis of Lackey's dilemma (Greco, 2012, 11f.; see also Gardiner, 2019): The concept of explanatory salience is well-suited to deal with Gettier-cases since only one factor in the causal history of an event can be *the most salient* one in its explanation, given a certain context. On the downside, a subject's cognitive abilities and the cognitive abilities of the helpful informants cannot both be the most salient factor in a given context. Consequently, KSA cannot account for even the most mundane cases of testimonial knowledge, as long as it appeals to *the most salient* factor. If, on the other hand, Greco only requires that the cognitive abilities of the subject are *among the salient explanatory factors*, he loses his grip on cases of epistemic luck.

3 A (difference-making) virtue epistemology

In this section, we will introduce a new and promising interpretation of the because clause in terms of difference-making. Hence, the resulting version of virtue reliabilism will be called "Difference-Making Virtue Epistemology" (DMVE).

An important lesson to draw from the diagnosis in the last section is that the dilemma KSA is confronted with only means trouble for Greco's appeal to explanatory *salience*, not for his suggestion that the proper relation between the exercise of cognitive ability and epistemic success is a *causal explanatory* one. Our proposal sticks to this last idea. Lackey's dilemma does not arise for it since our theory is not based on the concept of salience. Instead, DMVE appeals to the idea that in cases of knowledge the fact that the subject's cognitive abilities were exercised *makes a difference* for the occurrence of the subject's epistemic success.

We claim that a subject's success is creditable to the subject's abilities—that is, it is a personal *achievement*—just in case the fact that it was causally produced by the subject's abilities makes a difference for its being a success. This gives us the following theory of knowledge:

Difference-Making Virtue Epistemology (DMVE)

S knows that *p* if and only if

- (i) *S*'s true belief that *p* was causally produced by an exercise of a cognitive ability of *S* and

- (ii) the fact that S 's belief (concerning some question to which p is one possible answer) was causally produced by an exercise of a cognitive ability of S makes a difference for the fact that S has formed a true belief, rather than a false belief or no belief at all (concerning some question to which p is one possible answer).

Under which conditions does the occurrence of one causal factor make a difference for the occurrence of a specific effect? While we maintain that different answers to this question might do the job, we draw on the notion of *difference-making* by Christian List and Peter Menzies (2009; see also Menzies & List, 2010, 113ff.),⁴ since it highlights the close connection between DVME and the so-called 'safety principle' which plays an important role in the discussion about epistemic luck (see Sect. 4):

Truth conditions for making a difference: The presence of F makes a difference to the presence of G in the actual situation just in case (i) if any relevantly similar possible situation instantiates F , it instantiates G ; and (ii) if any relevantly similar possible situation instantiates $\neg F$ it instantiates $\neg G$. (List & Menzies, 2009, p. 483)

To explain how these conditions are to be understood, it is helpful to take a look at a prominent example from Steven Yablo who asks us to “[imagine] a pigeon, Sophie, conditioned to peck at red to the exclusion of other colors” (Yablo, 1992, 257). Imagine further that Sophie is now confronted with a scarlet triangle and that Sophie pecks. What *explains* the fact that Sophie pecks in this situation? Let us consider three candidates as *explanantia* of our *explanandum*: (1) There was a colored triangle, (2) there was a red triangle, and (3) there was a scarlet triangle. Intuitively, the fact that (1) *there was a colored triangle* does not explain Sophie's pecking very well. After all, the triangle could have easily been colored without Sophie pecking at all, e.g., if the triangle had been green. So, the fact that the triangle is colored is not specific enough. The fact that (3) *there was a scarlet triangle* is also no part of a good explanation for Sophie's pecking. After all, Sophie also would have pecked if the triangle had another shade of red, e.g., crimson. The fact that the triangle is scarlet is too specific. In contrast, fact (2) seems to be what we are looking for. That is because the explanation referring to the fact that (2) *there was a red triangle* gives us information about a factor whose presence really makes a difference to whether Sophie is pecking or not.

According to the above necessary condition for difference-making, the fact that (1) *there was a colored triangle* does not make a difference for Sophie's pecking since there are relevantly similar possible situations in which the triangle is colored (e.g., by being green) but Sophie fails to peck (condition (i) is not fulfilled). Similarly, the fact that (3) *there was a scarlet triangle* also does not make a difference to Sophie's pecking. After all there are relevantly similar possible situations in which

⁴ In contrast to List and Menzies, we think of difference-making as a feature of causal explanation rather than causation per se (cf. Marras and Yli-Vakkuri 2010, 134f.).

the triangle is not scarlet (e.g., by being crimson) but Sophie pecks as well (condition (ii) is not fulfilled). In contrast, the fact that (2) *there was a red triangle* fulfills both conditions for difference-making: (i) if in any relevantly similar possible situation there had been a red triangle, Sophie would have pecked, and (ii) if in any relevantly similar possible situation there had not been a red triangle, Sophie would not have pecked.

In order to apply the difference-making criterion to DMVE, we set the variables of the definition by List and Menzies, cited above, as follows:

F = S 's belief (concerning some question to which p is one possible answer) being causally produced by an exercise of a cognitive ability of S .

G = S forming a true belief, rather than a false belief or no belief at all (concerning some question to which p is one possible answer).

We think that DMVE provides us with necessary and sufficient conditions for knowledge. Consequently, a true belief amounts to knowledge iff, in addition to the causal clause, the because clause, understood in terms of difference-making, is fulfilled. As it turns out, though, condition (i) of making a difference suffices for the present purposes of showing that DMVE is superior to its alternatives. After all, the causal clause is part of every virtue reliabilist theory, and the second difference-making clause is almost trivially fulfilled in cases of knowledge. It says, roughly, that the belief in question would have been non-existent or false if it hadn't been caused by an exercise of cognitive ability and the situation is relevantly similar. So, to keep things simple, we will focus on the first necessary condition for making a difference. According to DMVE, this gives us the following necessary condition for S knowing that p :

if in any relevantly similar possible situation S 's belief (concerning some question to which p is one possible answer) is produced by an exercise a cognitive ability of S , S forms a true belief (concerning some question to which p is one possible answer), rather than a false belief or no belief at all.

Please note that we came to this requirement simply by asking what makes for a good explanation. Only in a second step, we defend this natural way of understanding the because clause cast in explanatory terms by showing that it helps virtue reliabilism to deal with a variety of problems posed by puzzle-cases.

A central question for assessing DMVE is which possible situations count as *relevantly similar*. A complete and detailed answer to this question would go beyond the scope of this paper. However, there is an emerging consensus that causal reasoning is crucially influenced by assumptions concerning descriptive and prescriptive normality (Halpern & Hitchcock, 2013, sec. 5; Halpern, 2016, sec. 3.1; Menzies, 2007, sec. 8.7; Henne et al., 2017 among others). Consequently, we hold that the same factors play a role in causal explanation and difference-making by making possible situations relevant.

First, which factors play roles in causal explanations is partly determined by what is descriptively normal. Here is an example. A causal explanation of why Liang's indoor plants withered can refer to the fact that Liang did not water them,

but not to the fact that some stranger did not break into Liang's house to water them. On the difference-maker account of causal explanation, this difference has to be captured by the fact that the possible situation in which Liang waters his plants is a relevantly similar one, while a situation where some stranger breaks into his house to do so is not. The difference between both situations lies in the fact that it is normal for persons to care for their plants, while it is not normal for a stranger to do so.

Descriptive normality amounts to defaultness, typicality and statistical normality (Halpern & Hitchcock, 2013, sec. 5; Halpern, 2016, sec. 3.1). If we are confronted with some actual situation, we can judge whether a close possible situation is descriptively normal by asking whether things behave in a default, typical, or statistically normal way. Defaultness is determined by assumptions about how things behave if no further information is given. We might, for example, have the default assumption that a dog is aggressive. In contrast, it is certainly not typical for dogs to be aggressive. We may answer questions about typicality (e.g., whether it is typical for mosquitoes to carry malaria) by considering the question of whether the corresponding generic generalization ("mosquitoes carry malaria") is true. Typicality has to be distinguished from statistical normality, which has to do with something being (close to) the statistical mode or mean. While it may be typical for mosquitoes to carry malaria, it is certainly not statistically normal since only one percent of mosquitoes carry the virus.

Second, the relevance of close possible situations is also determined by what is prescriptively normal. Let us consider the following situation:

(Huā County) Jiang lives in a Huā County. In Huā County it is taken for granted that people have to rescue plants in danger, just as it is customary in other parts of the world for humans and other animals. As every citizen of Huā County, Jiang has houseplants. Unfortunately, Jiang's plants withered.

What explains the fact that Jiang's plants withered? Even though the cases of Liang and Jiang are very similar, our intuition seems to differ. While we are happy to explain the withering of Liang's plants by appealing to the fact that he forgot to water them, the same explanation does not seem to suffice in the case of Jiang. After all, possible situations in which some other inhabitant of Huā County watered Jiang's plants have to be considered as well. Consequently, the norms in Huā County make certain possible situations salient and thereby relevant for applying the difference-making criterion. So, the fact that Jiang forgot to water his plants does not qualify as a difference-maker for their withering.

Concerning prescriptive normality, we can offer the following rough guide: If the actual situation takes place in a certain social setting (e.g., the community to which Liang belongs or Huā County) or refers to a social practice (e.g., playing chess), possible situations made salient by this setting or practice are prescriptively normal. In social settings, possible situations in which persons obey the laws, policies, or moral norms are salient. In Huā County, for example, it is a moral norm to rescue plants even though one has to break into other people's houses to do it. Consequently, the possible situation in which a person breaks into Jiang's house to water

his plants is prescriptively normal. For a given social practice, a situation is salient if it is the kind of situation that enables this practice to persist or proliferate.⁵ The practice of playing chess, for example, persists only because the players follow the rules of chess.

4 Difference-making virtue epistemology avoids Lackey's dilemma

In this section, we will show how DMVE deals with Lackey's dilemma. To this effect, we will first show that DMVE is able to handle Gettier-cases. Afterwards, we argue that DMVE also delivers plausible verdicts in cases of testimonial knowledge. First, we will discuss the following Gettier-case by Roderick Chisholm:

(Sheep in the Field) "A person takes there to be a sheep in the field [...]. The person, however, has mistaken a dog for a sheep and so what he sees is not a sheep at all. Nevertheless, it happens that there is a sheep in another part of the field." (Chisholm, 1989, p. 93)

In this case, the person, call her "Shaquana," does not know that there is a sheep in the field. DMVE warrants this result. Shaquana's (perceptual and sheep-identifying) abilities do not make a difference for her epistemic success. After all, there are similar possible situations in which Shaquana exercises her abilities but ends up with the false belief that a sheep is in the field, i.e., the situation in which there is no sheep hidden out of view. Furthermore, these situations are relevant since it is a default assumption and thereby descriptively normal that there is a dog present without a sheep being present in this very field.⁶ Consequently, DMVE yields the desired result that Shaquana lacks knowledge in this case.

For DMVE to deal with Lackey's dilemma, it also has to be able to classify cases of testimonial knowledge correctly. So, let us again take a look at the case of the Chicago visitor, Morris, who learns about the direction to the Sears tower by asking some passer-by for directions (Sect. 2). Unquestionably, Morris' true belief that the Sears tower is two blocks east of the train station was produced by his cognitive abilities. Consequently, the question of whether it amounts to knowledge depends on the question of whether the fact that it was produced in this way makes a difference for Morris' epistemic success. There is no reason to believe that the second condition on making a difference is not fulfilled. So, the question becomes whether the following first condition is: if in any relevantly similar possible situation Morris' belief about the location of the Sear's tower is produced by his cognitive abilities, it is true.

Let us assume that Morris has normal communicative abilities and is, therefore, able to ask the passer-by the right question, understand the given answer and

⁵ As we will see in Sect., this dependence of difference-making on prescriptive normality plays an important role when it comes to lottery cases.

⁶ Arguably, the same solution works for cases in which some epistemic bad luck is combined with an unusual environment that facilitates success such as the famous Temp-case (Pritchard 2012, p. 260) or the Rock and Hard Place-case by Miracchi (2015).

generate a corresponding belief. The exercise of these abilities qualifies as a difference-maker for Morris' epistemic success just as long as the passer-by provides him with truthful information about the location of the Sears Tower in every relevantly similar possible situation. Since the relevance of similar situations depends on them being normal, the question of whether the exercise of Morris' communicative abilities alone makes a difference for his epistemic success depends on whether untruthful passers-by are descriptively or prescriptively normal in Morris' situation. We can safely assume that Morris' informant is competent and honest in all prescriptively normal situations. After all, the whole practice of asking for directions depends on truthful information being given and honesty is also a moral norm. Whether false information given by a passer-by are descriptively normal, however, depends on Morris' environment. Consequently, we agree with Christoph Kelp (2011, p. 431) who argues that an epistemically friendly environment only demands basic communicative abilities from Morris for testimonial knowledge, while an environment with a significant amount of pathological liars, confused or directionally-challenged people additionally requires Morris to exercise his ability to discern helpful from unhelpful advice-givers.

As we have seen, DMVE can avoid Lackey's dilemma between Gettier cases and testimonial knowledge. In cases of testimony, the status of the subjects' abilities as a difference maker crucially depends on whether situations in which untruthful information is given are descriptively normal and on the subjects' ability to cope with these situations, e.g., by identifying reliable (viz. unreliable) informants. In Gettier-cases the subject's abilities' status as a difference maker for epistemic success notably is undermined by the fact that possible situations in which the subject acquires a false rather than a true belief by exercising the same abilities are descriptively or prescriptively normal and thereby relevantly similar. Consequently, the subjects in question can be said not to have knowledge.

5 Difference-making virtue epistemology and its anti-luck condition

Duncan Pritchard thinks that, in order to deal with cases of epistemic luck, we need to replace the because clause in virtue reliabilism with an independently motivated anti-luck condition, and thereby relinquish what we take to be one of the core ideas of the theory—i.e., that knowledge is definable as a specific kind of personal achievement. We will argue that Pritchard's theory is subject to the Greco's (2007, 299ff.) safety dilemma, while DMVE provides an adequate verdict in cases of epistemic luck. Furthermore, we will show that DVME is more parsimonious than Pritchard's anti-luck virtue epistemology and also has the advantage of offering a unificatory treatment of knowledge and other achievements. Consequently, DMVE is to be preferred to Pritchard's account.

Pritchard's theory of knowledge aims to capture two independent intuitions. First, the intuition that knowledge is a product of one's cognitive abilities. Second, the intuition that knowledge is incompatible with luck. Pritchard accepts that the first intuition is captured by the causal clause of strong virtue reliabilism. He denies, however, that there has to be some particular relation between the fact that the belief

was produced by ability and it being a true, rather than a false belief. Instead, he supplements the causal clause with a safety condition which is solely designed to capture the intuition that knowledge is incompatible with luck (see Pritchard, 2005, chap. 5.1 for a general discussion of luck). Thereby he reaches a view he calls anti-luck virtue epistemology (Pritchard, 2012, 272ff.):

Anti-luck Virtue Epistemology (ALVE)

S knows that p if and only if

- (i) S 's true belief that p was causally produced by an exercise of a cognitive ability of S and
- (ii) S 's belief that p is safe.

According to Pritchard (2012, 253), the safety clause (ii) can be understood along the lines of the following *safety principle*:

If S knows that p then S 's true belief that p could not have easily been false.

Unfortunately, however, Pritchard's safety principle is subject to a dilemma raised by John Greco (2007, 299ff.; see also Sosa, 2015, p. 119). This dilemma takes issue with the question of when a belief could easily have been false. Greco gives us two interpretations of Pritchard's safety principle:

Strong Safety: In close worlds, always if S believes p then p is true. [...]

Weak Safety: In close worlds, usually if S believes p then p is true. (Greco, 2007, 301)

Greco gives several counterexamples to show that strong safety is too strict. The following case, however, is from Ernest Sosa. He points to the fact that we can gain knowledge by calling a technical support hotline despite the fact that one in a million operators would have given us false information in a close possible world where she had been randomly assigned to us (bad Apple case; Sosa, 2015, sec. 119). One way to deal with these kinds of cases is to interpret the safety principle along the lines of weak safety. But this interpretation is in conflict with the fact that a lottery player simply does not know that she will lose, even though the belief that she does, will be true in the overwhelming majority of close possible worlds (lottery case; Greco, 2007, p. 300; Sosa, 2015, p. 119). ALVE cannot differentiate between both cases. After all, maintaining that the hotline caller knows the right solution to his or her problem while denying that the lottery player knows that she is about to lose must somehow be grounded in a difference concerning the closeness of worlds in which the subject has a false belief. But there is no reason to pose such a difference in the first place.

In contrast, DMVE has the resources to treat both cases differently. The reason is that a situation is relevant for difference-making not in virtue of its being close to the actual world, but in virtue of being descriptively normal or prescriptively normal. Normal worlds do not have to be close and vice versa. While a situation in which the agent acquires a false believe is not descriptively normal (default, typical or statistically normal) in the bad apple case as well as in the lottery case, there is a crucial

difference concerning prescriptive normality. If a player buys a lottery ticket, the situation in which this ticket wins is prescriptively normal. After all, the social practice of playing the lottery solely exists because people win sometimes. Consequently, the situation that the bought ticket wins is relevant and DMVE can deny that the player knows that she will lose. Contrast this with the hotline case. That service employees do give bad advice is not demanded by legal norms, moral norms or policies nor does any social practice require this. The situation in which a service employee provides false information is not prescriptively normal. Consequently, DMVE has no difficulty to ascribe knowledge in this case.

DMVE copes with Greco's dilemma since difference-making explanation is naturally sensitive to prescriptive normality. Of course, Pritchard could also achieve this result by incorporating prescriptive normality into his safety principle. Given the main idea behind this principle, however, he has good reason not to. As we have seen, the safety principle is designed as an anti-luck condition, and luck is not sensitive to prescriptive normality. To see this, consider the question of whether the person in the lottery case is luckier in her true belief than the caller of the support hotline in his true belief about the solution to the technical problem. Provided that a win in the lottery has the same probability as bad advice from the support hotline, there seems to be no difference concerning luck. Consequently, one would significantly deviate from the safety principle's main idea by incorporating prescriptive normality into it. But, let us, for the sake of argument, assume that the safety principle can be modified in a way so that ALVE and DMVE are both on par concerning their ability to classify cases of knowledge as such. We will now argue that DMVE should nevertheless be preferred on grounds of parsimony and unification.

The safety condition on knowledge is designed to deal with cases of knowledge-undermining luck. Even though a well-worked out version of it may help to classify these cases correctly, it is, at least from a purely virtue theoretic perspective, an ad hoc device just made for that purpose. Of course, this observation does not serve as a knockdown argument against the use of this principle. But all things being equal, a virtue theoretic theory without an ad hoc condition should be preferred on grounds of parsimony. DMVE is such a theory as it does not add a foreign element to the framework of strong virtue reliabilism and consequently takes the incompatibility of knowledge with certain kinds of luck just to be an outcome of knowledge being an achievement (this ad hocness charge against ALVE has also been made by Greco, 2010b, 227). Furthermore, DMVE delivers the intuitively plausible result that knowledge is, by its very nature, a kind of achievement and, therefore, more valuable than mere true belief. In contrast, ALVE cannot explain this intuition concerning the value of knowledge. It allows the virtue-theoretic component (clause i) to be satisfied without the resulting belief being safe (clause ii) and thus amounting to knowledge.

This brings us directly to the second reason for preferring DMVE to ALVE, even if both were to classify cases in the same way. DMVE does not only offer a theory of knowledge but also a way of understanding achievements more generally. According to DMVE, every achievement is a case of success for which the exercise of the subject's abilities makes a difference. This explains why all cases of achievements are, to a certain extent, incompatible with luck. ALVE, in contrast, is only a theory

tailor-made for knowledge without any commitments beyond. Consequently, DMVE is to be preferred to ALVE on grounds of unification.

But does DMVE really offer a plausible theory of achievements in general? As presently developed, DMVE implies that success is only an achievement if this success is guaranteed by the subject's use of her abilities. This consequence can be turned into an argument against DMVE (Carlotta Pavese, pers. comm.; see also Sosa, 2007, 28f.). After all, achievements are compatible with the fact that many possible attempts would be unsuccessful. In the case where a skilled archer hits the bulls-eye, there could have been a sudden gust of wind that blows the arrow off course. Similarly, a home run by an MLB player will be an achievement, even if most attempts of this player will fail to be home runs. As these examples show, achievements are compatible with the fact that success is not solely in the trying subject's hands. But since DMVE excludes this possibility, one might argue that it advances not only the claim that knowledge is an epistemic success due to cognitive ability but something more specific.

In response to this objection, we admit that knowledge does not exclude this kind of luck just in virtue of being an achievement. Especially the examples from sports show that there are achievements not wholly in the hands of the subject. We will call these achievements "weak achievements" and contrast them with "strong achievements".⁷ The difference between both lies in the subject's degree of control over her success (see Carter 2016, sec. 5 for a similar point). A goal scored by a skilled soccer player in a penalty kick is typically a weak achievement. Even the balls shot by world-class players will find the net only 80 percent of the time. In contrast, a shot by a skilled player two meters from the goal and without any obstacles will be a strong achievement. A skilled player who tries to score under these circumstances will almost certainly do so. The success of the player will be in her hands so that the factor of luck will be reduced to a minimum. In our view, cases of knowledge are strong achievements and, therefore, incompatible with luck in a strong sense.

The difference between weak achievements and strong achievements can be captured in terms of difference-making and, therefore, does not pose a problem to DMVE. In our view, a weak achievement is a success that is probabilistically explained by the agents' abilities, so that the exercise of the ability raises the probability of success (see Hirvelä, 2018, sec. 2 for a similar thought). Every account of difference-making has to handle probabilistic explanations. Let us, for example, modify the case of the pigeon above. The pigeon Alberto is conditioned to peck exclusively at red objects. In contrast to Sophie, however, he does not peck at every red object (e.g., due to a mild form of color blindness). Imagine that Alberto is confronted with a red triangle, and he pecks. Just as in Sophie's case, the best explanation of why Alberto pecks is the presence of a red triangle. In contrast to the case of Sophie, however, the truth conditions for making a difference are not fulfilled. After all, there are relevantly similar possible situations where there had been a

⁷ This distinction is similar to, yet different from, Carter and Duncan's (2017, 181f.) distinction between strong and weak cognitive achievements insofar as a weak cognitive achievement in their sense can be a strong achievement in our sense.

red triangle, but Alberto would not have pecked. Consequently, List and Menzies account of difference-making must be modified to allow for such probabilistic cases (as they acknowledge in 2009, p. 478).

We will not commit to any specific idea of how this modification has to look like. But a straightforward extension of the account above demands that the presence of *F* raises the probability of the presence of *G* in every relevantly similar *F*-situation in cases of probabilistic difference-making. The presence of a red triangle, for example, raises the probability that Alberto pecks in every relevantly similar possible situation, even if Alberto is not guaranteed to peck at every red triangle. Similarly, the fact that a player exercises her soccer abilities at a penalty shot will make a difference to her success because the exercise of these abilities raises the probability of a goal in every relevantly similar possible situation.

The fact that there are probabilistic and non-probabilistic difference-makers allows us to distinguish between weak and strong achievements. While weak achievements are compatible with a certain degree of luck, strong achievements are much less so. By maintaining that knowledge is a strong achievement, DMVE can deal with cases of epistemic luck. Furthermore, it does so without departing from the basic tenet of virtue reliabilism, in the sense that the nature of knowledge is primarily explained with recourse to properties of persons and considerations that fit snugly into this picture. As we have seen, this gives DMVE an advantage over ALVE since it offers a more parsimonious account that explains our intuitions concerning the value of knowledge and can be seen as part and parcel of a unificatory account of achievements in general.

6 Conclusion

According to Difference-Making Virtue Epistemology, a subject's true belief amounts to knowledge when the subject's abilities make a difference to her epistemic success. As we have shown, this theory captures the intuition that knowledge is incompatible with certain kinds of epistemic luck in an encompassing and economical way. This is because DMVE is well able to account for the fact that subjects in Gettier cases do not have knowledge and gives a clear and correct verdict in cases of environmental epistemic luck as well as testimonial knowledge. In contrast to Pritchard's ALVE, DMVE has this feature just in virtue of specifying the core idea behind virtue reliabilism. Consequently, DMVE is an auspicious way to spell out the basic tenet of strong virtue epistemology.⁸

⁸ Recently Beddor and Pavese (2020) have developed and defended Modal Virtue Epistemology, which can be summarized as follows:

Modal Virtue Epistemology (MVE) S knows that p if and only if in all relevantly close worlds, where S engages in a sufficiently similar performance, S's belief is true. (Beddor and Pavese 2020).

Despite their differences, the similarity between DMVE and MVE is striking, especially since Beddor and Pavese propose to spell the worlds relevant for MVE partly in terms of their normality and not only in terms of their similarity to the actual world. For the purposes of this paper, we take these similarities to be further evidence that DMVE is on the right track. Whether DMVE or MVE is the superior theory has to be decided elsewhere.

Acknowledgements We would like to thank Jaakko Hirvelä, Christian Nimtz, Carlotta Pavese, Peter Schulte, Rose Trappes and the audiences in Bielefeld, Cologne and Helsinki for valuable comments on previous versions of this paper.

Funding Open Access funding enabled and organized by Projekt DEAL.

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

References

- Beddor, B., & Pavese, C. (2020). Modal virtue epistemology. *Philosophy and Phenomenological Research*, 101(1), 61–79.
- Carter, J. A., & Duncan, P. (2017). Epistemic situationism, epistemic dependence, and the epistemology of education. In *Epistemic situationism* (pp. 168–191). Oxford University Press.
- Chisholm, R. M. (1989). *Theory of knowledge* (3rd ed.). Prentice-Hall.
- Gardiner, G. (2019). Virtue epistemology and explanatory salience. In H. Battaly (Ed.), *Routledge handbook of virtue epistemology* (pp. 296–308). Routledge.
- Goldman, A. I. (1976). Discrimination and perceptual knowledge. *Journal of Philosophy*, 73(November), 771–791.
- Greco, J. (2007). Worries about Pritchard's Safety. *Synthese*, 158(3), 299–302.
- Greco, J. (2010a). *Achieving knowledge: A virtue-theoretic account of epistemic normativity*. Cambridge University Press.
- Greco, J. (2010b). The value problem. In S. Bernecker & D. Pritchard (Eds.), *The Routledge companion to epistemology* (pp. 220–231). Routledge.
- Greco, J. (2012). A (Different) Virtue Epistemology. *Philosophy and Phenomenological Research*, 85(1): 1–26.
- Grundmann, T. (2020). Saving safety from counterexamples. *Synthese*, 197, 5161–5185.
- Halpern, J. Y. (2016). *Actual causality*. The MIT Press.
- Halpern, J. Y., & Hitchcock, C. (2013). Graded causation and defaults. *British Journal for the Philosophy of Science*, 66(2), 413–457.
- Henne, P., Pinillos, Á., & De Brigard, F. (2017). Cause by Omission and Norm: Not watering plants. *Australasian Journal of Philosophy*, 95(2), 270–283.
- Hirvelä, J. (2018). On virtue, credit and safety. *Grazer Philosophische Studien*, 95(1), 98–120.
- Hirvelä, J. (2019). Global safety: How to deal with necessary truths. *Synthese*, 196(3), 1167–1186.
- Jaster, R. (2020). *Agents' abilities*. De Gruyter.
- Kelp, C. (2011). In defence of virtue epistemology. *Synthese*, 179(3), 409–433.
- Kripke, S. A. (2011). Nozick on knowledge. In S. A. Kripke (Ed.), *Philosophical troubles. Collected papers* (Vol. I, pp. 262–224). Oxford University Press.
- Lackey, J. (2007). Why we don't deserve credit for everything we know. *Synthese*, 158(3), 345–361.
- Lackey, J. (2009). Knowledge and credit. *Philosophical Studies*, 142(1), 27–42.
- Lehrer, K. (1990). *Theory of Knowledge. Dimensions of Philosophy Series*. Westview Press.
- List, C., & Menzies, P. (2009). Nonreductive physicalism and the limits of the exclusion principle. *The Journal of Philosophy*, 106(9), 475–502.
- Marras, A., & Juhani, Y.-V. (2010). Causal and explanatory autonomy: Comments on Menzies and List. *Emergence in mind* (pp. 129–138). Oxford University Press.
- Menzies, P. (2007). Causation in context. In H. Price & R. Corry (Eds.), *Causation, physics, and the constitution of reality: Russell's republic revisited*. Oxford University Press.

- Menzies, P., & Christian, L. (2010). The causal autonomy of the special sciences. In C. McDonald & G. McDonald (Eds.), *Emergence in mind*. Oxford University Press.
- Miracchi, L. (2015). Competence to know. *Philosophical Studies*, 172(1), 29–56.
- Nimtz, C. (2013). Knowledge, abilities, and ‘because’-clauses. A critical appraisal of virtue-theoretic analyses of knowledge. In T. Henning & D. P. Schweikard (Eds.), *Knowledge, virtue, and action. Essays on putting epistemic virtues to work*. Routledge.
- Plantinga, A. (1988). Positive epistemic status and proper function. *Philosophical Perspectives*, 2, 1–50.
- Pritchard, D. (2005). *Epistemic luck*. Oxford University Press.
- Pritchard, D. (2008). Virtue epistemology and epistemic luck revisited. *Metaphilosophy*, 39(1), 66–88.
- Pritchard, D. (2012). Anti-luck virtue epistemology. *The Journal of Philosophy*, 109(3), 247–279.
- Riggs, W. D. (2002). Reliability and the value of knowledge. *Philosophy and Phenomenological Research*, 64(1), 79–96.
- Sosa, E. (1991). *Knowledge in perspective—Selected essays in epistemology*. Cambridge University Press.
- Sosa, E. (2007). *A virtue epistemology: Apt belief and reflective knowledge* (Vol. 1). Oxford University Press.
- Sosa, E. (2015). *Judgment & agency*. Oxford University Press.
- Turri, J. (2011). Manifest failure: The gettier problem solved. *Philosophers’ Imprint*, 11(8), 1–11.
- Yablo, S. (1992). Mental causation. *Philosophical Review*, 101(2), 245–280.
- Zagzebski, L. T. (1996). *Virtues of the mind: An inquiry into the nature of virtue and the ethical foundations of knowledge*. Cambridge University Press.
- Zhao, H. (2019). Better virtuous than safe. *Synthese*. <https://doi.org/10.1007/s11229-019-02501-5>

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