

Difficulty and Knowledge

Fernando Broncano-Berrocal

(Autonomous University of Madrid)

Abstract: Some things are more difficult to know than others. For example, proving the Poincaré conjecture is certainly more difficult than coming to know what does $2 + 2$ equal. However, as obvious as it seems, explaining that knowledge can be difficult in familiar epistemological terms (e.g., in evidentialist or simple reliabilist terms) is less straightforward than one could initially think. The aim of the paper is to show that virtue reliabilism (unlike virtue responsibilism) provides a promising framework for accounting for the relationship between difficulty and knowledge. However, it argues that virtue reliabilism first needs to get rid of the problematic assumption that cognitive abilities are reliable dispositions to form true beliefs in appropriate conditions. The reason is that this idea not only prevents the theory from explaining how knowledge relates to difficulty, but also renders its main tenet—the thesis that knowledge requires the manifestation of cognitive ability in appropriate conditions—false. To amend this problem, the paper connects the virtue reliabilist framework with recent work on the notions of achievement and difficulty. It then advances a positive proposal: the view that knowledge is a special kind of challenge with varying degrees of difficulty.

Keywords: difficulty; knowledge; ability; achievement; effort; virtue reliabilism; virtue responsibilism.

1 Difficulty in the theory of knowledge

Some things are more *difficult to know* than others. It is more difficult to come to know how many grains of sand are there in a dune than in a heap. Proving the Poincaré conjecture is certainly more difficult than coming to know what does $2 + 2$ equal. The larger the distance, the more difficult to recognize a facial expression is. The question that this paper will address is this: *How can the relationship between difficulty and knowledge be accounted for in epistemology?*

As obvious as it seems, explaining that knowledge can be difficult in familiar epistemological terms is less straightforward than one could initially think. Suppose with *evidentialism* that knowledge is a function of one's evidence. What is the difference in difficulty between coming to know how many sand grains are there in a dune and in a heap? One idea (in line with evidentialism) is that the more evidence is needed to know something, the more difficult to know it is. In order to know how many sand grains are there in a dune one certainly needs to gather more evidence, which explains the difference

in difficulty. But if the amount of evidence is the key factor that accounts for the relationship between difficulty and knowledge in general, what explains that solving a complex Sudoku puzzle is more difficult than solving a simpler one? Here the difference does not seem to lie in the quantity or quality of one's evidence—the evidence is roughly the same in both cases—, but in the *skill* required to find a solution.

Suppose now with *reliabilism* that knowledge is a function of the reliability of one's belief-forming processes. One idea (in line with reliabilism) is that the less reliable an agent in some domain of inquiry is, the more difficult it is for her to know propositions in that domain. Accordingly, the difference in difficulty between solving a complex Sudoku puzzle and a simpler one can be explained in terms of the fact that one needs to be more reliable in solving Sudoku puzzles to solve the former. However, while appealing to the notion of reliability may help account for cases in which the difference in difficulty is down to a difference in skill, it can hardly explain cases in which it is down to a difference in *effort*. For example, imagine a very rudimentary calculator with only three buttons '1', '+' and '=', which is designed to help kids learn natural numbers. There is no other way to add numbers but to press '1' and '+' as many times as needed and then press '=' to get a result. In addition, suppose that the calculator, being designed for kids, automatically corrects syntactic errors that may be produced in the course of typing, e.g., long strings of symbols. In such a device, calculating a large natural number is arguably *more difficult* than calculating a small one. However, this difference in difficulty neither lies in the reliability of the belief-forming method, nor in the skill needed. First, the calculator is as reliable when the result is a big number (e.g., 11,102,016) as when the output is a small one (e.g., 11). Second, since the calculator automatically corrects eventual syntactic errors, the only skill needed is that one is able to type '1' and '+' as many times as one wishes and that, at some point, one presses '=' to know the result. In this way, calculating 11,102,016 is more difficulty than 11 simply because it takes *more time and effort* to introduce the relevant values. The difference in difficulty is not a difference in reliability or skill.

These considerations do not show that the relationship between difficulty and knowledge cannot be accounted for in evidentialist or simple reliabilist terms, but serve to illustrate that explaining this relationship in familiar epistemological terms (such as evidence or reliability) is less straightforward than one could initially think.

Enter *virtue epistemology*. Virtue epistemology provides a promising theoretical framework for explaining and elaborating on the idea that knowledge can be difficult. The main reason is that it gives a central role to the *notion of skill* in the explanation of the nature of knowledge and, as we have seen, this notion seems to play a central role in the

understanding of the concept of difficulty. The *aim of the paper* is to explore these conceptual connections and account for the relationship between difficulty and knowledge in virtue epistemological terms. In so doing, I will make two relevant points about virtue epistemology. The first point will be negative: current views make certain assumptions concerning the role that situations play in the acquisition of knowledge that not only prevent them from explaining the relationship between difficulty and knowledge, but also lead them to false claims about the very nature of knowledge. The second point will be positive: the way to amend this problem is to adopt the idea that knowing can be difficult in the way I will suggest. In this manner, the paper aims to give insights both into the relationship between difficulty and knowledge and into the (virtue epistemological) nature of knowledge.

The paper is structured as follows. In §2, I will introduce the two main branches of virtue epistemology and the main assumptions that each make concerning the role that situations play in the acquisition of knowledge. The main point will be that *character-based virtue epistemology*, also known as *virtue responsibilism*, does not provide a suitable theoretical framework for explaining the idea that knowledge can be difficult; by contrast, *competence-based virtue epistemology*, also known as *virtue reliabilism*, is *prima facie* compatible with it. However, as I will argue in §3, a key assumption of virtue reliabilism—namely, the assumption that cognitive abilities are reliable dispositions to form true beliefs *in appropriate conditions*—not only prevents the theory from accounting for the relationship between difficulty and knowledge, but also renders its main tenet—the thesis that knowledge requires the manifestation of cognitive ability in appropriate conditions—false. In §4, I will connect the virtue reliabilist framework with recent work on the notions of achievement and difficulty. The aim will be to shed some light on the idea that knowledge can be difficult as well as on the virtue reliabilist thesis that knowledge is a cognitive achievement. Section §5 will follow up on the ideas introduced in §4 and will develop the idea that acquiring knowledge is a special kind of challenge with varying degrees of difficulty.

2 The situational assumptions of virtue epistemology

Virtue epistemology offers a simple answer to the question of what the *nature of knowledge* is: knowing that *p* is a matter of correctly believing *p* because of an exercise of intellectual virtue. Yet simple as it looks, there is some controversy about how to unpack this answer and especially about how to understand the relevant intellectual virtues.

Some virtue epistemologists think that the relevant intellectual virtues are stable reasons-responsive character traits that motivate or orient individuals towards the

attainment of epistemic goods, including knowledge.¹ Open-mindedness, curiosity, creativity, intellectual courage and intellectual perseverance are some examples. Crucially, just as character traits in general are supposed to produce *cross-situationally consistent behavior independently of normatively irrelevant situational factors*, the cognitive performances of agents possessing character-based intellectual virtues are supposed to be consistent in a great variety of situations and independently of such factors.² For example, the courageous person is supposed to exhibit courage in a wide range of dangerous or challenging situations, and more importantly, independently of, e.g., the presence of bystanders. In the same vein, to exhibit intellectual courage, one must consistently challenge widespread opinions regardless of bystanders. Likewise, to count as intellectually perseverant, one must carry on one's inquiry or pursue one's intellectual project whenever one's beliefs are contested, independently of whether one is tired, sad, or one's favorite show is live on TV.

To be clear, the claim is *not* that character-based intellectual virtues are situation-independent: one is not expected to be intellectually courageous or perseverant while sleeping or dead. Instead, the claim is that they are relativized to a *wide variety* of environments, in that intellectually virtuous agents should manifest them in all of them. In addition, if they are supposed to be independent of anything, it is of normatively irrelevant situational influences (e.g., the presence of bystanders, energy or mood elevators, and so on). These two features (cross-situational consistency across a wide range of environments and independence of normatively irrelevant factors) constitute what we may call *the situational assumption of character-based virtue epistemology* (also known as *virtue responsibilism*).

Other virtue epistemologists theorize that the relevant intellectual virtues are based on cognitive faculties such as perception, reasoning or memory. More specifically, they conceive them as stable cognitive dispositions—as opposed to motivational or conative character traits—to form true beliefs reliably. So understood, intellectual virtues are particular kinds of competences, skills or abilities (e.g., visual, inferential, memorial) whose goal is to form true rather than false beliefs—competences, skills or abilities in general are defined as stable dispositions to achieve goals reliably.³ Importantly, an assumption that epistemologists upholding this conception of intellectual virtues make is that abilities, in general, and cognitive abilities, in particular, are reliable dispositions to achieve goals in *appropriate conditions*.

¹ See Baehr (2011) for extensive discussion on this approach to intellectual virtues.

² See Alfano (2013) for relevant discussion on the cross-situational component of character traits and corresponding virtues.

³ See Greco (2009) and Sosa (2010) for helpful discussion on this way of understanding intellectual virtues. Throughout the paper, I will use the terms 'competence', 'skill' and 'ability' interchangeably.

*What are the appropriate conditions of an ability, on the reliabilist model?*⁴ One idea is that the appropriate conditions of an agent's ability to achieve a certain kind of goal are the *type* of conditions under which the agent, by exercising her ability, can achieve goals of that kind with a degree of reliability (e.g., a ratio of true-to-false beliefs or a ratio of successful-to-failed shots) that is enough to meet the relevant evaluative standard (e.g., knowledge-level reliability; gold-medal-level reliability). That cognitive abilities are relativized to so understood appropriate conditions is *the situational assumption of competence-based virtue epistemology* (also known as *virtue reliabilism*).

The previous is a *minimal characterization* of what appropriate conditions are, but at least one that most authors in the virtue reliabilism literature should accept. There are, of course, further options. One may opt for understanding appropriate conditions as *normal conditions* and then give an account of what kind of conditions count as 'normal', e.g., in terms of *statistical normality* (the kind of conditions that are typical), *etiological normality* (the kind of conditions where the ability was learned or selected), *teleo-functional normality* (the kind of conditions that enable an ability to comply with its causal-functional role, if any) or even in terms of *practical relevance* (conditions that are suitable for the practical task at hand).⁵ One can also bypass the identification of appropriate conditions with normal conditions and yet use some of the previous considerations to explain what appropriate conditions are. One can also perhaps renounce to give a general account of appropriate conditions for all abilities and use different senses of appropriateness for different types of abilities. Thus, for *innate abilities*, appropriate conditions might be those to which we are adapted as a species. For example, appropriate conditions for human visual abilities plausibly include a set of light conditions ranging from full sunlight to dim light, but excluding pitch dark—a species with infrared night vision would have different appropriate conditions. For *learned abilities*, appropriate conditions might be those where the ability was learned in the first place. For example, appropriate conditions for the ability to hit baseballs plausibly include a set of wind conditions ranging from no wind to relatively windy conditions, but excluding hurricanes—no one learns how to hit baseballs in hurricane conditions. For *extended abilities*, such as the ability to navigate the environment using a smartphone app like Google Maps, appropriate conditions might include the kind of conditions under which the relevant artifact can comply with the function it is designed for (e.g., being charged).⁶ All these theoretical alternatives state different ways to specify conditions under which one would exercise a given ability with

⁴ Thanks to John Greco for urging me to clarify this and for helpful suggestions.

⁵ See Greco (2012) for relevant discussion of these options.

⁶ For useful discussion of how active externalism (including the extended cognition hypothesis) relates to epistemology, see Carter et al. (2014).

enough reliability for the relevant goal and evaluative standard and, in this way, they are compatible (in fact, specifications of) the minimal characterization of appropriate conditions offered above.

What conditions are inappropriate for an ability, then? Given this minimal characterization, it seems that *inappropriate conditions* are whatever type of conditions in which an agent's otherwise reliable ability would not be reliable enough. In this way, on the reliabilist model, factors such as the presence of bystanders, one's mood, one's energy levels may render an environment *inappropriate* for exercising a cognitive ability so long as they may reduce the reliability of one's abilities in that environment.⁷

However, note that whatever account one gives of the kind of conditions that prevent an agent from exercising her abilities reliably enough ('inappropriate conditions'), one needs to supplement it with an explanation of why some conditions count as inappropriate *despite the fact that one would exercise them with enough reliability*. For example, a skillful archer shooting metallic arrows at a giant magnetic target seems to be in inappropriate conditions for manifesting her archery skills *despite getting close to infallibility*.⁸ The reason is that her conditions are *unduly helpful*, in that a terrible archer would hit the target almost infallibly. In a sense, such conditions *mask* her archery skills.⁹

2.1 Virtue responsibilism: a problematic framework

There is a tension between the situational assumption of virtue responsibilism and the idea that knowledge may involve different degrees of difficulty. On the one hand, it is intuitively *more difficult* to manifest epistemically appropriate behavior in some circumstances than in others. Consider the virtue of intellectual courage. It was certainly more difficult for Galileo Galilei to challenge geocentrism at his time and thus exhibit intellectually courageous behavior than for Albert Einstein to challenge Newtonian mechanics—Einstein's life was not at stake. Yet, the situational assumption of virtue responsibilism demands that intellectually virtuous agents consistently manifest appropriate epistemic behavior across a wide range of circumstances.

One way character-based virtue epistemologists may bring the two ideas together (the idea that situations feature varying degrees of difficulty and the idea that intellectual traits must exhibit cross-situationally consistent epistemic behavior) is to argue, first, that character-based intellectual virtues *come in degrees*; second, that the degree to which one is

⁷ See Alfano (2015) and Turri (2017) for an elucidation of why this version of virtue epistemology is not in trouble with such situational influences.

⁸ See also Sosa's guardian angel case below.

⁹ As we will see below, this special kind of inappropriate conditions bear on the relationship between difficulty and knowledge.

intellectual virtuous depends on the *range of conditions* across which one would manifest the relevant responsibility trait. Thus, circumstances outside that range are such that it is *more difficult* to behave virtuously in them.

To see this, suppose that Galileo and Einstein were equally intellectually courageous, in that they would have manifested intellectually courageous behavior across the same range of conditions, e.g., Einstein would have publicly defended the theory of relativity in Galileo's times, and the other way around. Now consider a contemporary scientist, Cowardeo, who is much less courageous than they were. In particular, in Einstein's circumstances, Cowardeo would manifest intellectual courage, but in Galileo's, he would not. What explains then why Galileo's circumstances are more demanding than Einstein's for Cowardeo is that the former are not part of the range of action of his intellectual virtue.¹⁰

Some virtue responsibilityists might not agree with this externalist characterization of character-based intellectual virtues, but at least helps explain how *learning* intellectually virtuous behavior is possible: one learns how to be intellectually virtuous by including more and more types of circumstances in the set of circumstances over which one's intellectual trait is defined (i.e., by expanding its range of action) until one gets close to what an *ideal intellectually virtuous agent* would do. For instance, in order to learn how to become (or be more) intellectually courageous one must expose oneself to a *variety* of situations in which holding one's ideas publicly puts one at risk of incurring *different types of personal costs* (this is what intellectual courage consists in). Accordingly, exhibiting intellectually courageous behavior may be more or less difficult when one is at risk of (in decreasing degree of difficulty) being murdered, incarcerated, stripped of citizenship, or deprived of one's dinner.

However (and here comes the objection to virtue responsibilityism), not all conditions serve for expanding the range of action of one's intellectual traits. The former are certainly the sort of circumstances in which one is supposed to manifest intellectually courageous behavior if one is to be correctly attributed this virtue, i.e., the type of circumstances that serve to define what intellectual courage *is*. By contrast, the presence of bystanders, being in a good mood or rested are *normatively irrelevant situational influences*, in that intellectually courageous people are supposed to consistently manifest virtuous behavior independently of whether or not such factors are in place. The *problem* is that such apparently irrelevant situational influences *do* have a bearing on whether or not we end up manifesting intellectually virtuous behavior, as the *situationist critique to character-based virtue*

¹⁰ Thanks to Chris Kelp for the suggestion.

epistemology has plausibly shown. Consider the intellectual virtue of creativity. Some empirical studies show that cognitive tasks that require creativity are better accomplished when subjects are in a good mood, but being creative-while-in-a-good-mood is for sure not an intellectual virtue.¹¹

Now, if our epistemic behavior is subject to normatively irrelevant situational influences such as mood elevators, mood depressors, the presence of bystanders or even sounds and smells, it is easy to see that these factors also have a bearing on *how difficult* it is for us to manifest virtuous epistemic behavior. For example, being intellectually courageous while happy is arguably easier than when one is depressed. However, it follows from the notion of intellectual courage that, if one is in its possession, one ought to manifest it *independently of such factors*. In other words: such factors *should* pose no difficulty when it comes to manifesting intellectually courageous behavior, but the problem is that they *do*. Consequently, it is unclear how the relationship between difficulty and knowledge can be account for in virtue responsibilist terms.

To put it differently, one could argue that to learn how to become (or be more) intellectually courageous one must not only expose oneself to different situations in which holding one's ideas publicly puts one at risk of incurring different types of personal costs (as pointed out earlier), but also to situations in which the aforementioned normatively irrelevant factors are in place. However, accepting this idea entails casting doubt on the very *nature* of intellectual courage. While behaving intellectually-courageously-while-at-risk-of-incarceration is the kind of thing intellectually courageous people are supposed to do, behaving intellectually-courageously-while-in-a-good-mood is not. On a character-based model, then, being in a good or a bad mood cannot be part of the explanation of why it can be more or less difficult to manifest (or to learn how to manifest) intellectually virtuous behavior and hence to acquire knowledge.

2.2 Virtue reliabilism: a promising framework

Virtue reliabilists are happy to accept that intellectual virtues and in particular *cognitive abilities are situation-dependent even on normatively irrelevant factors* (irrelevant for the manifestation of responsibilist traits).¹² After all, unlike intellectually virtuous character traits, which are supposed to consistently manifest appropriate epistemic behavior across the board and, more importantly, independently of normatively irrelevant situational

¹¹ For a very helpful monograph on moral and epistemic situationism, see Alfano (2015).

¹² Epistemologists defending virtue reliabilist accounts or virtue reliabilist conditions on knowledge include Broncano-Berrocal (2017; 2018), Carter (2014), Greco (2010; 2012), Kelp (2013), Jarvis (2013), Pritchard (2012), Turri (2011), and most prominently Sosa (2007; 2015).

influences, cognitive abilities are expected to operate reliably under a much more restricted set of circumstances. These include factors that are normatively irrelevant for manifesting responsibility traits but that are normatively relevant for manifesting cognitive abilities, insofar as *they may have a bearing on how reliably they are exercised*.¹³

This is best illustrated by empirical evidence, which shows that *cognition is very sensitive to situational influences* (e.g., if a statement is printed in an easier-to-read color, we are more likely to believe it; if we are tired, we are more likely to overestimate the distance of an object; if we hear spoken arguments that are uttered quickly, we are more likely to believe their conclusions).¹⁴ Although some might worry that this speaks against the existence of cognitive abilities,¹⁵ the upshot, as several virtue reliabilists have argued, is that the range of action our cognitive abilities is not broad but *narrow*.¹⁶

Accordingly, virtue reliabilists can in principle argue that knowing a proposition is difficult to the extent that one attempts to form a belief in it in inappropriate circumstances, i.e., when such epistemically disrupting situational influences are in play, and easy (or easier) otherwise. If we add to this possible explanation of the relationship between difficulty and knowledge the fact that virtue reliabilism defines knowledge in terms of the same concept that serves to define difficulty (namely, the concept of skill), there does not seem to be any *prima facie* tension between the competence-based framework and the idea that knowing can be difficult. However, as we will see next, virtue reliabilism can neither account for the relationship between difficulty and knowledge, nor for the nature of knowledge unless it drops its key situational assumption first.

3 Virtue reliabilism and the appropriateness requirement

Virtue reliabilism gives the following simple answer to the question of what propositional knowledge is: knowing that *p* is a matter of believing *p* truly because of an exercise of cognitive ability. This simple definition of knowledge is backed up by a more complex and explanatorily powerful normative framework for evaluating beliefs, which is in turn based on the assumption that to hold a belief—or alternatively, to form a belief—is a sort of performance, namely a *cognitive performance*.¹⁷ The guiding idea is that, since beliefs are kinds of performances, they can be evaluated along the same normative dimensions of assessment that serve to evaluate performances in general. In this sense, virtue reliabilists think that (at least) four dimensions of assessment are relevant to account for knowledge:

¹³ Recall: the kind of circumstances that sustain enough-levels of reliability (or truth-conduciveness) is what virtue reliabilists call *appropriate conditions*.

¹⁴ See Olin & Doris (2014) for an overview of the relevant empirical evidence.

¹⁵ Olin & Doris (2014) share this skepticism.

¹⁶ See Sosa (2017) and Turri (2017) for this reply.

¹⁷ See Chrisman (*forthcoming*) for relevant discussion of this assumption.

(1) *success*, (2) *competence* (or *ability* or *skill*), (3) *success because of competence*, and (4) *achievement*.

Success. Performances can be *successful* or *unsuccessful*. A performance (e.g., making a recipe) is successful if and only if its aim is accomplished (e.g., just in case the resulting dish is tasty). Analogously, beliefs (i.e., cognitive performances) are (epistemically) successful if and only if their aim is accomplished (viz., just in case they are correct).¹⁸

Competence. Performances can be also deemed *competent* or *incompetent*. A performance is competent if and only if a competent or skillful agent (i.e., skillful in the performance's domain) produces it (e.g., a dish can be considered competent, or competently made, when it has been made by a competent cook). Likewise, an agent's belief is competent if and only if the agent uses her cognitive abilities to form it.

What is the relationship between the competence and success dimensions of assessment? They are related in that competent performances are often successful (and the other way around), but they are ultimately *independent*. First, if a performance is competent, it is not necessarily successful. A competent chef's magistral dish may be ruined by the hair of a careless kitchen helper. Likewise, beliefs that are the product of cognitive ability may turn out false. One must recall that abilities and cognitive abilities are reliable, not infallible, which means that they can produce unsuccessful outputs from time to time. Second, if a performance is successful, it is not necessarily competent. A chef's apprentice might end up making a delicious cake by mixing random ingredients in random proportions and baking the compound in the oven. Analogously, epistemically incompetent agents may form true beliefs by luck. For example, from time to time fortune-tellers get things right about the future by reading Tarot cards.

Success because of competence: Thus, performances (making a cake, forming a belief) may be successful (delicious, true), yet they may fail to be successful *because of* (or *due to*, or *in virtue of*) skill or ability, but because of luck. One popular way to understand this is as follows: to succeed in some endeavor because of ability (cognitive or otherwise) is for that success to *manifest* the exercise of such an ability. By way of comparison, a delicious cake made by a well-prepared pastry chef manifests the pastry-making skills of the chef. However, if an apprentice succeeds in making the same cake by randomly mixing ingredients, the fact that it is delicious does not manifest any pastry-making skill. Analogously, finding out the solution to a complex calculation task by properly following

¹⁸ Performances can have more than one aim, e.g. a cook can aim at making a delicious *and* a beautiful dish. Likewise, at least according to Sosa (2011: 15), beliefs may be aimed at truth but also at some practical goal (e.g., self-relief).

the right steps manifests mathematical competence. In contrast, correctly guessing the solution manifests none.

Recall that virtue reliabilists think that knowledge is a matter of believing a true proposition p because of an exercise of cognitive ability. This now translates as the thesis that one knows p just in case one's cognitive success (i.e., the fact that one gets things right about p) manifests the exercise of one's cognitive abilities. Crucially (and here is where the *situational assumption of virtue reliabilism* sneaks in the theory of knowledge), precisely because cognitive abilities are conceived as dispositions to form true beliefs in *appropriate conditions*, virtue reliabilists assume (without much argument) that an agent's cognitive success manifests her cognitive abilities only if exercised in such conditions (understood in terms of the minimal characterization of appropriate conditions given in §2). This means that knowing a proposition requires that one's cognitive success manifests the exercise of one's cognitive abilities in so understood appropriate conditions. Let's call this the *appropriateness requirement*.¹⁹

The appropriateness requirement is not an innocent assumption. Virtue reliabilists use it to rule out cases (i.e., counterexamples) as a cases of knowledge. For example, according to Kelp (2013) if one consults a thermometer that is fluctuating randomly and one (luckily) forms a true belief about the temperature on this basis, one does not possess knowledge because the situation is not appropriate for the ability of placing trust on reliable devices such as thermometers.²⁰ According to Greco (2012), the driver of the famous fake barn case does not manifest visual competence when forming the belief that there is a barn in front of him when looking at a real barn because the presence of nearby fake barns makes the environment inhospitable for his typically reliable visual capacities. Finally, according to Sosa (2015), if a competent archer shoots an arrow and an unexpected gust deviates it but a guardian angel puts it back on track, the archer does not manifest her shooting abilities because the circumstances are inappropriate—note that the archer's circumstances are inappropriate in a different sense than the circumstances in the previous two cases: this time the likelihood of success is high; what happens is that the

¹⁹ Virtue reliabilists (e.g., Turri 2011) typically use the analogy with physical dispositions (e.g., solubility) to make this point: in the same way as salt manifests its solubility only in appropriate conditions (e.g., when stirred into water but not into gasoline), the success of an agent's performance is taken to manifest the agent's abilities only if the conditions are appropriate. But there is an important respect in which physical and agential dispositions differ. Whereas it is physically impossible for salt to dissolve into gasoline—the probability of this happening is 0—, it is not impossible but difficult for agential abilities to be manifested in inappropriate conditions—the probability of this happening is greater than 0. See §3.1 for more on this point.

²⁰ See the Temp case in §5.

circumstances are too friendly in that shots would be successful no matter how terribly one shot.²¹

Achievement. The fourth dimension of assessment is that of *achievement*. Achievements are a subset of successful performances. In particular, as understood in the virtue epistemological literature, they are successes because of ability, viz., successes that *manifest* ability. For example, the goal of a striker is an achievement if she scores because of her abilities, i.e., if the goal manifests them. A delicious cake is an achievement if its deliciousness manifests the pastry-making abilities of the chef. An excellent dissertation is an achievement if its quality manifests the PhD student's academic skills.

Why do virtue epistemologists appeal to the notion of achievement if the success-because-of-competence dimension of assessment already serves to account for knowledge? Because it allows them to give the following simple solution to the value problem, i.e., the problem of explaining why knowledge is more valuable than mere true belief. The argument, as reconstructed by Pritchard (et al. 2010: 31) goes like this: (1) An achievement is a success because of ability, i.e., a success that manifests ability; (2) knowledge is a cognitive success that manifests cognitive ability; therefore, (3) knowledge is a cognitive or intellectual achievement (*the cognitive achievement thesis*); (4) since achievements are finally valuable, (5) so is knowledge.

3.1 Against the appropriateness requirement

There is no doubt that the virtue reliabilist explanation of the final value of knowledge is simple and elegant, but the cognitive achievement thesis, when combined with the aforementioned appropriateness requirement, leads to what we can call *the problematic cognitive achievement thesis*: the thesis that knowledge is a cognitive achievement, understood as a cognitive success that manifests the agent's cognitive abilities *only if the conditions are appropriate* (as characterized above).

It does not take much argument to see what is problematic with the problematic cognitive achievement thesis: *it is precisely when the circumstances are inappropriate that human abilities may be most manifested (if successfully exercised) and therefore when the greatest achievements may be accomplished*. Indeed, many successes count as great achievements precisely because the circumstances in which they are skillfully attained are inappropriate. Thus, what is problematic with the thesis is *the appropriateness requirement*.

²¹ Accordingly, we can distinguish two senses of inappropriateness: inappropriate conditions for achieving success and inappropriate conditions for manifesting ability. Most inappropriate conditions are inappropriate in both senses, but sometimes, as Sosa's case illustrates, they come apart.

By way of illustration, consider a feat such as the first successful ascent to Mount Everest by Tenzing Norgay and Edmund Hillary in 1953. Three decades of failed attempts are a good indication that the conditions in Mount Everest were not the most appropriate for the climbing skills of mountaineers at that time. After all, appropriate conditions are the sort of conditions in which one's abilities are reliable enough and hence conditions in which it is likely that one will succeed by exercising such abilities. But Norgay and Hillary succeeded against all odds (i.e., *in spite of* being in inappropriate conditions) *thanks to* their climbing skills. The same can be said about Annette Fredskov who, after being diagnosed with multiple sclerosis, ran 366 marathons in 365 days between 2012 and 2013, despite conditions of extreme tiredness and muscle soreness. Another illustrative feat is that of Fernando Parrado and Roberto Canessa, two of the sixteen survivors of the 1972 airplane crash in the remote Andes, who after two months of extreme survival conditions decided to embark on an expedition to look for help and save the rest. It took them 10 days through the merciless peaks and valleys of the Andes, without mountain clothing or proper equipment, in conditions of thin oxygen, extreme exhaustion, sickness and bitter cold, to reach the end of the snowline. Soon after, they found a horseman who contacted the authorities, who immediately set up the rescue.

In these success stories great abilities are manifested *in inappropriate conditions*—indeed, the magnitude of these achievements is greater precisely because of it—, which not only invalidates the idea that abilities and skills can be *only* manifested in appropriate conditions, but also the corresponding claim that achievements are successes that can be *only* accomplished in such conditions.²²

The same considerations apply to the intellectual domain. Consider *intellectual feats* such as decisive breakthroughs in science. Take the Poincaré conjecture. This Millennium problem defeated the efforts of mathematicians for almost a century until Grigori Perelman proved it in 2002. The inappropriateness of the conditions when trying to find proofs for problems of such a caliber does not obviously lie in the harshness of the environment (like Mount Everest or the Andes), but in the fact that one may need to connect fields of mathematics that may apparently have little to do with each other, or resort to complex

²² Two caveats are in order. First, succeeding in inappropriate conditions does not necessarily mean that the relevant success is an achievement, let alone a feat, as one may succeed completely by fluke in such circumstances. Second, inappropriate conditions are typically such that they make success unlikely, but we have also seen that some circumstances may be inappropriate precisely by making success unduly likely or by masking competent success. Sosa's guardian angel case and the case of the archer shooting metallic arrows at a giant magnetic target are two examples. Unlike the usual kind of inappropriate conditions (conditions in which success is unlikely), this unusual kind of inappropriate conditions can make things *easier*, and thereby lead to lesser achievements.

mathematical knowledge that is beyond the understanding of very skillful mathematicians, or develop techniques or concepts that have not been developed yet.²³

Such enormously difficult problems are not unfamiliar to philosophers either. Consider a hard problem: how physical processes in the brain give rise to the conscious experience of the mind and the world? To answer this question one may need to connect different disciplines that apparently have little to do with each other, such as neuroscience and philosophy, or use knowledge about the world or the brain that is beyond current scientific knowledge (e.g., how can we manipulate physical matter so as to give rise to conscious states?), or perhaps develop rather radical ideas (e.g., the idea that consciousness is a fundamental property of the world, just as space, time or mass). Given this inappropriate intellectual conditions (i.e., the current state the art), it is very unlikely that philosophers or scientists, by exercising their reliable intellectual capacities, will come up with a satisfactory solution to the hard problem in the short term (at least one that would be accepted by all parties in dispute). However, suppose that someone did: that person's solution would be considered an extraordinary intellectual feat that we would be happy to call 'knowledge'.

The bottom line is that some decisive breakthroughs count as cognitive or intellectual achievements *despite* being made in inappropriate conditions, i.e., in conditions where intellectual success is unlikely given the exercise of otherwise reliable cognitive abilities. Therefore, the cognitive achievement thesis that we have previously labelled 'problematic'—i.e., the thesis that knowledge is a cognitive achievement, understood as a cognitive success that manifests cognitive ability *only if the conditions are appropriate*—is false because the appropriateness requirement it assumes (namely, the assumption that knowledge requires the manifestation of cognitive ability in appropriate conditions) is false.

In order to avoid possible misunderstandings, let me make more explicit my concerns about this flawed assumption of virtue reliabilism. The appropriateness requirement says that an agent's cognitive success can only manifest her reliable cognitive abilities in appropriate conditions, i.e., in conditions that sustain or guarantee the threshold of reliability that is presumably enough for knowledge and hence conditions that make cognitive success likely. As one might have noticed, this requirement introduces a

²³ A different source of inappropriateness when pursuing an intellectual endeavor is not its complexity, but adverse social or cultural circumstances. For instance, the outstanding discoveries of Marie Curie (who holds two Nobel prizes for her research on radioactivity and the discovery of new chemical elements such as radium and polonium) are even more remarkable in view of the fact that her scientific career was often *hindered* by male colleagues simply because of being a female researcher. This kind of inappropriate conditions do not make the relevant success unlikely given the exercise of reliable intellectual abilities. Instead, what they make unlikely is the very exercise of such reliable abilities.

qualitative (viz., an all-or-nothing) criterion that splits circumstances into two types: (1) *appropriate conditions*, i.e., those in which cognitive abilities are reliable enough, cognitive success is likely and hence knowledge can be acquired, and (2) *inappropriate conditions*, i.e., those in which cognitive abilities are not reliable enough, cognitive success is unlikely and knowledge *cannot* be acquired—because *no* eventual cognitive success can manifest the relevant cognitive abilities when exercised in such circumstances. However, it is actually the other way around: *it is precisely when circumstances are inappropriate that our intellectual abilities may be most manifested (if successfully exercised) and the greatest breakthroughs in knowledge may be accomplished.* Note that the use of examples of intellectual feats such as decisive breakthroughs in science is no coincidence: an intellectual feat is the best example of a cognitive achievement, and if anything deserves to be called ‘knowledge’, this is a breakthrough in science. Yet, these are things that are sometimes accomplished in impoverished epistemic conditions.

The final reason to drop the appropriateness requirement is that splitting circumstances into ‘appropriate’ and ‘inappropriate’ offers an incorrect representation of the relationship between difficulty and knowledge. The requirement says that *no* cognitive success can manifest the exercise of cognitive ability (and hence amount to knowledge) when the relevant conditions are inappropriate. Therefore, knowledge in inappropriate circumstances is *not just difficult* to achieve (as one would naturally expect), *but impossible*.²⁴ But then it is unclear how *human learning* can be explained. In order to expand the range of action of our abilities (the conditions in which we would be reliable), we often need to expose ourselves to conditions outside that range, i.e., to conditions in which our performances are less reliable (even very unreliable) and hence conditions where success is unlikely. This is how we typically come to learn new things, i.e., how we come to acquire new practical and propositional knowledge. In this way, the learning process is such that we incorporate new circumstances to the range of action of our abilities. But this is for sure something that does not happen automatically, but gradually—hence why it seems misguided to split circumstances *categorically*, i.e., into circumstances where one *can* know and circumstances where knowledge (learning) is *impossible*. The reality is that at some point during the learning process we may succeed cognitively despite being something unlikely. In some such cases, we consider the relevant cognitive successes achievements, even feats, and indeed, some of them (such as extraordinary breakthroughs in science) may be properly called ‘knowledge’.

²⁴ Note that if φ -ing is difficult for S, the probability that S will φ may be low, but greater than zero, whereas if φ -ing is impossible for S, the probability that S will φ is zero.

4 Knowledge as a (more or less) difficult cognitive achievement

The key tenet of virtue reliabilism that knowledge is a cognitive achievement (a cognitive success that manifests cognitive ability) turns out false on the assumption of the appropriateness requirement. This requirement also leads to an incorrect representation of the idea that knowledge can be difficult. In that follows, and to amend the latter problem, I will combine the virtue reliabilist framework with recent work on the notions of achievement and difficulty.²⁵ This will help understand how knowledge relates to difficulty without appeal to considerations about appropriate or inappropriate conditions and hopefully shed some light on the nature of knowledge as well (especially on the thesis that knowledge is a kind of cognitive achievement).

If the cognitive achievement thesis is correct, then the extent to which knowledge is difficult corresponds to the extent to which it is difficult for a cognitive success to qualify as a cognitive achievement in the way virtue reliabilists understand this notion. In what sense can a belief be considered a more or less difficult achievement? Whereas difficulty is a *property of activities, performances or processes* and then derivatively of the states that result from them, achievements are *states*, namely the results of successful performances. Consider Usain Bolt's breaking the 100m world record in the 2009 Berlin World Championships. The activity/performance/process of running faster than anyone else in an official competition is very difficult. Having crossed the finish line in 9,58 seconds, i.e., the state of being beyond the finish line in 9,58 seconds, is an achievement. It is an achievement because (in keeping with the virtue reliabilist understanding of the notion of achievement) it is a success that is due to an exercise of ability. Likewise, beliefs are states that (according to virtue reliabilism) count as intellectual or cognitive achievements if their correctness (i.e., their 'success') is due to the exercise of cognitive ability. This is the sense in which beliefs can be considered achievements.²⁶

4.1 Achievements: their (class-relative) nature vs. their (value-relative) magnitude

In a narrow sense, by 'achievement' we mean an exceptional accomplishment, i.e., a feat, or as Bradford (2015) puts it a 'capital-A achievement'. The first ascent to Mount Everest, the proof of the Poincaré conjecture or Marie Curie's discoveries are practical and intellectual achievements in this sense: they are successes because of an exceptional display of ability

²⁵ Three authors who have recently cast light onto the notions of achievement and difficulty are Gwen Bradford (2015), Hasko von Kriegstein (2017; *forthcoming*), and Alexander A. Guerrero (2017). This section is largely indebted to their work.

²⁶ Because I am interested in investigating how knowledge, *qua* cognitive achievement, can be more or less difficult, here I am assuming that achievements can be *more or less* difficult, i.e., I don't consider being difficult (i.e., difficult to a high degree) a necessary component of achievements. There is some controversy concerning this point, though. See Bradford (2015) and von Kriegstein (*forthcoming*) for extensive discussion.

and effort. In a broader sense, however, we call an achievement any success because of ability, even if mundane, such as tying a shoelace, carrying a heavy shopping bag, threading a needle or finding out whether it's raining outside. What is the difference between capital-A and mundane achievements? In a nutshell: their *magnitude*. The magnitude of an achievement typically has to do with the effort invested and the skill needed to attain it. But very effortful or skillful performances can still result in mundane achievements. For example, counting the grains of sand in a heap one by one is no doubt a very effortful success. But it is not a capital-A achievement. Burping at 110 decibels is a very skillful success, but again, not a capital-A achievement. The reason is that the resulting knowledge (in the former case) and the resulting loudness (in the latter) have little significance or value. This suggests that the *magnitude* of achievements features a *significance* or *value condition*: the more significant or valuable a practical or intellectual success is, the more of an achievement it is.²⁷

Knowledge can be a capital-A intellectual achievement, such as the proof of the Poincaré conjecture, the (eventual) discovery of the ultimate cure for cancer, or an (eventual) satisfactory solution to the hard problem of consciousness. But for the most part, knowledge is a mundane achievement, such as deducing the answer of a problem in a maths exam or knowing the departure time of the next train at the station. Like the magnitude of practical achievements, the magnitude of intellectual achievements depends to a large extent on how significant or valuable they are.

However, the value or significance of a cognitive success (*viz.*, of the true proposition believed) is not relevant to whether or not it counts as an intellectual achievement.²⁸ The crucial element is whether it manifests the cognitive abilities of the agent as well as the effort invested. To see why value considerations are not important, consider two scenarios. *First scenario*: to pass the time, someone correctly counts the grains of sand in a heap one by one. *Second scenario*: finding out how many grains of sand are there in a heap (with equal number of grains) is the only way to quench the thirst of blood of a furious god and save humanity from extinction; a hero succeeds in counting the grains, one by one. In both scenarios, the relevant performances lead to the same cognitive success, *i.e.*, a true belief about the number of grains, which manifests ability and effort to the same

²⁷ Suppose that on three different occasions I sign three cheques for different fundraising events for cancer research: one for \$1, one for \$2, and the third for \$1,000,000. On the present proposal, there is an increase in the *magnitude* of the achievement, even if its nature is the same in the three cases (the effort and the skill manifested are the same). Donating \$1,000,000 is certainly more of an achievement than donating \$1 or \$2 (in the sense of having more personal costs and benefits for the benefactor and more impact on the target group or society).

²⁸ I use the terms 'cognitive' and 'intellectual' interchangeably. That said, the term 'intellectual' denotes successes or achievements that result from higher forms of inquiry (*e.g.*, scientific research), whereas 'cognitive' denotes successes or achievements that result from more basic cognitive faculties, such as perception.

degree. The difference between the two cases lies in the significance or value of the resulting knowledge. Does this mean that the relevant cognitive success in the low-stakes scenario doesn't count as a cognitive achievement? No, it is certainly *less of an achievement* (not a capital-A achievement for sure), but a cognitive achievement after all. In other words, whereas the value of an achievement decreases or increases its *magnitude*, it does not affect its *nature* or status *qua* achievement.²⁹ Since the aim is to investigate how difficult it is that a cognitive success *qualifies as* a cognitive achievement in the relevant sense and thus as knowledge, in what follows I will omit considerations about the value of cognitive achievements.

Suppose with virtue reliabilism that knowledge is a cognitive achievement (which may have different magnitudes). There is a problem in the offing. On the one hand, *knowledge* is considered a *universal phenomenon*, at least universal in the sense that it is attributed to and may be instantiated by *many* creatures. On the other hand, the notion of *achievement* is always relativized to a reference class, which may include only one individual (i.e., successes that count as achievements only *for* those who produce them). For example, a baby's first steps is an achievement *for* the baby as an individual and perhaps also *for* the entire baby collective, but such a remarkable success might not be considered an achievement for healthy adults (at least not in the capital-A sense of achievement).³⁰ How can knowledge, *qua* achievement, be considered universal if the notion of achievement is relative?

The first thing to note is that some achievements are considered universal at least in the sense that the relevant reference class is humanity as a whole. Consider for instance the eradication of infectious diseases or a stable world peace: they are considered universal accomplishments of humanity. We also tend to take this more comprehensive stance towards relevant successes that are accomplished for the first time in history, such as the first ascent to Mount Everest or the first time a human landed on the moon. But the sense in which knowledge is a universal achievement is not the same the sense in which these achievements are considered universal: the latter are universal because they are feats, i.e., because of their magnitude, but as we have seen knowledge can be mundane.

The reason why knowledge constitutes a universal achievement is the fact that it is a *distinctive* state that can *in principle* be possessed by *any* being capable of holding doxastic states with propositional content or at least capable of having contentful mental representations. This includes humans, other higher-order mammals, many animals with

²⁹ See von Kriegstein (2017; *forthcoming*) for helpful discussion on the distinction between the magnitude and the nature of achievements.

³⁰ Again, see Bradford (2015) and von Kriegstein (2017; *forthcoming*) for extensive discussion on this point.

less sophisticated nervous systems and complex forms of AI.³¹ Of course, this does not mean that beings in this list *actually* have knowledge, because their methods of belief formation may be unreliable. The idea is rather that knowledge is a kind of state that we would universally attribute to doxastic agents *only*, i.e., to agents with belief capabilities.

An unreliable doxastic agent (e.g., a brain in a vat) may not be credited with *actual* knowledge, but could *in principle* come to acquire knowledge. By contrast, we would *not* say the same thing about a stone, a chair or a mountain: they could not in principle come to acquire knowledge—for some relevant sense of ‘could’. To see this more clearly, suppose that the actual world contains two types of things that by their constitution cannot have knowledge: stones and brains in vats. It takes much more departure from the actual world to change the constitution of a stone so it can have knowledge than to change the constitution of a brain in a vat. After all, the latter may be restricted by constitution to fairly unreliable ways of forming beliefs (barely above chance), but unlike the stone brains in vat satisfy the pre-condition that makes knowledge universal: having the capacity to form contentful mental representations (beliefs). This is why the brain in vat, unlike the stone, can *in principle* have knowledge.

What reference class is relevant to evaluate knowledge as a distinctive cognitive achievement then? Unsurprisingly, the relevant reference class that is *assumed by default* in non-contextualist epistemology (at least in virtue epistemology) is the class of agents capable of forming contentful doxastic representations.³² This reference class may be of course subdivided into a narrower reference class when it comes to evaluate whether there is actual knowledge (e.g., the class of agents who form contentful doxastic representation in reliable ways, who in doing so manifest cognitive ability, agents who follow their evidence, and so on). In epistemology, there are of course many competing views about how to define this reference subclass of knowers *in a context-independent way*.

Why should this reference subclass *not* be fixed by context? Because if the reference class relative to which an agent’s cognitive success qualifies as a *knowledge-level cognitive achievement* included only individuals that are salient in the context of assessment, then too many cognitive successes would count as knowledge. For example, suppose that your preferred way of forming beliefs is wishful thinking. Suppose that you get things right this time. Is your cognitive success a cognitive achievement? Given that most of the times you

³¹ What exactly makes knowledge *distinctive* is an open question: virtue reliabilist might argue that only true beliefs that have the ‘knowledge’ trademark manifest cognitive ability; others might argue that knowledge is distinctive because it serves to flag good informants or stop inquiry. The way we fill in the details is not as relevant as the general point that, for all beliefs possibly held by doxastic agents, there is a subset that we distinctively call ‘knowledge’.

³² When the kind of knowledge is reflective knowledge, the typically assumed reference class excludes human infants and non-human animals.

form false beliefs, and given the class-relativity of the general notion of achievement, this *does* count as a cognitive achievement *for you* in your present context (after all, albeit luckily, you get things right). But relative to a relevant *non-contextually fixed* reference subclass of doxastic agents (e.g., reliable doxastic agents, doxastic agents whose cognitive successes manifest their cognitive abilities, doxastic agents who follow their evidence, and so on) your cognitive success does not represent a (knowledge-level) cognitive achievement. To be clear, *some cognitive successes count as cognitive achievements but do not amount to knowledge* because there are always contextually-fixed reference classes of individuals following bad epistemic practices (e.g., wishful thinking) for which such cognitive successes count as cognitive achievements. In particular, their less epistemically defective practices (e.g., lucky wishful thinking) can be considered *achievements for them*. However, once again, relative to the relevant non-contextually fixed reference subclass of doxastic agents that serves to evaluate whether there is actual knowledge bad epistemic practices are what they are: bad epistemic practices. In this way, the class-relativity of the notion of achievement is no impediment to conceptualizing knowledge as a (special kind of) cognitive achievement.

4.2 Difficulty: likelihood, skill, effort, and resources

Similar considerations apply to the notion of difficulty. Difficulty is always relative: when we deem something difficult, we always think of it as difficult *for* the agents of a relevant reference class. This is obviously relevant to the question of how knowledge can be difficult. How is the relevant class of agents fixed in each case? Here is Bradford:

The relevant class of agents for any particular instance of difficulty is determined, I think, fairly straightforwardly. Like many things of this nature, it is determined by context. Whatever activity we are interested in evaluating, it seems that we can quite easily pick out the relevant class of agents for whom it may or may not be difficult. The default may be the class of adult human beings, but of course many activities are difficult for four-year-olds, but not for the typical adult. Whatever class it is that is relevant for the evaluation of whether or not some activity is difficult is generally apparent from the context of the activity, or can be discerned or stipulated without much controversy. (Bradford 2015: 28)

How does this bear on the explanation of the idea that knowledge can be difficult and particularly once knowledge is theorized as a special kind of cognitive achievement? As we have just seen, when we evaluate whether a given cognitive success amounts to a *knowledge-level cognitive achievement* (i.e., a cognitive achievement in the relevant sense and not only *for* the agent in question), the reference class that epistemologists assume *by default*

(i.e., not contextually) is a subclass of the class of agents capable of holding doxastic states. The idea is that if relative to this reference subclass a cognitive success counts as a cognitive achievement, then it can be universally considered knowledge (universally for all doxastic agents). However, when it comes to the assessment of whether or not a certain proposition is *difficult to know*, things are different. If Bradford is right, here the relevant reference class is fixed *by the context of assessment*, not by default. This means that knowing can be difficult relative to a reference class that may include from all agents capable of holding beliefs to one single agent. To see this, take a simple cognitive success, such as answering correctly what does $2 + 2$ equal. Is it difficult to know what does $2 + 2$ equal? It depends. If the relevant reference class only includes educated adults, then it is not difficult. If it only includes two-year-olds, then it is difficult. It is a matter of what sort of individuals the salient reference class includes.

We can generalize these ideas for any practical or intellectual endeavor E in the following way: if the most skillful agents in a contextually-fixed reference class are unlikely to succeed in E , E is difficult for all members of the reference class. Conversely, if the less skillful agents in a reference class are likely to succeed in E , E is easy for all members of the reference class. Comparatively, E is more difficult for members of reference class R_1 than for members of reference class R_2 if the less skillful agents in R_2 are more likely to succeed in E than the most skillful agents in R_1 . In this way, finding out what does $2 + 2$ equal is more difficult for two-year-olds than for educated adults because the most skillful two-year-olds are less likely to find out the correct answer than the less skillful adults.³³

So far I have been assuming a *probabilistic conception of difficulty* according to which how difficult (or easy) is it for an agent to reach a goal corresponds to how likely the agent

³³ The present project is to explore how the concept of difficulty relates to intellectual endeavors in general and especially in connection with the notion of achievement. A different question which cannot be answered from the armchair is how to *measure difficulty* in each domain. In some domains, there are objective metrics, such as videogames or high jump. In the latter case, for example, the higher the bar is, the more difficult to jump over it without dislodging it is. In some other domains, there are also metrics for difficulty, but they are subjective. For example, in bouldering there are detailed scales that measure the physical difficulty of a 'boulder problem' (i.e., the route that a climber takes to complete a climb). The first person who sets the route makes a subjective assessment of the relevant grade, which may be then slightly modified by consensus. In some other domains, there are no scales or metrics, but there is a shared implicit agreement on what is easy and what difficult in the domain. For example, in the study plan of a maths degree linear algebra is typically taught before complex analysis because it is easier, just like propositional logic is taught before modal logic in philosophy degrees. In other areas, it is not obvious what is more difficult than what, e.g., some might find Frege an easy read, but some might find it difficult and opt for Husserl instead. The same kind of considerations apply to the *measurement of skill levels*. Sometimes, it is easy to pin down the features that make someone an expert in the domain, but sometimes this is less obvious. Empirical research often helps shed light on the issue. For example, in the case of bouldering one study concluded that "more advanced ability climbers make greater use of foot holds, with associated lowering in physiological response (oxygen uptake and heart rate) across all slope inclinations" (Baláš et al. 2014: 75). In the domain of physics, another study concluded that a key factor that distinguishes experts from novices in the representation of physics problems is that "the experts initially abstract physics principles to approach and solve a problem representation, whereas novices base their representation and approaches on the problem's literal features" (Chi et al. 1981).

is to reach it.³⁴ In the foregoing paragraph, I have also assumed that the factor with most weight on an agent's chances of success with respect to some goal—and therefore on how difficult for the agent to reach that goal—is how skillful the agent in the relevant domain is. But as Bradford (2015) and von Kriegstein (2017; *forthcoming*) have suggested (and as pointed out before), another relevant factor that determines the difficulty of reaching a goal is the amount of *effort* invested. What is the relationship between invested effort, exercised skill and likelihood of success? von Kriegstein (2017) nicely summarizes how they are related in the following passage:

Actions generally involve a combination of effort and skill, and this combination (together with facts about the goal or intention) determines the likelihood of success. As a general rule skill and effort can be substituted for each other to some degree: a less skilled agent can make up for her lack of talent with great effort, and a lazy agent may compensate her lack of effort through great skill. Keeping skill level constant, then, we should generally expect a correlation between effort and contribution to likelihood of success. That means that effort, via this contribution, has an achievement-enhancing effect. (von Kriegstein 2017: 43)³⁵

How is effort to be understood in this context? A plausible view, advanced by von Kriegstein, understands effort as a matter of *dedicating resources*. Von Kriegstein conceptualizes the relevant resources as *internal*, but does not say much about their nature (only that they can be physical or mental). Let me elaborate on the notion of resource then.

The term 'resource' may be used to refer to whatever method one has when it comes to accomplishing a certain task (φ -ing). However, you may have a method that is completely useless for φ -ing. When that happens, and if you have no other method at your disposal, we would correctly say that you have *no* resources for φ -ing. It seems, then, that the notion of resource features some sort of *reliability condition*. One plausible idea is that one's resources for φ -ing are one's *abilities*, where these need not be internal, as there may be abilities based on extended cognitive processes.

However, this might not yet suffice for capturing the notion of resource: it also seems to feature an *availability condition*. Suppose that you want to hunt a rabbit for dinner and that you are very skillful at hunting rabbits with your bow, but someone has stolen it. In a sense, we might say that you have the resources for hunting a rabbit: you are a skillful hunter after all. In another (I think more plausible) sense, you don't have the resources:

³⁴ See von Kriegstein (*forthcoming*) for this account.

³⁵ Guerrero (2017) makes a similar point by distinguishing two kinds of difficulty: effort and skill-related.

while you have the relevant abilities, you cannot *exercise* them. It seems then that in order for an agent to have resources for φ -ing, she not only needs to have the ability to φ but also be in a position to φ , i.e., be in conditions that *enable* her to exercise her ability to φ . In other words, exercising an ability to φ is something that needs to be available in the circumstances of φ -ing.

The fact that certain enabling conditions for exercising an ability to φ are in place does not entail that the conditions are *appropriate* for φ -ing. Conceptually, these are two different things. In the previous example, you cannot hunt a rabbit because you don't have your bow. Having a bow is a condition that *enables* you to exercise your ability to hunt rabbits and, consequently, when you have a bow at your disposal, you have the resources for hunting a rabbit—the presence of rabbits is also an enabling condition. But the fact that you are in conditions that enable you to exercise your hunting skills is compatible with the fact that you are in 'inappropriate' conditions for hunting, i.e., conditions in which hunting a rabbit successfully is difficult or unlikely despite your abilities (e.g., strong winds, poor light conditions, and so on). These latter kind of conditions *determine the success (or failure) of your performances* rather than *enable the exercise of your abilities*. Having resources for φ -ing only requires that your enabling conditions are in place. It does not require that you are in favorable determining conditions (i.e., conditions that make it likely that you will φ).

A clarification point is in order, as some might have noticed an ambiguity between two competing conceptions of the notion of resource in the previous paragraph. On the present understanding of the notion of resource, you may have resources to φ even in conditions in which φ -ing is unlikely (unfavorable determining conditions). Call this the *permissive conception*. But it could be argued that if φ -ing is unlikely in certain conditions, you don't have the resources to φ in such conditions. Call this the *restrictive conception*. I think the ambiguity originates in a prior ambiguity regarding the notion of ability. On a restrictive conception, you don't have the ability to φ if your conditions are such that φ -ing is unlikely. Having an ability to φ , in this restrictive sense, consists in having a *complete ability* to φ , i.e., in not only possessing whatever constitutional basis makes you disposed to φ in the relevant way, but also in being adequately situated for φ -ing (both internally and externally). On a permissive conception, by contrast, you can have the ability to φ by having the relevant constitutional basis even if you are not adequately situated for φ -ing, i.e., even if you are unlikely to φ .

What conception of resources should we adopt? There are reasons to opt for the permissive understanding, as the restrictive reading leads to counterintuitive results. Suppose that two agents have the same constitutional bases for the disposition to φ but are

in different situations. Do we want to say that they have different resources for φ -ing? A more intuitive claim seems that they have the same resources but they are in different situations. The restrictive reading does not allow to say that, however. If this is not convincing enough, consider another reason. Imagine a quite skilfull agent in unfavorable conditions and a not so skillful agent in favorable conditions so that the likelihood that they will φ is the same. The restrictive conception must say that they have the same resources for φ -ing, but that seems wrong. A trained mathematician has objectively more resources for solving a mathematical problem than a primary school student, even in a given situation the likelihood that they will solve it is the same.

With this account of the notion of resource in place, what is effort then? According to von Kriegstein, it is the degree to which one's resources are put to use. As a general rule, the more resources one has at one's disposal, the more likely it is that one will succeed in accomplishing a task and, therefore, the easier the task will be. But this needs qualification. First, as we have seen, the likelihood of success also depends on the circumstances, which may be more or less demanding. Second, as von Kriegstein (2017: 44) points out, "potential resources simply have no bearing on the likelihood of success as long as they remain untapped". Given the previous characterization of the notion of resource, to 'tap' a resource for φ -ing is just to exercise the corresponding abilities in conditions that enable such an exercise. With these two caveats in place, we may say that what determines the likelihood of success and therefore the degree of difficulty in reaching a goal is how many resources one actually uses of those at one's disposal as well as the kind of *determining* conditions one is in.

What is the *time span* in which one can use one's available resources for reaching a goal?³⁶ von Kriegstein (2017: 45) helpfully sheds light on the issue with another proposal: the relevant time span is "the interval between the adoption of the goal and the last moment when it is possible to reach it". According to von Kriegstein, two types of factors can set a limit beyond which it is not possible to reach a goal with one's resources. The first factor is the very *content of the goal*. Some goals, such as filing a legal document or answering the questions of an exam, have deadlines, i.e., they do not count as 'reached' unless they are accomplished before a certain time. One can of course complete these tasks after their deadlines, but one would be thereby reaching different goals. The second factor that can set a limit beyond which it is no longer possible to reach a goal are the *circumstances* in which one tries to reach it. For example, suppose that the goal is to prepare

³⁶ Time should intuitively count as a resource too, one might think. That is probably right. For the sake of explanation, however, I follow von Kriegstein in treating time as something separately from one's resources, so we can say that there is a *time span* for using one's available resources which is not itself part of one's resources.

a delicious dish with the fresh ingredients in the fridge. This goal can be only accomplished while the ingredients stay fresh, otherwise the dish will not be delicious no matter how well one cooks. Finally, as von Kriegstein points out, the limiting case of a deadline being set by the circumstances is the very death of the agent. Complex mathematical problems are often like this: mathematicians die before solving them.

The temporal dimension is important to understand the relationship between difficulty and knowledge. According to the main tenet of virtue reliabilism an agent knows just in case her cognitive success manifest her cognitive abilities. But this can be now qualified as follows: when an agent knows at t , her cognitive success manifests the exercise of her cognitive abilities in a given time span up to t , which means that what is manifested is not only how competent she is, but also the degree to which she puts her cognitive abilities to use in that interval of time, i.e., her effort.

The reason why the notion of effort should be included in the virtue reliabilist analysis of knowledge is that, as we have seen, ability and effort can be substituted for each other. The implication of this is that the same piece of knowledge can manifest varying degrees of cognitive ability as long as a higher degree of effort offsets a lower degree of ability. By way of illustration, suppose that an expert solves a complex Sudoku puzzle easily and quickly. Her cognitive success (i.e., her knowledge of the solution) manifests great cognitive ability and little effort. In comparison, a novice is less likely to solve the puzzle, i.e., it is more difficult for her. But with dedication and effort she can solve it, and if she solves it her cognitive success can be considered knowledge with the only difference that, unlike the expert's cognitive success, it will manifest more effort than cognitive ability. This interplay between cognitive ability and effort in increasing an agent's chances of cognitive success is at the core of the relationship between difficulty and knowledge.³⁷

5 Knowledge as a (more or less) difficult challenge

This section will condense the foregoing discussion in a simple account of the relationship between difficulty and knowledge. The account is based on two main points. The first point is that, although they often go together, the difficulty of succeeding in a cognitive endeavor—which depends on how competent the agent is and on how much effort she puts in—is different from the difficulty of manifesting competence and effort. There are cases in which they come apart. The second point is that the difficulty of accomplishing a cognitive task (viz., acquiring knowledge) does not require the intentional adoption of a

³⁷ In both cases, the relevant cognitive successes count as cognitive achievements. In the expert case, the relevant cognitive achievement presumably has a greater magnitude because the expert solves the puzzle more efficiently (she needs to invest less effort). This gives rise to an interesting question concerning the value of knowledge: is it more valuable to know something using a greater degree of skill or a greater degree of effort?

goal, because knowledge can be conceived as a challenge of the situation that can be met unintentionally. Let's start with the second point.

Cognitive tasks are diverse: they range from effortful forms of inquiry (e.g., solving complex mathematical problems) to automatic information processing (e.g., perceiving the color of an object). While we may intentionally aim at accomplishing a cognitive task (e.g., solving a differential equation), most cognitive tasks do not need the intentional adoption of cognitive goals to the extent that most of the times we simply have no propositional attitude towards the proposition believed prior to believing it. This marks an important difference with success in the practical domain, which is typically the product of intentional action. That being so, we cannot understand the difficulty of knowing in terms of intentionally adopted cognitive goals being reached. Instead, as I will propose next, we can better understand it in terms of *situational challenges* being intentionally or unintentionally met.

The general idea, in a nutshell, is the following: in each situation, circumstances or context, there are potentially many challenges that an individual in the situation can meet.³⁸ Each challenge poses different degrees of difficulty for the individual in the situation, where how difficult it is for an individual to meet one of these situational challenges depends on the type of challenge in question, her level of skill and the amount of effort invested.

By way of illustration, consider a football game. There are potentially many challenges that a striker can meet: scoring a goal, scoring a hat-trick, making two assists, dribbling past three defenders in a single play, following the orders of the coach, avoiding being shown a yellow card, not getting injured, and so on. All these challenges have different degrees of difficulty for the striker, which may be different from the degrees of difficulty for other players exhibiting different levels of ability and effort. For example, it is certainly easier for a striker to score a goal or to dribble past another player than for a defender, because strikers are in general more competent at those tasks than defenders. Likewise, it is presumably more difficult for a striker to avoid getting injured than for a defender (strikers tend to be more exposed).

In a given situation, there are challenges for the agents involved that are irrelevant to them or the domain. In football, for instance, it is irrelevant whether a player does a perfect *déboulé* or lands a jump in fifth position, but this is very relevant in ballet. In addition, impossible challenges, because they are impossible, are also typically deemed irrelevant (e.g., the challenge of running faster than light in a 100-meter race). The bottom

³⁸ I use the terms 'situation', 'context' and 'circumstances' loosely and interchangeably.

line is that what challenges are relevant or salient in a situation is determined by context or by the domain itself.

Crucially (and here comes one of the main points of the section), an agent in a certain context or situation *does not need to intentionally adopt a goal* to meet the challenges that are relevant or salient. For example, a striker's main aim may be to score a goal, but she may end up scoring a hat-trick, in which case she overcomes two challenges by reaching one goal. The same applies to cognitive tasks. Consider perceptual tasks. Suppose that an agent enters a room with only two objects. There are several perceptual challenges available to this agent in the situation that she can accomplish using her perceptual abilities. For example, she can find out how many objects are there in the room, what is their color, their shape, their location, and so on—richer perceptual environments obviously feature many more and more complex perceptual challenges. The point is that, no matter how many perceptual challenges are available in the situation, one does not need to intentionally adopt the goal of meeting them to actually meet them. It is possible to overcome them spontaneously or unintentionally by carrying out other perceptual tasks as long as these challenges are available in the situation. Another way to put it is this: as in the case of football, one can meet cognitive challenges that are different from one's initial cognitive goals. For instance, when an ophthalmologist asks you to tell the letters on the eye chart, you may intentionally adopt this perceptual goal. But suppose that while you are trying to identify the letters, you notice that the eye chart is white and has a rectangular shape. In this way, you overcome two perceptual challenges (what is the shape of the eye chart? what is its color?) in addition to the perceptual challenge you were trying to meet in the first place (which are the letters in front of you?).

What is the *time span* in which an agent can put her resources to use in order to meet a challenge in a given situation? As we have seen in §4.2, this can depend on the content of the relevant challenge (e.g., if it involves a deadline) or on the circumstances themselves. For example, the challenge of scoring three goals in a game has a fixed time span: the duration of the game. By contrast, the time span of complex intellectual challenges often corresponds to the time span of the lives of those trying to overcome them: it is not rare that a scientist devotes her entire academic life to challenges such as the cure for HIV or a very hard mathematical problem. For more mundane cognitive challenges, such as simple perceptual tasks, the time span typically depends on the circumstances. For example, the time span for perceiving the location of an object is the time the object is at that location.

My proposal, then, is that knowing is a matter of meeting a more or less difficult cognitive challenge in a given situation. However, knowing is a special kind of challenge in

that it does suffice that one succeeds cognitively, i.e., it is not the challenge of succeeding cognitively *simpliciter*, but the challenge of succeeding cognitively *in such a way that one's cognitive success manifests one's cognitive abilities and effort*.

Let's generalize these ideas for any practical and intellectual endeavor. Let φ -ing be any practical or intellectual task (e.g., performing an action, believing a proposition). In a given situation, we can distinguish *three general types of challenges*:

(C1) The challenge of *successfully* φ -ing at t : an agent overcomes this challenge if and only if she is successful in φ -ing at t (e.g., just in case she forms a true belief).

(C2) The challenge of *competently* φ -ing at t : an agent overcomes this challenge if and only if she uses her abilities to φ at t (e.g., just in case her belief is produced by her cognitive abilities).

(C3) The challenge of *successfully* φ -ing at t *because of one's abilities and effort*: an agent overcomes this challenge if and only if her successful φ -ing at t *manifests* the exercise of her abilities and effort, i.e., if and only if her success is a *achievement* (e.g., just in case the fact that she comes to believe a true proposition is due to a dedicated exercise of cognitive ability).

Acquiring knowledge is a challenge of the latter type: it is not the mere challenge of succeeding cognitively in a competent way but in a way that makes one's cognitive success qualify as a *cognitive achievement* in the relevant sense.

Meeting C3 entails meeting C1 and C2, but the converse does not hold. Someone can get things right by wishful thinking (i.e., she meets C1) but this cognitive success does not manifest cognitive ability. Someone can also use her cognitive abilities to form a true belief (i.e., she meets C2) but the fact that she get thing rights may still be by luck (as it happens in Gettier cases), which means that her cognitive success is not because of her cognitive abilities (i.e., she doesn't meet C3). In this way, precisely because C3 entails C1 and C2 but not the other way around, meeting C3 is (at least in principle) more difficult than meeting C1 or C2.

Sosa's guardian angel case is an illustrative example of this mismatch in difficulty. When the competent archer shoots an arrow and an unexpected gust deviates it but a guardian angel puts it back on track, the shot is successful (C1 is met), the archer uses her abilities to shoot (C2 is met), but the shot does not manifest the archer's exercise of ability: it manifests the intervention of the guardian angel instead (C3 is not met). In these circumstances, meeting C1 is very easy because success is likely independently of how

competent the archer is: no matter how badly she shoots, the guardian angel will correct the trajectory of the arrow and make it hit the mark. What is difficult is precisely to hit the mark manifesting archery competence.

The difficulty of acquiring knowledge is like this. *In most situations*, if we exercise our cognitive abilities to form beliefs (and thus meet C2), the resulting beliefs are not only true (C1 is met), but their correctness also manifests our cognitive abilities and effort (C3 is met). Sometimes, however, *these challenges come apart* and so does the degree of difficulty in overcoming them. To see this more clearly, consider a case that is analogous to Sosa's guardian angel situation but in the intellectual domain:

[Temp] forms his beliefs about the temperature in his room by consulting a thermometer on the wall. Unbeknownst to Temp, however, the thermometer is broken and is fluctuating randomly within a given range. Nonetheless, Temp never forms a false belief about the temperature by consulting this thermometer since there is a person hidden in the room, next to the thermostat, whose job it is to ensure that whenever Temp consults the thermometer the temperature in the room corresponds to the reading on the thermometer. (Pritchard et al. 2010: 48-49)

In this case, Temp is presumably competent at reading and trusting reliable devices (C2 is met), he forms true beliefs (C1 is met), but his cognitive success does not manifest this competence, but the intervention of the hidden agent (C3 is not met). As in Sosa's case, meeting C1 is very easy in Temp's circumstances, because he is very likely to form true beliefs about the temperature in the room (no matter how reliable or unreliable he is). What is difficult is overcoming C3: Temp is unlikely to meet this challenge because the intervention of the hidden agent systematically thwarts the manifestation of his cognitive abilities.

Why does the intervention of the agent thwart the manifestation of Temp's abilities? Traditional virtue reliabilism answers this question in terms of the appropriateness requirement (i.e., the idea that an agent's cognitive success can manifest her cognitive abilities only in appropriate conditions). The traditional, unsatisfactory answer is simply that Temp does not manifest his cognitive abilities because the intervention of the agent makes the conditions *inappropriate*. As we have seen, however, the appropriateness requirement is incorrect (and should therefore be dropped) for three related reasons. First, it implies that manifestation of cognitive ability and hence knowledge is something that cannot happen in inappropriate conditions, but this is not true: it is when the conditions are inappropriate that our cognitive abilities may be *most* manifested (this is the case of

intellectual feats, such as decisive breakthroughs in science). Second, and relatedly, it offers an incorrect representation of the relationship between difficulty and knowledge: it makes knowledge impossible (rather than unlikely) in inappropriate conditions. Third, it is unclear how learning can be explained if we assume the appropriateness requirement.

Therefore, we need an alternative explanation of what thwarts manifestation of ability. A possible one goes as follows. Two factors prevent an agent's cognitive success from manifesting her cognitive abilities. The first is *luck*. The second is *masking circumstances*. Let's start with the first. Recall that there are two kinds of inappropriate conditions for the exercise of cognitive ability. The first type make cognitive success unlikely (e.g., poor light conditions, psychedelic drugs, and so on). In such conditions, you may use your cognitive abilities (e.g., your visual faculties) and you may even succeed cognitively by forming true beliefs against the odds—i.e., challenges C1 and C2 can be met in them. However, those conditions are such that one *typically* gets things right because of *luck*, not because of the exercise of one's abilities—i.e., challenge C3 cannot typically be met in them.³⁹

Let's consider the second factor: *masking circumstances*. Temp's circumstances—just like the archer's circumstances in Sosa's guardian angel case or when shooting metallic arrows at a giant magnetic target—are different: cognitive success is likely, to the point that it may be even *guaranteed*. This feature makes the circumstances *masking*, in the sense that no matter how unreliable one's belief-forming methods are, one will end up forming true beliefs. In such circumstances, a cognitive success that is produced by the exercise of cognitive ability does not manifest such an exercise. In a sense, that exercise is masked by the circumstances. Take Temp's case. Temp has the ability to trust reliable devices. Suppose that another agent, Tempo, lacks this ability and often trusts unreliable devices. In Temp's circumstances, it does not matter whether one uses Temp's or Tempo's method: the intervention of the hidden agent *guarantees* cognitive success no matter what. In this way, the circumstances prevent Temp's cognitive success from manifesting his normally reliable cognitive abilities by masking its exercise. In Sosa's guardian angel case, or in the case of the competent archer shooting metallic arrows at a giant magnetic target, the circumstances are also masking: it does not matter whether the archer makes competent or deliberately incompetent shots. Any shot is successful no matter what.

In conclusion, two relevant situational factors influence the difficulty of knowing: factors that make cognitive success unlikely and factors that, by making cognitive success likely, mask the manifestation of cognitive ability.

³⁹ Why 'typically'? Because, as we have seen, there are exceptions. This is the case of intellectual feats, such as some decisive breakthroughs in science.

Concluding remarks

In view of the foregoing discussion, we are now in a position to gloss the relationship between difficulty and knowledge as follows: how difficult it is for S to know that *p* at *t* in circumstances C corresponds to how likely S is to meet the challenge of forming a true belief in *p* at *t* in C *in such a way that the fact that S comes to believe p truly at t manifests S's cognitive abilities and effort up to t*. In other words, the difficulty of knowing is simply the difficulty of attaining a cognitive success that counts as a cognitive achievement in the relevant sense (a true belief that *manifests* cognitive ability and effort). This completes the explanation of the idea that knowledge can be difficult.⁴⁰

⁴⁰ There is a possible objection lurking around that might have crossed the mind of the interested reader: what about fake barns? In the perennial *fake barn case*, a driver crosses an area full of indistinguishable fake barns and is lucky enough to look at the only real barn in the area. The objection is that, while the driver ignores that there is a barn in front of him (at least this is what many epistemologists think), it is seemingly very easy for him to manifest cognitive ability: all he needs to do is to look at the barn and form the corresponding true belief that there is a barn in front of him. But why is that in the normal case in which there are no fakes around it is *as easy* to manifest visual competence and yet the driver would acquire knowledge as a result?

In reply, I have elsewhere argued (Broncano-Berrocal 2017; 2018; *forthcoming*) that this diagnosis of the fake barn case (and similar cases) rests on the false assumption, often made by virtue reliabilists, that manifesting cognitive ability is just a matter of manifesting *belief-forming* ability. As I have extensively argued, we also count with a range of abilities that do not play a belief-forming but a *precautionary role*. Whereas belief-forming abilities are abilities to form true beliefs reliably, what I call *precautionary cognitive abilities* are abilities that prevent the activation of one's belief-forming abilities in circumstances where it is unlikely that the latter will produce correct outputs (i.e., true beliefs). An example of a precautionary visual ability is the disposition not to form visual beliefs about the color of objects when the illumination is poor.

The crux of the matter is that in the fake barn scenario the driver *does manifest belief-forming ability* (after all, he counts with a reliable visual system that automatically produces the true belief that there is a barn when he sees it), but *fails to manifest precautionary visual ability*, in that he should have stopped belief formation. Why? Because the circumstances (*viz.*, the fake barn area) are such that too easily he would have formed a false belief. In other words, that is something that the driver should have done because it is a real possibility in his circumstances (unlike, e.g., the possibility of being a brain in a vat). Of course, he does form a true belief against all odds, but this is precisely why he fails to manifest *reliable* precautionary ability. Once knowledge is understood in terms of the manifestation of *all* epistemic resources (i.e., not only of the agent's belief-forming abilities and effort, but also her precautionary cognitive abilities), we can easily explain why the driver lacks knowledge: he only manifests part of the epistemic resources he ought to manifest.

In this way, knowing in circumstances where there are *visually indistinguishable* fakes is extremely difficult because our visual system is wired in such a way that we automatically take them to be real. Consequently, manifesting precautionary cognitive ability (i.e., manifesting the disposition to stop visual belief formation when cognitive success is unlikely) is something extremely difficult to accomplish in such circumstances. In such cases, in order to meet challenge C3—i.e., the challenge that we can now better understand as the challenge of succeeding cognitively because of *all* epistemic resources—we often need to resort to external aid to acquire the missing precautionary disposition. For example, in the fake barn situation the driver could use a fake barn detector: the information provided would reliably dispose him not to form beliefs about the presence of barns when in fact there are none. More mundanely, someone could tell the driver that he is in a fake barn area thus endowing him with the disposition to suspend judgment altogether. However, this sort of crucial information is not typically available in such unusual circumstances and this is the reason why it is so difficult to acquire knowledge in them: one needs to be epistemically cautious and one does not even know this to be the case.

The same considerations apply in general. In order to succeed and meet difficult challenges, sometimes the best we can do is to stop the current course of action, reassess and decide whether or not the current strategy is going to work. But this is not an easy task, especially if changing strategy is what we ought to do and we ignore it. In this sense, *reliably* applying caution constitutes an ability in itself, namely a precautionary ability—*cf.* von Kriegstein's interesting related discussion on wasted effort (von Kriegstein 2017).

References

- Alfano M., (2015). *Character as moral fiction*. Cambridge University Press.
- Baehr, J. S. (2011). *The inquiring mind: On intellectual virtues and virtue epistemology*. Oxford University Press.
- Baláš, J., Panáčková, M., Jandová, S., Martin, A. J., Strejcová, B., Vomáčko, L., Draper, N. (2014). The effect of climbing ability and slope inclination on vertical foot loading using a novel force sensor instrumentation system. *Journal of Human Kinetics* 44: 75–81.
- Bradford, G. (2015). *Achievement*. Oxford University Press.
- Broncano-Berrocal, F. (2017). A robust enough virtue epistemology. *Synthese* 194 (6).
- (2018). Purifying impure virtue epistemology. *Philosophical Studies* 175 (2):385-410.
- (forthcoming). Epistemic dependence and cognitive ability. *Synthese*.
- Carter, J. A. (2014). Robust virtue epistemology as anti-luck epistemology: A new solution. *Pacific Philosophical Quarterly*. doi:10.1111/papq.12040.
- Carter, J. A.; Kallestrup, J.; Palermos, S. O. & Pritchard, D. (2014). Varieties of externalism. *Philosophical Issues* 24: 63-109.
- Chi, M. T. H., Feltovich, P. J. and Glaser, R. (1981), Categorization and representation of physics problems by experts and novices. *Cognitive Science*, 5: 121–152
- Chrisman, M. (forthcoming). Epistemic normativity and cognitive agency. *Noûs*.
- Greco, J. (2009). Knowledge and success from ability. *Philosophical Studies* 142: 17–26.
- (2010). *Achieving knowledge*. Cambridge University Press.
- (2012). A (different) virtue epistemology. *Philosophy and Phenomenological Research* 85: 1–26.
- Guerrero, A. 2017. Intellectual difficulty and moral responsibility. In J. W. Wieland & P. Robichaud. *Responsibility: The epistemic condition*. Oxford University Press.
- Kelp, C. (2013). Knowledge: The safe-apt view. *Australasian Journal of Philosophy* 91: 265–278.
- Jarvis, B. (2013). Knowledge, cognitive achievement, and environmental luck. *Pacific Philosophical Quarterly* 94: 529–551.
- Olin, L. & Doris, J. M. (2014). Vicious minds. *Philosophical Studies* 168: 665-692.
- Pritchard, D. (2012). Anti-luck virtue epistemology. *Journal of Philosophy*, 109, 247–279.
- Pritchard, D. Haddock, A. & Millar, A., (2010). *The nature and value of Knowledge: Three investigations*. Oxford University Press.
- Sosa, E. (2007). *A virtue epistemology*. Oxford University Press.
- (2010). How competence matters in epistemology. *Philosophical Perspectives* 24: 465–475.
- (2015). *Judgment and agency*. Oxford University Press.
- (2017). Virtue theory against situationism. In M. Alfano & A. Fairweather (Eds.), *Epistemic situationism*. Oxford: Oxford University Press.
- Turri, J. (2011). Manifest failure: The Gettier problem solved. *Philosophers' Imprint*, 11, 1–11.
- (2017). Epistemic situationism and cognitive ability. In M. Alfano & A. Fairweather (Eds.), *Epistemic situationism*. Oxford: Oxford University Press.
- von Kriegstein, H. (2017). Effort and achievement. *Utilitas* 29: 27-51.
- (forthcoming). On being difficult: towards an account of the nature of difficulty. *Philosophical Studies*. doi.org/10.1007/s11098-017-1005-3.