

A ROBUST ENOUGH VIRTUE EPISTEMOLOGY

Fernando Broncano-Berrocal

(published in *Synthese*)

Abstract

What is the nature of knowledge? A popular answer to that long-standing question comes from robust virtue epistemology, whose key idea is that knowing is just a matter of succeeding cognitively—i.e., coming to believe a proposition truly—due to an exercise of cognitive ability. Versions of robust virtue epistemology further developing and systematizing this idea offer different accounts of the relation that must hold between an agent's cognitive success and the exercise of her cognitive abilities as well as of the very nature of those abilities. This paper aims to give a new robust virtue epistemological account of knowledge based on a different understanding of the nature and structure of the kind of abilities that give rise to knowledge.

Keywords: robust virtue epistemology; ability; cognitive ability; aptness; safety.

What is the nature of knowledge? A popular answer to that long-standing question comes from robust virtue epistemology, whose key idea is that knowing is just a matter of succeeding cognitively—i.e., coming to believe a proposition truly—due to an exercise of cognitive ability. Versions of robust virtue epistemology further developing and systematizing this idea offer different accounts of the relation that must hold between an agent's cognitive success and the exercise of her cognitive abilities as well as accounts of the very nature of those abilities. This paper aims to give a new robust virtue epistemological account of knowledge based on a different understanding of the nature and structure of the kind of abilities that give rise to knowledge.

To motivate these alternative accounts —of ability and knowledge— and to compare them with current views of the notion of cognitive ability as well as with the theories of knowledge that result from them, I will find it useful to introduce

and discuss Jennifer Lackey (2007, 2009) and Duncan Pritchard's (2012) infamous creditworthiness or attributability dilemma against robust virtue epistemology. The dilemma —as well as the two independent problems posed by it— will not only serve to show the shortcomings of rival views but also, and more importantly, to highlight the need of a systematic account of the notion of ability if one is the business of accounting for knowledge just in terms of cognitive ability.

The plan of the paper is as follows. In section 1, I will explain what robust virtue epistemology is and how it relates to performance normativity. In section 2, I will introduce Lackey and Pritchard's dilemma and explain what the real problem behind is. In section 3, I will discuss attempted solutions to the first of the two independent problems posed by the dilemma —the problem of accommodating epistemically relevant factors beyond cognitive agency— and argue that, once the relation between cognitive success and exercise of cognitive ability is properly understood, the problem vanishes. In section 4, I will discuss attempted solutions to the second problem —the problem of eliminating knowledge-undermining error possibilities— and argue that none is satisfactory because they give rise to new dilemmas concerning the safety condition for knowledge and fine-grained conceptions of the notion of cognitive ability. In section 5, I will propose an account of the notion of ability that distinguishes two sorts of abilities: abilities to complete tasks —or tasks-completion abilities— and abilities to prevent the use of task-completion abilities in inappropriate circumstances —or precautionary abilities. In section 6, I will build an original virtue epistemological account of knowledge on that basis and argue that the resulting view is superior to existing views in a number of aspects.

1 Robust Virtue Epistemology and Performance Normativity

Robust virtue epistemological accounts are typically based on the assumption that belief formation is a way to perform cognitively. The reason why they make that assumption is the following simple reasoning: (1) belief formation is a cognitive performance; (2) we commonly evaluate performances in general along certain normative dimensions; (3) cognitive performances are just kinds of performances; (4) then, we can plausibly evaluate belief formation along the same dimensions we use to evaluate performances in general in order to define the

concept of knowledge. Analyzing knowledge using performance normativity automatically gives a solution to the problem of identifying the source of normativity of the concept of knowledge: the normativity of knowledge is just performance normativity.¹

A popular framework in the literature is the AAA structure of performance normativity, introduced by Ernest Sosa, the founder of virtue and performance epistemology. According to Sosa (2007: 22-3), any performance —e.g., hitting a target— may be evaluated along three normative dimensions: *accuracy* —is the performance successful, i.e., does the performance attain its aim?, e.g., does the arrow hit the target?—; *adroitness* —is the performance the product of competence or ability?, e.g., has the arrow been delivered skillfully?—; and *aptness* —is the performance successful because of competence or ability, i.e., accurate because of adroit?, e.g., has the arrow hit the target because of being skillfully delivered —instead of being caused by a lucky gust of wind?²

In order for one's performance to count as apt, i.e., in order for the success of the performance to be because of one's abilities, a certain relation must hold between success and ability, namely the success must be *attributable* to the exercise of ability —as I use it here, the term “aptness” refers to the property that is instantiated whenever that relation holds. In the literature, the attributability relation and hence aptness is cashed out in two main ways:³

EXPLANATION APTNESS (EA): *S*'s successful performance is apt iff the exercise of *S*'s abilities *explains* the success of *S*'s performance.

MANIFESTATION APTNESS (MA): *S*'s successful performance is apt iff *S*'s success *manifests* *S*'s abilities.

EA simply says that the exercise of ability must figure in the total set of causal

¹ Initial virtue epistemological accounts of knowledge —e.g., Sosa (1991)— did not appeal to performance normativity. The advantage of using performance normativity is that one can bring together in the same normative framework not only an account of the nature of knowledge with an account of the source of its normativity, but also with an account of its value.

² The evaluative norms that are relevant are not moral. In other words, robust virtue epistemologists are not concerned with whether performances are right or wrong from a moral point of view but with whether they are good or bad performances *qua* performances. Torturing people is repugnant. But a torturer's performance might be assessed —with moral disgust— for accuracy, adroitness and aptness. For example, one can judge whether the torturer attains his atrocious goal of inflicting pain in his victim, whether he does so with skill, and whether his despicable success is due to that skill.

³ See Greco (2012) for an exhaustive review of all the current interpretations of the attributability relation.

factors that explain the success of the performance. As we will see, whether it should be the most important or just an important part of that set is a matter of controversy.⁴

On the other hand, manifestation of skill, competence or ability —i.e., the core idea of MA— is considered an intuitive primitive notion. One can grasp it by considering the minimal thesis that when an outcome manifests one’s abilities, the outcome is indicative of the exercise of ability in a way that shows that it has not been caused independently of it—in section 3, I will explain the notion of manifestation in more detail.⁵

Robust virtue epistemologists think that the problem of accounting for the nature of knowledge can be easily addressed by an application of the triple A structure of performance normativity to beliefs directly or to belief formation and derivatively to beliefs.⁶ The idea is to evaluate beliefs —the outputs of cognitive performances— using the same properties of performances in general: *accuracy* —is the agent’s belief true?—; *adroitness* —is the agent’s belief the product of her reliable cognitive abilities?—; *aptness* —is the agent’s belief true because of or attributable to the exercise of her reliable cognitive abilities? Given this normative framework, one way to define knowledge is as:

(AA) Accurate and adroit belief.

AA roughly amounts to views such as agent reliabilism —a development of process reliabilism.⁷ AA is a good start, but more is needed. It is widely acknowledged that views only based on reliabilist conditions do not suffice for

⁴ Greco (2010) upholds the stronger reading. Pritchard (2012) opts for the weaker one.

⁵ Authors advocating MA are Kelp (2013), Palermos (2013), Sosa (2011, 2015) and Turri (2011, *forthcoming*).

⁶ Sosa assumes that cognitive performances are beliefs, but that seems mistaken: while performances are dynamic, beliefs are stative —see Chrisman (2012) for this line of criticism. Instead of identifying cognitive performance with the doxastic state resulting from an instance of belief formation, cognitive performance seems to be better conceptualized as belief formation itself. By doing so, one can attribute the properties of being accurate, adroit or apt to the way agents form their beliefs and derivatively to their beliefs —although then the question is how can one account for their value. On a different note, belief revision might also count as a way to perform cognitively. For example, one can judge whether an agent has a competent belief in virtue of whether she has competently maintained her true belief in the presence of misleading defeaters.

⁷ See Greco (2000) for a seminal defense of agent reliabilism, according to which a true belief is justified and constitutes knowledge just in case it is the product of a reliable belief-forming process and the process is grounded on a stable disposition to form true beliefs that is integrated in the agent’s cognitive character. A predecessor of that view may be found in Goldman (1979), where he offers a classic defense of process reliabilism, according to which a true belief is justified —and thus a candidate for knowledge— just in case it is the product of a reliable belief-forming process.

knowledge because Gettiered beliefs might be the product of globally reliable belief-forming processes or cognitive abilities. Standard Gettier-style cases are cleverly structured in such a way that when the unfortunate Gettiered agent is about to form her belief using her reliable cognitive abilities and, when the almost already produced belief is about to be false due to a stroke of bad luck, a stroke of good luck corrects the situation making the resulting belief true.⁸ Epistemologists sometimes call the kind of knowledge-undermining epistemic luck involved in those cases *intervening luck*, because it is as if luck intervened in the way beliefs are formed.⁹ A more popular way to define knowledge is as:

(AAA) Apt belief.

Analyzing knowledge in terms of apt belief —whether interpreted in terms of EA or MA— successfully rules out standard Gettier-style cases as cases of knowledge. The reason typically given is that what general cases of intervening luck show is that a successful performance that is the product of ability —or adroit— might be such that its success is independent of —in the sense of not being attributable to— the exercise of ability, as when a skillfully delivered arrow hits the target because of a double gust of wind. When it comes to Gettier-style cases, the aptness condition on knowledge excludes intervening epistemic luck precisely because it links the relevant cognitive success to the exercise of cognitive ability by means of the attributability requirement.¹⁰

⁸ See Zagzebski (1994) for the view that standard Gettier-style cases have this “double luck” structure. See Blouw, Buckwalter & Turri (*forthcoming*) for a different interpretation of Gettier-style cases.

⁹ The term has been popularized by Pritchard (e.g., in Haddock, Millar and Pritchard 2010: 36). But see McKinnon (2014) for an argument to the effect that the terminology is not intelligible, although it remains a useful metaphor.

¹⁰ But see Greco (2012), who argues that inferential Gettier-style cases put pressure on AAA views precisely because the inferred beliefs meet that requirement. In those cases, the relevant Gettiered beliefs are true because of the reasoning abilities of the agent. Although I share the spirit of Greco’s criticism, I think that inferential cases need separate treatment, mainly because they involve belief-dependent cognitive processes —namely, inferences—, which means in turn that the lack of inferential knowledge might not have to do with a failure of the aptness condition in the transition from the premise belief to the conclusion, but with the premise belief’s poor epistemic status. My own view is that in order to account for inferential knowledge we need an aptness condition on evidence selection for the premise belief, and that condition fails in inferential Gettier-style cases. However, since the issue requires careful treatment, I will leave it for another occasion.

2 The Attributability Dilemma: The Real Problem Behind

Not all is good news for AAA theorists. They face a pressing dilemma due to Lackey (2007, 2009) and Pritchard (2012). The dilemma is known as the *creditworthiness* or *attributability dilemma* and is composed of two different types of cases:

NEW YORK

By asking for directions at a New York airport King comes to know the location of the Empire State from the testimony of Kong, a New Yorker.¹¹

FAKE BARNS

While traveling by train, Barnaby averts his eyes from the screen of his cellphone to admire the landscape of the barn county. Unbeknownst to Barnaby, a construction company recently demolished all but one of the many barns in the county and replaced them by barn façades. As luck would have it, the first object with the appearance of a barn he looks at is the only real barn in the county. He consequently forms the true belief that the object in front of him is a barn.¹²

The cases pose independent problems to AAA accounts. On the one hand, we commonly acquire testimonial knowledge in situations like NEW YORK. But AAA accounts seem to rule out those cases as cases of knowledge because the hearer's cognitive success —King's— does not seem to be solely attributable to his linguistic capacities or to his ability to select good informants but also to the speaker's cognitive abilities —Kong's—, such as his ability to retrieve the relevant information from memory.

On the other hand, AAA accounts seem to lead to the counterintuitive verdict that there is knowledge in cases of knowledge-undermining error possibilities such as FAKE BARNS. In FAKE BARNS, Barnaby's actual visual success seems attributable to the exercise of his visual capacities in the same way as in the normal perceptual case in which there are no fake barns around. But Barnaby's true belief does not amount to knowledge because, unlike in the normal case, Barnaby could easily have looked at a fake barn so that he would have formed a false

¹¹ The original case is in Lackey (2007: 352).

¹² The original case is in Goldman (1976: 772-3).

belief as a result —i.e., FAKE BARNS is a case in which the relevant error possibilities are knowledge-undermining. To put it in terms of epistemic luck, the reason why Barnaby’s true belief does not amount to knowledge is that it is *environmentally lucky*, in the sense that it is by luck that Barnaby has looked at the only real barn in the area and formed a true belief as a result.

The *attributability dilemma*, as presented by Pritchard, connects the two problems as follows: if AAA accounts are weakened to account for testimonial knowledge, then they cannot rule out cases like FAKE BARNS as cases of knowledge —*first horn*—, but if AAA accounts are strengthened to rule out cases like FAKE BARNS as cases of knowledge, then they cannot account for testimonial knowledge —*second horn*. In this way, the dilemma pulls AAA views in opposite directions.¹³

Before entering into further details, it is important to explain what the real problem behind the dilemma is as well as what kind of methodological approach the present project follows. At first sight, the main problem posed by the dilemma seems yet one more instance of the classic challenge of avoiding counterexamples that has driven post-Gettier epistemology so far. However, we should keep in mind the kind of views the dilemma is a challenge for: accounts that attempt to define knowledge *solely* on the basis of a proper understanding of both the notion of cognitive ability and of the relation between exercising cognitive ability and succeeding cognitively.

Therefore, what the dilemma shows to anyone in the business of analyzing knowledge solely in virtue-theoretic terms is the need of undertaking at least one of the following two non-mutually exclusive tasks: one should either revise one’s account of the notion of cognitive ability, or interpret the attributability relation in different terms, or both. In other words, the real problem behind the dilemma is a twofold problem: (1) settling what the correct understanding of the notion of cognitive ability is and (2) giving an adequate interpretation of the attributability

¹³ The way Lackey (2007, 2009) and Pritchard (2012) formulate the dilemma is different in two relevant respects. On the one hand, Lackey formulates it in terms of whether the agent *deserves credit* for her cognitive success, i.e., in terms of whether her cognitive success is praiseworthy —hence the name the *creditworthiness* dilemma. By contrast, Pritchard formulates it in terms of whether or not the agent’s cognitive success is *attributable* to her cognitive abilities —hence the name the *attributability* dilemma. On the other hand, for Lackey the first horn of the dilemma concerns Gettier-style cases in general, while for Pritchard it only concerns cases of environmental luck, such as FAKE BARNS. The reason why I use Pritchard’s presentation of the dilemma is that it makes more clear how the two independent problems relate to robust virtue epistemology and in particular to EA and MA.

relation.

While (2) seems to be a desideratum of virtue epistemology only, (1) is—or at least should be—of interest to epistemologists in general. The intuition that knowledge depends on exercising one’s cognitive faculties is very widespread both in and outside epistemology, so developing an adequate account of the nature and structure of the faculties that give rise to knowledge would not only be an achievement in virtue epistemology, but also in the theory of knowledge more generally.

The reason why (1) and (2) are so difficult to achieve—and this is exactly what the dilemma shows—is that whatever account of cognitive ability or of the attributability relation one puts forward, it should be compatible with there being *factors beyond one’s cognitive agency*—such as a speaker—contributing in crucial way to the acquisition of knowledge, while at the same time it must help eliminate *knowledge-undermining error possibilities*—or environmental luck—in circumstances in which one exercises one’s cognitive faculties in seemingly the same way that typically makes one gain knowledge.

My hypothesis—and here comes the methodological approach of the paper—is that we do not really need to come up with new ways of interpreting the attributability relation.¹⁴ MA is a sufficiently satisfactory reading, not only because it helps avoid counterexamples, but also because it makes very easy to think of abilities in a familiar way: the way we think of dispositions in general. Consequently, what seems to be in need of rigorous philosophical work is the very notion of cognitive ability.

More specifically, it would be very positive to count with a fully-fledged account of the notion of ability from which one can derive a corresponding account of the notion of cognitive ability and of the functional roles that different dispositions play in the acquisition of propositional knowledge. The idea is to combine the results of that account with a plausible reading of the attributability relation—MA—in an aptness condition on knowledge. Hopefully, the resulting robust virtue epistemological account of knowledge will be able not only to solve the infamous attributability dilemma, but also to predict common intuitions about familiar cases in epistemology.

¹⁴ Greco (2012), for example, reinterprets the attributability relation in pragmatic terms.

3 Epistemically Relevant Factors Beyond Cognitive Agency

The first step towards that account of knowledge is to show how robust virtue epistemology can easily accommodate epistemically relevant factors beyond cognitive agency simply by reading the attributability relation and hence the aptness condition in terms of MA. Consider again what the problem is: the core idea of robust virtue epistemology is that knowing is just a matter of exercising one's cognitive abilities, but how can that be reconciled with there being factors that have nothing or little to do with one's cognitive abilities and that play a crucial role in the acquisition of knowledge? In what follows, my focus will be on cases of testimonial knowledge because they constitute an excellent example of knowledge acquisition partly in virtue of factors that are beyond the knower's cognitive agency—namely, a speaker. But similar points could be made concerning cases of knowledge from instruments.¹⁵

3.1 Attempted Solution I: Weakening the Explanatory Reading

In cases of testimonial knowledge, the hearer's cognitive abilities do not *fully* explain her cognitive success, because the speaker's cognitive abilities also play an important role in that explanation. That indicates that the problem of accommodating epistemically relevant factors beyond cognitive agency is especially worrisome for AAA theories based on EA, the explanatory reading of aptness.

The first solution to the problem consists in weakening EA so that testimonial beliefs can count as apt. This route is followed by Pritchard (2012), who interprets EA so that it requires that the agent's cognitive abilities play a salient or an important role in the explanation of her cognitive success, but not necessarily the *most* salient or important role. In testimony cases more specifically, Pritchard thinks that the fact that hearer exercises her ability to select reliable informants—together with her linguistic capacities—is sufficient for her cognitive success to count as apt. The reason why he thinks this is that, although the hearer's linguistic and informant-selecting abilities are certainly not the most salient factors in the explanation of why she succeeds cognitively—considering the role played by the speaker in passing accurate information—, they still play an im-

¹⁵ See Pritchard (2010) for relevant discussion.

portant role —e.g., in the target cases of testimonial knowledge, the hearer does not select her interlocutor by random.

A so weak interpretation of EA obviously leaves the aptness condition in an even worse position than the stronger reading when it comes to handling cases of knowledge-undermining error possibilities such as FAKE BARNS. To amend the problem, Pritchard opts for introducing an independent non-virtue-theoretic condition. In particular, he opts for introducing a safety condition, i.e., a condition to the effect that knowledge requires that one must not form true beliefs that could easily have been false.

The problem with that move is, firstly, that the resulting view —called anti-luck virtue epistemology— is no longer purely virtue-theoretic, i.e., it renounces to explain knowledge only in terms of cognitive ability, which is obviously not an option for robust virtue epistemologists. Secondly, the view is insufficient for knowledge precisely because its ability condition is too weak. As I argue in Broncano-Berrocal (2014a), there are plausible counterexamples in which (1) one’s belief is safe, (2) one’s cognitive abilities are a salient factor in the explanation why one comes to believe a true proposition —and even in the explanation of why one’s true belief is safe—, but (3) one still does not know.¹⁶ Thirdly, Pritchard’s anti-luck virtue epistemology falls prey to the dilemma that I will propose in section 4.2.

3.2 Attempted Solution II: Social Competences

An alternative solution to the problem of accommodating epistemically relevant factors beyond cognitive agency acknowledges that the hearer’s cognitive success is not solely attributable to her individual cognitive abilities. With that in place, it goes on to argue that, unlike perception, memory or reasoning, testimony is a social source of knowledge, which means that the kind of abilities or competences that give rise to testimonial knowledge are not individual but social. In particular, the proposal is that in cases of testimonial knowledge what is exercised is a social ability or competence that is partly *constituted* by the individual cognitive abilities of the hearer and partly by the individual cognitive

¹⁶ In such cases, the safety of one’s beliefs is partly explained by the unnoticed intervention of an agent —which triggers the ignorance intuition— and partly by fact that one is able to eliminate defeaters that would easily have made one believe false propositions —which is sufficient to make the conditions of anti-luck virtue epistemology obtain.

abilities of the speaker. Testimonial knowledge is accordingly explained in the following way: when H comes to know that p from the testimony of S , H 's cognitive success is *fully* explained by the social competence that is seated on H 's and S 's individual cognitive abilities.¹⁷

The problem with explaining testimonial knowledge in terms of socially seated competences is that testimonial knowledge from uncooperative speakers is fairly common, but knowledge in those cases does not seem to be explained by the exercise of a joint competence. Suppose that S is a compulsive liar and that H is reasonably competent at detecting lies. After scrutinizing S 's facial expressions for a while, H detects cues indicating that S is lying, so that H non-inferentially believes—and comes to know—that what S is saying is false.¹⁸ It does not seem that H 's testimonial knowledge can be explained by an exercise of a social competence composed of H 's and S 's individual cognitive abilities. After all, all of S 's efforts are directed towards deceiving H , which means that S 's cognitive abilities—if any—play no role in the explanation of H 's cognitive success. Yet, H acquires knowledge from testimony.¹⁹

3.3 A Simple Solution

There is a simple solution to the problem of accommodating epistemically relevant factors beyond cognitive agency, one that neither leads to the introduction of non-virtue-theoretic conditions nor to a rather controversial revision of the notion of cognitive ability. The solution consists in simply adopting MA, i.e., the idea that an apt—cognitive or non-cognitive—performance is a performance whose success *manifests* the agent's—cognitive or non-cognitive—abilities. Before spelling out the solution, it is important to pause for a moment and explain how to understand MA. The explanation will not only serve to motivate the promised solution, but also to avoid misunderstandings concerning the account of ability offered in section 5.

¹⁷ See Sosa's 2007 book (*passim*) for this kind of solution, where he gives a similar explanation of knowledge from instruments.

¹⁸ Readers who are familiar with gambling games involving deception—e.g., poker—will agree that this sort of lie-detection abilities are not only possible but real.

¹⁹ It might be argued that cases of uncooperative speakers are not genuine cases of testimony, but whatever way to individuate testimony one favors, it should be compatible with there being cases of easy and difficult acquisition of testimonial knowledge. That is, the fact that there is more effort on the part of the hearer in a given case should not prevent the case from counting as a genuine case of testimony.

MA explains aptness in metaphysical terms —instead of explanatory— by means of the manifestation relation between success and ability. As stated earlier, manifestation is considered a primitive notion, in the sense that it remains undefined, but certainly not unexplained. To understand better what is for a success to manifest ability, it is useful to consider dispositions in general —abilities in epistemology are typically thought to be dispositions of agents.

Firstly, the sense in which the *unsuccessful* shot of a skillful archer *does not manifest* her shooting abilities when shooting in a hurricane is the same sense in which salt does not manifest its solubility when stirred into gasoline —salt does not dissolve in gasoline. Salt manifests its solubility only if the solubility conditions are normal or appropriate. In the same way, an agent's success —e.g., hitting a target— manifests her abilities only if the agent's *internal conditions* and the *environmental conditions* in which the success obtains are normal or appropriate —i.e., in the same way as salt manifests its solubility when stirred into water, the *successful* shot of a skillful archer *manifests* her shooting abilities when shooting in good shape and normal wind conditions. In a slogan, aptness —when understood according to MA— requires appropriateness of circumstances.

Secondly, does this mean that all successes in normal circumstances automatically manifest ability when the agent and the performance are in fact skillful? To put it differently, is the appropriateness of the circumstances sufficient for aptness? No, there might be cases of normal circumstances in which a successful and skillful performance is inapt —i.e., not successful because of skillful. For example, the successful shot of a skillful archer does not manifest her shooting abilities if after the skillful delivery of the arrow a double gust of wind deviates the arrow and puts it back on track thus making the arrow hit the target. One could argue that the case is just another example of inapt performance due to inappropriate circumstances. The problem with that idea is that if we do not allow minimal risky variations within normal circumstances we might end up with the view that normal or appropriate circumstances are the set of invariant circumstances that guarantee success. The defender of that view, however, would have a hard time accounting for the fallibility of human abilities. A more reasonable view of appropriateness has it that shooting in a hurricane does not count as appropriate circumstances but allows for double gusts of wind to be part of the normal range of action of an archer's shooting ability.

Explaining cognitive abilities and their manifestation by cognitive successes

using the model of dispositions and their manifestations by certain outcomes is theoretically fruitful when it comes to epistemology. In particular, it provides a simple solution to the problem of accommodating epistemically relevant factors beyond cognitive agency.

The solution starts by adopting MA as the correct reading of the attributability relation and the dispositional model as the correct approach to conceptualizing cognitive abilities.²⁰ One thing that the literature on dispositions teaches us is that a certain outcome might manifest two or more different dispositions without that making the dispositions involved be less manifested. Think about salt. The solubility of salt is no less manifested when stirred into water because the water of the glass also manifests its disposition to dissolve salt. Indeed, it is a common idea in the literature on dispositions that dispositions *need* reciprocal disposition partners to be manifested.²¹ In order to be both manifested, salt needs the disposition of water to dissolve salt and water reciprocally needs the disposition of salt to be dissolved in water. Therefore, keeping the analogy between dispositions and abilities in mind, it is not unnatural to claim that a cognitive success might manifest more than one set of cognitive dispositions or abilities.

That is precisely the case of testimonial knowledge. The cognitive success of the hearer not only manifests her ability to select reliable informants —e.g., to select people not looking like tourists— but also the reciprocal cognitive abilities of the speaker, whatever they are. In the same way as the manifestation of the disposition of water to dissolve salt does not make the solubility of salt be less manifested when stirred into water, the intervention of a speaker in a given conversation does not make the hearer's cognitive abilities be less manifested when competently asking for information.

That is compatible with there being cases of testimonial knowledge from uncooperative speakers. In those cases, the kind of ability exercised by the speaker, *S*, is not the usual ability that a cooperative speaker exercises when willingly passing accurate information, but an ability to lie. Does that mean that the cognitive abilities of the hearer, *H*, lack a reciprocal disposition partner? No, be-

²⁰ See Fara (2008), Greco (2012) and Vihvelin (2004) for dispositionalist accounts of the notion of ability, which contrast with the classical conditional analysis according to which an agent is able to φ just in case she would φ if she tried or chose to. See Maier (2010) for a different account of abilities in terms of having the option to φ .

²¹ See Heil (2005: 350).

cause as stated earlier, those are cases in which *H* acquires knowledge *because* she is competent at detecting lies. Accordingly, *H*'s cognitive success manifests her ability to detect lies. But it would be odd to say that *H* manifests an ability to detect that her interlocutor is lying if there is no interlocutor lying to *H*. That means that *S*'s disposition to lie is the reciprocal disposition partner of *H*'s ability to detect lies. The difference with cases of cooperative testimony is that, while in the latter *H*'s and *S*'s individual cognitive abilities are manifested by *H*'s cognitive success, in cases of uncooperative testimony *H*'s cognitive success only manifests her disposition to detect lies. In order for *S*'s disposition to lie to be manifested *H* needs to come to believe something false.

In conclusion, (1) interpreting the attributability relation in terms of manifestation seems to be better than in explanatory terms and (2) understanding cognitive abilities along dispositionalist lines seems to be better than introducing social competences. The simplicity with which the cases are explained provides excellent motivation for making these two theoretical choices.

4 Knowledge-Undermining Error Possibilities

Knowledge is commonly thought to have a modal dimension. A common way to exemplify the idea is by means of cases in which close possibilities of error undermine actual knowledge, as it happens in the widely discussed FAKE BARNS. Different accounts in the epistemological literature try to capture the modal dimension of knowledge using different theoretical resources. But how can AAA-based accounts explain the intuition of ignorance in those cases? After all, the relevant cognitive abilities involved are exercised in seemingly the same way that typically produces knowledge. A natural conclusion one can draw is that robust virtue epistemology lacks the theoretical resources to capture the modal dimension of the concept of knowledge. Another way to phrase the worry is to point out that since robust virtue epistemology fails to exclude the kind of luck instantiated in FAKE BARNS and cases alike, it does not constitute an adequate anti-luck epistemology.

The literature on virtue epistemology counts with several solutions to the problem, but they basically boil down to two strategies: (1) introducing a safety condition in virtue-theoretic terms —i.e., as a modalized ability condition— and (2) revising the notion of cognitive ability so that agents lack the relevant cogni-

tive abilities —or do not exercise them— in the problematic cases. I will discuss these two strategies in turn and argue that the most promising is the latter.

4.1 Attempted Solution I: Safety as an Ability Condition

John Turri (2011, *forthcoming*) follows the strategy of introducing safety as an ability condition. In particular, he uses MA to extend Sosa’s AAA framework with another normative property: *amplitude*. According to Turri, ample performances are not only safe —i.e., successful in the actual and in close possible worlds—, but also their *safety* manifests the agent’s abilities. Imagine a successful shot by an archer. Ask: does the arrow hit the target because of being skillfully delivered —instead of being caused by a lucky gust of wind— and in close possible worlds the arrow would still hit the target because of being so delivered? If the answer is yes, the shot is ample. The idea then is to define knowledge as:

(AA⁺) Ample belief.

Amplitude is a modalized version of aptness. When a success is ample, not only in the actual but also in close possible worlds the success manifests the exercise of ability —i.e., it is apt. Most cases of apt performance are also cases of ample performance because the relevant success and the skillful way to achieve it typically stay stable across close possible worlds —e.g., consider your present apt and ample belief that you are right now reading “right now reading”. In other words, aptness and amplitude tend to coincide because most of the times we just happen to be in friendly circumstances.

But there are cases in which amplitude and aptness come apart. For example, an archer might skillfully hit the only target in the field that, unbeknownst to her, is not protected with invisible force fields that would repel incoming arrows.²² Given Turri’s understanding of amplitude, the shot is not ample because it is unsafe. The same reasoning applies to cases of knowledge-undermining error possibilities such as FAKE BARNS: in those cases the relevant beliefs are not ample because of being unsafe.

Sosa (2015) also opts for introducing a safety condition in virtue-theoretic terms, but in a less straightforward way than Turri. In keeping with his previous work (2007, 2009, 2011), Sosa distinguishes between *animal knowledge* —i.e.,

²² The example is from Pritchard (2006: 15).

apt belief—, *reflective knowledge* —i.e., meta-apt belief or apt belief aptly noted—, and a better quality of knowledge: *knowledge full well*. According to Sosa (2015), in cases of environmental luck such as FAKE BARNS agents have animal knowledge, might even have reflective knowledge, but fail to know full well. Let us reconstruct Sosa’s treatment of the case.

As we have seen, Barnaby’s first-order belief —the belief that there is a barn in front of him— is apt: it is true because of an exercise of cognitive ability. In addition, Barnaby’s environment does not prevent him from having a meta-apt belief, namely an explicit second-order belief or an implicit presupposition consisting in an apt judgmental affirmation that his first-order belief is apt. In order for Barnaby’s second-order belief or presupposition to be apt —and hence in order for his first-order belief to be meta-apt— it must be true because of an exercise of cognitive ability, namely of reflective competence —Sosa’s requirement for reflective knowledge is in this sense a second-order aptness requirement.

Therefore, Barnaby has animal knowledge and can even have reflective knowledge if Barnaby’s second-order belief or presupposition is true because of reflective competence. This result helps Sosa explain the widespread intuition that in cases of environmental luck such as FAKE BARNS the true beliefs in question are epistemically better off than the true beliefs of standard Gettier-style cases, i.e., than beliefs that are true because of intervening luck: only the former are apt and even meta-apt, and in that sense they have a better positive epistemic status.

What Barnaby definitively lacks, according to Sosa, is knowledge full well. The idea is that, although Barnaby’s second-order apt belief entails that his first-order belief *is* apt, it does not entail that it *would* also reliably or *likely enough* be apt, or to put it in terms of close possible worlds, that it would continue to be apt *in close possible worlds*. Fulfilling that safety requirement at the second-order level is what is needed for knowing full well when one already has animal and reflective knowledge. More specifically, knowing full well requires full-well belief, namely:

(FW) Second-order apt belief or presupposition that one’s first-order belief is safely apt —i.e., that it would not easily be inapt.

Since in close possible worlds Barnaby’s first-order belief would not be apt, Barnaby’s second-order belief that his first-order belief is safely apt is false. This result

helps Sosa explain the intuition that a relevant sense of knowledge is missing in cases of knowledge-undermining error possibilities or environmental luck.

4.2 The Safety Dilemma

In this way, a virtue of Turri's and Sosa's respective solutions to the problem of eliminating knowledge-undermining error possibilities —or environmental luck— is not only that they are able to account for the intuition that Barnaby fails to know in a relevant sense, but also for the widespread intuition that environmentally lucky beliefs are better off than beliefs that are true because of intervening luck. However, introducing safety requirements in virtue-theoretic terms leads to an even more pressing dilemma for robust virtue epistemology than the attributability dilemma.

When one takes cases of environmental luck at face value, there is a strong temptation to endorse solutions that make use of safety principles. After all, cases of environmental luck are paradigmatic cases of unsafe belief—all ample and full-well beliefs are safe. The problem is that not all close error possibilities are knowledge-undermining, and if one's safety principle is unable to discriminate between epistemically harmless and epistemically harmful error possibilities, one runs the risk of undesirably ruling out good cases as cases of knowledge. In fact, cases involving epistemically harmless error possibilities have been used to argue that safety is not a necessary condition for knowledge. Consider two cases in the literature:

DACHSHUND

Snoop knows that the animal before his eyes—a dachshund—is a dog despite the fact that he could easily have encountered one of the many wolves in the area that he would have erroneously taken to be dogs.²³

CLOCK

Mr. Punctual knows that it is three o'clock by looking at the most accurate clock on Earth despite the fact that an undecayed isotope could easily have decayed disrupting the clock's internal mechanism

²³ The original case is in Goldman (1976: 779).

thus upsetting Mr. Punctual for forming false beliefs about the time.²⁴

Amplitude does not hold in DACHSHUND and CLOCK, which are cases of knowledge, because in close possible worlds where Snoop forms beliefs about the presence of dogs or Mr. Punctual about the time, the beliefs are false. Therefore, it is not true that knowledge is true belief whose safety manifests cognitive ability.

The same reasoning applies to full-well belief. In DACHSHUND and CLOCK, the relevant first-order beliefs, although true and apt in the actual world, would fail to be apt —because of being false— in close possible worlds. Therefore, the corresponding second-order beliefs or implicit presuppositions are false.

Sosa's reaction could be that this treatment of the cases misconstrues his requirement that full-well belief entails that the relevant first-order beliefs would reliably or likely enough be apt. In particular, he could argue that the requirement only entails actual reliability or likelihood of truth relative to the actual situation, not to close possible worlds, so that DACHSHUND and CLOCK are cases of knowledge full well. But if this is the correct way to understand the notion of full-well belief, why would we keep saying that Barnaby fails to know full well —i.e., that he fails to achieve the highest quality of knowledge possible? After all, Barnaby's *actual* situation is just as good as a normal perceptual case in which the object of perception is right in front of him: relative to the actual situation, everything is epistemically all right.

If we generalize the previous rejoinder to Sosa's possible reply, we can easily notice why introducing safety as an ability condition leads to a subtler dilemma for robust virtue epistemology than the attributability dilemma. Call it *the safety dilemma*: if one's robust virtue epistemological view does not make safety a necessary condition for knowledge, then the view is too weak to rule out cases of knowledge-undermining error possibilities —e.g., FAKE BARNS— as cases of knowledge —*first horn*—, but if it does, then it is too strong to account for cases of knowledge involving epistemically harmless error possibilities such as DACHSHUND or CLOCK —*second horn*. In this way, the dilemma pulls safety-based robust virtue epistemology in opposite directions. Accounts that introduce safety as an independent non-virtue theoretic condition on knowledge —e.g.,

²⁴ The original case is in Bogardus (2014: 300-1).

Pritchard's anti-luck virtue epistemology— also fall prey to dilemma.²⁵

A better strategy —fully in the spirit of robust virtue epistemology— consists in understanding the notion of cognitive ability in a way that epistemically harmful and harmless cases of error possibilities respectively correspond to cases in which the agents in question lack and have the relevant cognitive abilities. As we will see next, achieving that result using the notion of ability as the only theoretical resource to maneuver with is anything but easy.

4.3 Attempted Solution II: Fine-Grained Cognitive Abilities

Following the strategy of revising the notion of cognitive ability, John Greco (2010) conceives cognitive abilities in a fine-grained way. More specifically, Greco narrows down the range of environments relative to which we have the ability to discriminate object *A* from object *B*. For example, if *A* is a barn, one has the ability to recognize barns in environments in which *B* is a skyscraper, but not in environments in which *B* is a barn façade. The fact that Barnaby lacks such an ability explains why he does not know in FAKE BARNS.

In a similar fashion, Alan Millar (in Haddock, Millar and Pritchard 2010) also opts for narrowing down the range of our cognitive abilities, but instead of explaining cases such as FAKE BARNS in terms of failing to possess the required ability, he explains them in terms of failing to exercise it. In such environments, one fails to exercise what he calls a perceptual-recognitional ability —i.e., an ability to tell of an object that is of a certain kind from the way it looks, sounds, and so on— because environments with fakes, Millar argues, are not favorable to things being recognized by the way they appear to be. Crucially and derived from that, Millar thinks that one does not exercise perceptual-*recognitional* ability unless one thereby acquires knowledge. This requirement turns Millar's view into a knowledge-first view.²⁶

²⁵ A possible way out consists in reinterpreting one's safety condition as method-safety and then proposing a principle for individuating methods of belief formation that restricts safety to close possible worlds in which there are no dog-looking wolves (DACHSHUND) or broken clocks (CLOCK). I have tried that strategy elsewhere in order to defend safety from counterexamples to its necessity (see Broncano-Berrocal 2014b). But although I think that it is the best strategy a safety theorist and therefore a defender of safety-based robust virtue epistemology can adopt, giving an adequate principle for method individuation proves to be extremely difficult. See Bogardus and Marxen (2014) for some objections.

²⁶ Most robust virtue epistemologists find that move unwelcome, as they still think that there is hope for a reductive analysis of knowledge.

4.4 The Ability Dilemma

Greco and Millar's approach has found some resistance in the literature. For example, Pritchard (2012) argues that the default way to think of abilities is as relativized to a wide range of environments —i.e., in a coarse-grained way—, and only if there is special reason to do so, we relativize abilities to specific environments —i.e., in a fine-grained way.

To add some extra force to Pritchard's worries, it is not clear how a fine-grained relativization to environments is going to help to account for cases of epistemically harmless error possibilities such as DACHSHUND or CLOCK. For example, if —following Greco— the reason why Barnaby does not know in FAKE BARNS is that he lacks the ability to tell barns from fake barns, then why does Snoop have knowledge in DACHSHUND if he lacks the ability to tell dogs from wolves?

Millar's way out could be simply to point to the fact that, unlike in FAKE BARNS, in DACHSHUND there is knowledge. But how explanatory would be to appeal to the mere absence or presence of knowledge to tell the difference between the cases? That one favors a knowledge-first view —as Millar does— does not mean that one is exempt from giving an informative explanation of what makes widely discussed cases different. To do so, knowledge-first theorists have to appeal to necessary conditions for knowledge —that is allowed insofar as they remain skeptic about there being non-circular jointly sufficient conditions for knowledge. But what could those necessary conditions be? If the answer is safety, one goes back to the safety dilemma.

A similar dilemma serves to illustrate Millar and Greco's delicate position. Call it *the ability dilemma*: if one's robust virtue epistemological view does not understand cognitive abilities in a fine-grained way, then the view is too weak to rule out cases of knowledge-undermining error possibilities —e.g., FAKE BARNS— as cases of knowledge —*first horn*—, but if it does, then it is too strong to account for cases of knowledge involving epistemically harmless error possibilities such as DACHSHUND or CLOCK —*second horn*. In this way, the dilemma pulls views endorsing a fine-grained individuation of cognitive abilities in op-

posite directions.²⁷

We have thus arrived at the following dialectical situation: simple AAA accounts are too weak to eliminate knowledge-undermining error possibilities. Introducing safety as an ability condition (AA⁺, FW) results in too strong views, namely views that also exclude knowledge in cases of epistemically harmless error possibilities. Strengthening AAA accounts with fine-grained conceptions of cognitive ability leads to the same problem.

What we need is an account that is in-between simple AAA views and accounts based on AA⁺ and FW on the one hand, and AAA accounts based on fine-grained abilities on the other. In other words, we need a version of robust virtue epistemology that is neither too weak nor too strong, just *robust enough*. I will give that account in the next two sections. Like Greco and Millar, I will follow the strategy of revising the notion of cognitive ability. In this sense, the proposed account will be fully in the spirit of robust virtue epistemology. Unlike Greco's and Millar's respective views, my view will be compatible with there being coarse- and fine-grained cognitive abilities.

5 An Account of the Notion of Ability

5.1 Abilities and Tasks

Pritchard is right in thinking that we generally adopt a coarse-grained stance towards abilities, but —consistently with what he says— that does not mean that fine-grained abilities do not exist. The question is: how fine-grained should cognitive abilities be? No knockdown argument will prevent a robust virtue epistemologist from analyzing the notions of ability and cognitive ability either —exclusively— in a coarse-grained or in a fine-grained manner. But any such an attempt to settle the question is probably going to encounter problematic cases in the opposite direction the attempt intends to favor. For that reason, a better strategy is to avoid taking sides by making one's preferred account of the notions of ability and cognitive ability compatible with both theoretical options. An ac-

²⁷ In Millar's case, the dilemma goes as follows: if one's knowledge-first view does not understand cognitive abilities in a fine-grained way, then it is not clear what necessary condition for knowledge fails to obtain in cases of knowledge-undermining error possibilities such as FAKE BARNS —*first horn*—, but if it does, then the view cannot explain the difference between those cases and cases of epistemically harmless error possibilities such as DACHSHUND or CLOCK —*second horn*.

count following that strategy will be given in a moment. First, it will be helpful to introduce a natural way to think of abilities and explain on that basis why different theorists have different intuitions about the range of action of our abilities. This will also help motivate and introduce the promised account.

By definition, abilities are abilities to do something, namely *tasks*. The tasks that abilities serve to complete might have different degrees of complexity often involving several —many— subtasks. Katherine Hawley (2003) sheds some light on the issue. The external and internal conditions of an agent, she explains, often serve to identify the relevant subtasks related to the task generally associated to an ability. Consider one of Hawley’s examples: the ability to walk up the stairs. This general ability might involve different tasks depending on what type of conditions the conditions of the agent are —e.g., walking up the stairs with a broken leg, in the dark, in high heels or backwards. In view of that variety of tasks, Hawley thinks that a plausible way to conceive abilities is as involving “families” of tasks where the family members are tokens of “doing the same thing” —e.g., getting upstairs, driving, cooking— relative to different types of internal and external conditions or parameters.²⁸ Therefore, a plausible and intuitive *criterion for individuating abilities* is by the tasks they serve to complete in certain types of internal and external conditions or relative to certain types of parameters.

Hawley anticipates that we will not find a sharp criterion for what counts as the same task and what does not, because that would require a sharp distinction between what counts as the same conditions or parameters and what does not. Something we can say —somehow vaguely— is that two tasks are of the same type just in case the conditions or parameters that define them are roughly the same. For example, an archer’s token task of hitting an apple off someone’s head when riding a bike is the same type of task as a token task of hitting an object with the same properties —e.g., same shape, same size, and so on— when riding a similar bike in the same conditions —e.g., same distance, same wind and light conditions, and so on. If we also think that abilities are dispositions, we can define abilities as dispositions to complete tasks of a certain type successfully. We can call abilities so conceived *task-completion abilities*.

Once we start thinking of abilities in this way —as individuated by the tasks

²⁸ Relevantly, Hawley also remarks that “although some tasks in a family may be more difficult than others, (...) not every family is neatly ordered in this way” (Hawley 2003: 21).

they serve to complete in certain internal and external conditions— we can explain why different commentators have different intuitions concerning how coarse- or fine-grained abilities should be. Since, for any given case, the criterion for individuating abilities is silent on what is the correct —if any— level of generality of the type of tasks that individuate the relevant ability in that case, whatever choice one makes will determine how coarse- or fine-grained one takes the ability to be. For example, if one thinks that the type of cognitive task in FAKE BARNS is to discriminate barns from fakes —i.e., objects instantiating the properties often associated to barns—, the resulting cognitive ability will be much more fine-grained than if one thinks that the type of cognitive task is simply to discriminate barns from buildings with very different properties, such as skyscrapers. Finally, an even more fine-grained way to individuate an ability is by the token task it serves to complete in the circumstances —e.g., the ability to discriminate *that* barn from *those* other barn-looking objects. The general conclusion is that an ability will be more or less fine-grained depending on what type of task or token task one thinks the ability serves to complete.²⁹

5.2 Task-Completion Abilities

Let us take a more systematic look at this alternative way to conceptualize abilities in terms of tasks. A task-completion ability can be defined more precisely as follows:

ABILITY: *S* has an ability to complete tasks of type *T* iff *S* has a disposition to complete tasks of type *T* when in certain internal and environmental conditions, with degree of reliability *D*, where *D* is above a certain threshold.³⁰

Cognitive abilities are defined in the same way. The only difference is that the relevant tasks are cognitive, such as the task of coming to believe the truth in a field of propositions or the task of continuing to believe what is true when one already believes a true proposition.

²⁹ The distinction that I will make between broad and narrow abilities will mirror this discussion.

³⁰ The degree and threshold of reliability of an ability obviously depends on the kind of tasks the ability serves to complete and the domain to which the ability pertains. See Greco (2012) for a similar but relevantly different definition of the notion of ability.

Not all conditions are appropriate for completing tasks of type *T*. Very roughly, in the case of archery appropriate internal conditions include having the physical constitution required for shooting arrows, being in good shape, sober, awake, and so on, while appropriate environmental conditions include a working bow, arrows, normal winds, targets at a proper distance, good light conditions, and so on. Also very roughly, in the case of visual beliefs appropriate internal conditions include a working visual apparatus, being awake, sufficiently attentive, and so on, while appropriate environmental conditions include objects at a proper distance, good light conditions, and so on. We can talk of the *normal range of action* of an ability as follows:

NORMAL RANGE OF ACTION: the normal range of action of *S*'s —cognitive or non-cognitive— ability to complete tasks of type *T* is the set of appropriate conditions relative to which *S*'s ability is reliable to degree *D*, where *D* is above a certain threshold.

We can distinguish three types of internal and environmental conditions:³¹

LOCAL INTERNAL AND ENVIRONMENTAL CONDITIONS: the internal and environmental conditions in which *S* is when *S* completes a certain task of type *T* at *t* —roughly: the actual world.

REGIONAL INTERNAL AND ENVIRONMENTAL CONDITIONS: the internal and environmental conditions *S* might easily have been at *t* or at times close to *t* —roughly: close possible worlds.

GLOBAL INTERNAL AND ENVIRONMENTAL CONDITIONS: the internal and environmental conditions in which *S* typically is when completing tasks of type *T*.

Note that *local* and *regional conditions* might or might not be part of *global conditions*. The motivation for introducing these distinctions is that they allow to distinguish two types of abilities, *narrow* and *broad abilities*:

NARROW ABILITY: *S* has a narrow —cognitive or non-cognitive— ability to complete a certain token task iff *S* has a disposition to complete that task only in certain local and regional internal and environmental conditions, with degree of reliability *D*, where *D* is above a certain threshold.

³¹ I borrow the terminology from Kallestrup & Pritchard (2014), but I modify their distinctions slightly.

Narrow abilities are less rare than one might think. We can all imagine drugs that temporarily enable people to overcome their shyness and give speeches in public, that increase their strength or endurance substantially when it comes to performing certain physical tasks, or that temporarily fix dysfunctional medical conditions.³² But the standard way to think of abilities corresponds more to the following:

BROAD ABILITY: *S* has a broad ability to complete —cognitive or non-cognitive— tasks of type *T* iff *S* has a disposition to complete tasks of type *T* in a set of global internal and environmental conditions, with degree of reliability *D*, where *D* is above a certain threshold.

The normal range of action of many of our broad abilities —especially the cognitive ones— has been shaped by evolution for the most part. One just has to think about how human prehensile thumbs have evolved to grab objects of a certain size. On the other hand, the normal range of action might correspond to the set of learning conditions. For example, children learn how to hit baseballs under normal wind conditions, but not during hurricanes. Plausibly, many broad abilities are such that learning processes combine with natural selection to shape the relevant set of appropriate global conditions.³³

5.3 Precautionary Abilities

Natural selection and learning processes that endow agents with reliable abilities to complete tasks typically also endow agents with some sort of sensitivity to the conditions under which their task-completion abilities have a lower ratio of success. Doing that competently is itself an ability. We can call abilities of that type *precautionary abilities* —alternatively, *protective abilities*—, in that they are abilities to prevent the use of corresponding task-completion abilities when the circumstances are not in the range of action of the latter —i.e., when the circumstances are inappropriate for completing tasks.

By way of illustration, professional golfers are sensitive to wind conditions, so that a golfer's precautionary ability minimally prevents her from exercising her ability to put the ball into the hole during hurricanes, tornadoes, blizzards

³² There might even be *very narrow* abilities, i.e., dispositions to complete token tasks only under *local* conditions.

³³ Proper functionalists —see for instance Graham (2012)— and philosophers of biology —see for instance Millikan (1984)— have largely discussed these issues.

and fog conditions. The latter conditions are obviously not the conditions under which an ability to put the ball into a hole is considered reliable in golf. In the same way, we also tend to mistrust our vision when the light conditions are not good —e.g., when it is too dark. In that case, one exercises a precautionary cognitive ability to prevent the use of one's vision. These cases show that the normal range of action of a precautionary ability is different from the normal range of action of its corresponding task-completion ability, which means in turn that the two types of abilities may operate with different degrees of reliability. The following schematic definition accounts for that:

PRECAUTIONARY ABILITY: *S* has a precautionary ability [precautionary relative to *S*'s —cognitive or non-cognitive— ability to complete a certain token task —or tasks of type *T*— that is reliable to degree *D* in a certain set of conditions *C*] iff *S* has a disposition *not* to exercise her ability to complete that task —or tasks of type *T*— in a certain set of conditions *C** —and only in *C**, where *C** does not belong to *C*—, with degree of reliability *D**, where *D** is above a certain threshold.

It is important to highlight that precautionary abilities are not reliable in completing tasks —e.g., forming true beliefs—, at least not directly. Precautionary abilities are reliable in preventing the activation of a corresponding task-completion ability —e.g., a belief-forming ability— when the circumstances are *inappropriate* for that ability and only in such circumstances. This does not mean that a reliable precautionary ability must prevent belief formation in *all* circumstances that are inappropriate for deploying task-completion ability. Precautionary abilities need not be infallible in their operation, but reliable. To put it differently, they need to be reliable *in their own normal range of action*, which includes some proper subset of all the circumstances that are inappropriate for belief formation.

What if the circumstances are appropriate for a task-completion ability —e.g., if the circumstances are appropriate for belief formation? Then, the corresponding precautionary disposition, if reliable, must not prevent the activation of the task-completion ability —e.g., belief formation must not be prevented if the circumstances are appropriate. Recall that precautionary abilities are reliable in preventing the activation of task-completion abilities in a proper subset of all the circumstances that are inappropriate for the latter and *only in that subset*. Therefore, if one's dispositions prevent the activation of one's task-completion abilities

in circumstances that are appropriate for completing tasks, then those dispositions are not reliable in the stipulated sense and therefore cannot be considered precautionary *abilities*.³⁴

This means that precautionary abilities typically remain inactive or dormant in circumstances that are appropriate for completing tasks. This marks an important difference with the role that Sosa assigns to reflective competences in his account of knowledge: while reliable precautionary cognitive abilities typically remain inactive when the circumstances are appropriate for belief-formation, Sosa's reliable reflective competences are supposed to form second-order —implicit or explicit— beliefs when the circumstances are appropriate for belief formation —where the content of those second-order beliefs is that the corresponding first-order beliefs are, not only apt, but reliably so.

On the other hand, there are different factors that might make restrain oneself from using one's task-completion abilities, which amounts to saying that there are *different types of precautionary abilities* —this also marks an important difference with Sosa's reflective competences. For example, precautionary abilities might be *constitutively built in task-completion abilities*. This is often the case of abilities that have been naturally selected. In the cognitive case, we are endowed with all sorts of metacognitive mechanisms that prevent the activation of our belief-forming processes when the conditions are inappropriate —e.g., when fluency is interrupted we tend to reflect on the cognitive task we are engaged in while suspending judgment.³⁵

But precautionary abilities might be *contingently related to task-completion abilities*. For example, suppose that one is warned that it is very probable that the object that one is going to see is a fake barn. By competently putting that information into action, one can prevent the formation of the belief that the object in front of one is a barn. In the same way, an archer might restrain herself from shooting if she is told that the target she is aiming at is protected with a force field. Relevant information and the disposition to put it into action thus play the role of a —narrow— precautionary ability.

³⁴ In this way, an agent who does not form any belief in any condition —out of pathological fear or because of being maximally cautious— fails to exhibit reliable precautionary cognitive ability, because she prevents belief formation when the circumstances are appropriate for belief formation. Thanks to an anonymous referee for giving the example.

³⁵ Fluency refers to the easiness and quickness with which a cognitive task —e.g., a perceptual task— is completed. See Kahneman (2011) and Proust (2013) for excellent introductions to the psychological and philosophical issues surrounding metacognition.

Finally, a *task-completion ability* might play a *precautionary role* for other task-completion ability. For example, a sniper's ability to use a wind sensor might prevent failure when it comes to shooting.

In view of the variety of things that can *function* as a precautionary ability, the distinction between precautionary and task-completion abilities is not a distinction in kind —e.g., reflective versus non-reflective; second-order versus first-order. Instead, the best way to draw the distinction is *in functional terms*: all abilities are task-completion abilities because they are directed towards successfully completing tasks, but some of them are also considered precautionary because the task they serve to complete consists in preventing the use of other task-completion abilities in inappropriate circumstances —by succeeding in doing that they prevent the failure of the latter.

5.4 Sensitivity to Conditions and Resources

Consider the case in which the light conditions are bad —e.g., when it is too dark— and one's precautionary visual ability prevents one from forming a visual belief. Or consider the case of the archer who does not shoot because she realizes that the wind is too strong. In such cases, we can talk of the relevant precautionary abilities being *sensitive to local conditions*, in the sense that they correctly prevent the agents in question from using their task-completion abilities in their local conditions.

Now consider the case in which the actual light conditions seem to be good but they are such that they could easily be misleading. For example, suppose that

[y]ou see a surface that looks red in ostensibly normal conditions. But it is a kaleidoscope surface controlled by a jokester who also controls the ambient light, and might as easily have presented you with a red-light+white-surface combination as with the actual white-light+red-surface combination (Sosa 2007: 31).

Or consider the case of the archer who hits the only target that, unbeknownst to her, is not protected with an invisible force field that would repel any incoming arrow. In such cases, we can talk of the relevant precautionary abilities failing to be *sensitive to regional conditions*, in the sense that they fail to prevent the agents in question from using their task-completion abilities in their regional conditions

when the correct thing to do is not to use them —because such circumstances are inappropriate for the task-completion abilities.

Finally, when it comes to completing a certain task in certain local conditions we can talk of the agent's *resources* for completing that task in the following way:

RESOURCES: *S*'s resources when it comes to completing a certain task in certain local conditions include (1) *S*'s narrow and broad —cognitive or non-cognitive— abilities to complete that task —or tasks of that type— and (2) *S*'s precautionary abilities —precautionary relative to *S*'s task-completion abilities.

When the relevant abilities and tasks are cognitive we can talk of the agent's *epistemic resources*.

6 Robust Enough Virtue Epistemology

6.1 Partial and Complete Aptness

The second horn of the attributability dilemma makes salient the intuition that cases of environmental luck —e.g., FAKE BARNS, the kaleidoscope case, the case of the targets with force fields— involve apt successes. That is the reason why simple AAA accounts are insufficient for knowledge: they define knowledge as apt belief. However, the preceding account of the notion of ability and in turn of the notion of cognitive ability reveals that aptness understood in terms of MA is just part of the story, because it defines aptness only in terms of part of the agent's resources, but not in terms of all of them. We can rephrase MA as follows:

PARTIAL APTNESS: *S*'s successful performance is partially apt iff *S*'s success manifests *S*'s task-completion abilities.

A simple AAA account amounts to the view that knowledge is *partially apt* belief. With that distinction in place, we can now explain the intuition that environmentally lucky successes are apt: an environmentally lucky success is apt, but only partially, in that it manifests only part of the agent's resources, namely her task-completion abilities. But we have seen that task-completion abilities are not the only abilities we have. Aptness should be better conceived as manifestation of the whole resources when completing a task. If one lacks the relevant precautionary abilities or if one's relevant precautionary abilities are not sensitive to

regional conditions, at best one can achieve partial aptness. What one will not achieve is *complete aptness*:

COMPLETE APTNESS: *S*'s successful performance is completely apt
iff *S*'s success manifests *S*'s resources.

Robust enough virtue epistemology is the view that knowledge is completely apt belief. The following schema represents where exactly robust enough virtue epistemology stands in comparison to the other versions of robust virtue epistemology. They respectively define knowledge as:

(AA) Accurate and adroit belief.

(AAA) Partially apt belief.

(AAA^c) Completely apt belief.

(AA⁺) Ample belief.

(FW) Full-well belief.

As I will argue next, understanding knowledge in terms of AAA^c is incompatible with cases of environmental luck being cases of knowledge because environmentally lucky true beliefs, while partially apt, fall short of complete aptness. In other words, robust enough virtue epistemology not only escapes the second horn of the attributability dilemma, but it also accounts for the intuitions it elicits, namely that aptness, understood as success because of —task-completion—ability, is compatible with luck. Let us analyze the cases.

6.2 Eliminating Knowledge-Undermining Error Possibilities

In FAKE BARNS, the cognitive task at issue is to form beliefs about the presence of barns. Barnaby's actual cognitive success manifests his ability to form such beliefs —i.e., his visual task-completion ability— when looking at the only real barn in the county, so Barnaby's belief is partially apt. However, he fails to manifest his visual precautionary abilities because he is not sensitive to his regional conditions. Why is he not sensitive to his regional conditions? Because the regional conditions —the conditions in close possible worlds— are inappropriate for the manifestation of Barnaby's visual *precautionary* abilities: unlike the case in which it is too dark and one aptly stops belief formation, fakes of objects in

general are not part of the set of conditions relative to which a visual precautionary ability is reliable in its task of stopping belief formation. But, as we have seen, a cognitive success manifests cognitive ability —precautionary or non— only under appropriate circumstances. So Barnaby’s belief is not completely apt and therefore does not amount to knowledge.

The explanation of the kaleidoscope case is the same: while the local conditions are appropriate for manifesting the ability to form visual beliefs —a task-completion cognitive ability—, the regional conditions, i.e., the red-light+white-surface combination, are not appropriate for manifesting visual *precautionary* ability. That combination of light and color is simply not in the range of action relative to which the relevant visual precautionary ability is typically reliable. Sensitivity to regional conditions fails for that reason and in this way the target actual true belief that the surface is red does not amount to knowledge.

Analogously, in the case of the archer who luckily hits the only target that is not protected with an invisible force field, the archer’s shot is partially apt because it manifests the archer’s ability to shoot arrows, but the archer fails to be sensitive to regional conditions because those conditions are not in the range of action of her shooting *precautionary* abilities, and in this sense they are inappropriate for the manifestation of such abilities: unlike the case in which the wind is too strong and the archer aptly stops shooting, force field protected targets in general are not part of an archer’s range of action, i.e., they are not part of the set of conditions relative to which a shooting precautionary ability is reliable in its task of stopping shooting. So the archer’s shot is not completely apt.

It is worth noting that agents without reflective capacities can form completely apt beliefs. As explained above, possessing a precautionary ability relative to a certain task might be a matter of having a very narrow ability acquired in virtue of relevant information about the local or regional conditions. The simple automatic and unreflective disposition to put that information into action can make an agent stop belief formation when the conditions are not appropriate. In such cases, the reliability of that contingent ability might be sufficiently high, but only locally, and that is all that is required. It is not a requirement of the theory that agents must possess broad higher-order abilities guiding their first-order cognitive abilities.

To illustrate the point further, consider the following example from Kelp (2013). Consider two unsophisticated agents who are cognitive duplicates. One,

A, is in a normal barn county, the other, *B*, is in the fake barn county. Intuitively, *A* knows but *B* does not. Does it follow that *A* has some precautionary ability that *B* lacks?³⁶ No, *A* and *B* are equals concerning their precautionary abilities: neither *A*'s nor *B*'s capacities are reliable in preventing belief formation when situated in the regional conditions of the fake barn county. Why does *A* but not *B* know then? Because whatever precautionary abilities *A* and *B* have, they are inactive or dormant when the circumstances —e.g., the regional conditions— are appropriate for belief-formation, and only *A*'s regional conditions are appropriate for belief formation.

6.3 Dissolving the Safety and the Ability Dilemmas

The question that naturally arises is the following: how sensitive should one be to regional conditions in order for one's successful performance to manifest one's cognitive precautionary abilities and thus be completely apt? To answer this question I will explain how the account can accommodate two kinds of cases. This will not only yield important insights into how to think about complete aptness, but it will also serve to show why the safety and the ability dilemmas pose no problem to robust enough virtue epistemology —the two kinds of cases discussed are the kind of cases that serve to formulate the dilemmas.

In the first kind of case, a skillful archer is about to shoot at a target under what be normal circumstances —e.g., the archer is in good shape, the target is at a reasonable distance, the light conditions are good, and so on. The peculiarity of the case is that the archer is in a submarine whose hull could be fatally damaged at any moment by incoming missiles, of which the archer is unaware. The resulting explosion would let water in and the archer would not be able to exercise her shooting ability. She would die. But the archer feels very confident, shoots quicker than usual and hits the target before the explosion.³⁷ Does the fact that the archer is not aware of the incoming missiles entail that she fails to be sensitive to regional conditions and that her shot is therefore only partially apt?

An objector could argue that. However, sensitivity to regional conditions is a more nuanced notion than the notion that the possible objector has in mind. The objector could argue that the submarine case is the same kind of case as the archery case in which most targets are protected with invisible force fields,

³⁶ Thanks to an anonymous referee for suggesting the example.

³⁷ See Pritchard (2012: 268-9) for a similar case with a pianist.

because in both cases the shots in regional conditions—for the present purposes, in close possible worlds— would be unsuccessful.

However, we must distinguish two ways in which an ability might fail to produce an output. One way is that the ability is removed from the agent before producing the output. This is what happens in the submarine case: in close possible worlds water enters the chamber, kills the archer and prevents the shot. In the force field case, however, close possible worlds are such that the archer does not lose her ability, she exercises it, but it is masked by the protective force fields.³⁸ The cases are not analogous in that respect.

The difference is important because it marks a distinction between conditions that enable the exercise of an ability—*enabling conditions*— and conditions that determine the success or failure of that exercise—*determining conditions*. In the archery case, the effect of the force fields on the incoming arrows is clearly a condition that determines the failure of the shots. By contrast, in the submarine case, while the size, shape and distance at which the target is, the light conditions, and so on, are conditions that determine the success of the archer's performance, the fact that the missiles have not reached the submarine yet is a mere enabling condition for the exercise of shooting ability.

Only determining conditions that are inappropriate can prevent a success from manifesting precautionary ability. In other words, the kind of sensitivity to regional conditions that is required by complete aptness is sensitivity to determining conditions, not to enabling conditions. That is why in the submarine case the archer's shot is completely apt, but only partially apt when there are targets protected with force fields around.

CLOCK can be explained along the same lines. Cases of knowledge using instruments or epistemic devices—such as clocks or thermometers—are cases in which the reliability of the agent's belief-forming ability is *conditional* on the reliability of the instrument, so that disrupting the latter entails disrupting the former. In CLOCK, the reliability of Mr. Punctual's belief-forming ability is conditional on the reliability of the clock. So whatever disrupts the latter, it also disrupts the former. That indicates that the fact that the isotope has not yet decayed in local conditions is merely an enabling condition for belief formation. CLOCK is analogous to the submarine case in that regard: in the same way as the fact that the missiles have not reached the submarine yet is an enabling condition

³⁸ See Fara (2008: 846-7) for relevant discussion on masking.

for the exercise of shooting ability, the fact that the isotope has not decayed yet is an enabling condition for the exercise of Mr. Punctual's belief-forming ability. But since complete aptness does not require sensitivity to enabling conditions in regional conditions, Mr. Punctual's actual belief is completely apt.³⁹

The second kind of case that will help refine the view is DACHSHUND. DACHSHUND is a case of knowledge, so according to robust enough virtue epistemology it is a case of completely apt belief. However, Snoop would easily come to believe the same proposition he believes in his local conditions falsely—that the animal before his eyes is a dog. Does this mean that he lacks sensitivity to the relevant determining factors in regional conditions and hence that his actual belief is not completely apt?

To answer that question, we need to reflect a bit further on what is the relevant cognitive task at issue. The relevant cognitive task is not to recognize objects as animals. Snoop presumably has that ability in local and regional conditions—i.e., both when looking at the dachshund and the wolves. Instead, the relevant cognitive task consists in recognizing animals as dogs, but not any animal. As I stated before, two tasks are of the same type just in case the parameters or conditions that define them are roughly the same. Thus, in the same way as hitting an apple off someone's head is not the same task as hitting a mosquito—because the targets have different sizes—, visually recognizing a dachshund as a dog is not the same cognitive task as recognizing a wolf as —not being— a dog.

The cognitive tasks of recognitional abilities are similarly individuated by the parameters that define them and the properties instantiated by the categorized objects—e.g., their shape, size, weight, color, texture, smell, and so on. In general, a categorization task consists in applying concept *C* to objects instantiating properties $p_1, p_2 \dots p_n$. In this sense, one completes a perceptual categorization task when one perceives those properties in an object—or an enough number of them—, the object instantiates them, and one categorizes it as being *C*. For ex-

³⁹ One could modify the submarine case so that close possible worlds—i.e., regional conditions—are such that the archer is able to take the shot but a missile hits the submarine thereby preventing the shot from being successful, while the archer survives the whole episode retaining her shooting ability. This modification makes the case analogous to the force fields case: in both cases the circumstances are very demanding—because the relevant precautionary abilities of an archer do not typically include being sensitive to incoming missiles or invisible force fields. But all the cases show is that being sensitive to *determining* factors is sometimes more difficult than usual and, in this sense, achieving complete aptness might be a difficult task without external aid—e.g., without relevant information about the environment, i.e., information that would locally endow one with an adequate precautionary disposition for that kind of unusual circumstances.

ample, in DACHSHUND Snoop correctly classifies the dachshund as a dog on the basis of perceived properties such as *<short-legged>*, *<long-bodied>*, *<droopy ears>*, and so on. One performs the *same type* of categorization task—but fails to complete it—if one perceives those properties—or an enough number of them—in an object that does not instantiate them but that one categorizes as being *C* nevertheless. This is the case of fakes in general: one perceives in them the typical features of the object of which they are fakes. In this sense, recognizing a dachshund as a dog counts as the same categorization task as recognizing a fake dachshund as a dog.

However, one *does not* perform the same categorization task if the object to categorize instantiates—an enough number of—different properties $q_1, q_2 \dots q_n$, one perceives them, and one classifies the object as being *C*. *C* might be a concept that correctly applies to objects with very different properties—this is the case of the concept of dog, which correctly applies to both dachshunds and German shepherds. DACHSHUND is thus a case in which the same dog-beliefs are produced on the basis of the agent perceiving two different sets of properties that correctly fall under the concept of dog. Specifically, Snoop classifies wolves as dogs because they have the features of, say, German shepherds. But those features are different from the features of dachshunds. That is why they count as different categorization tasks.⁴⁰

In conclusion, by sensitivity to regional conditions is not only meant sensitivity to *determining factors* in regional conditions, but also sensitivity to factors that determine the success or failure of the *same type of task* one performs in local conditions. DACHSHUND fails to show that sensitivity of that kind fails when an agent knows. Complete aptness must be understood accordingly.

Concluding Remarks

Robust virtue epistemology is built on the widespread intuition that the concept of knowledge can be explained by how the exercise of cognitive ability contributes to succeeding cognitively. This intuitive idea has proven to be difficult to develop in a systematic way. The reason is that developments of the idea clash with two independent problems connected by a dilemma that intends to show that the exercise of cognitive ability is at the same time too strong and too

⁴⁰ See Goldman (1976) for a similar treatment of DACHSHUND.

weak to account for knowledge, but never robust enough to explain all the cases. In this paper, I have shown that the intuitive idea on which robust virtue epistemology is built can be developed in a systematic way without falling prey to that dilemma—and to related dilemmas that jeopardize other virtue epistemological accounts. The way I have approached the problem is by taking a step back in order to examine the very notions of ability and cognitive ability. The result has been a particular view—robust enough virtue epistemology—that has proven to be better than its rivals in solving a long-standing problem in epistemology: the problem of accounting for the nature of knowledge.

Acknowledgements

Thanks to Chris Kelp, two anonymous reviewers for *Synthese*, and the audiences of the LEG seminar (Leuven), the University of Copenhagen, and the Virtue Epistemology Conference (Leuven).

References

- Blouw, P., Buckwalter, W., & Turri, J.** (forthcoming). Gettier cases: a taxonomy. In R. Borges, C. de Almeida, & P. Klein (eds.), *Explaining knowledge: new essays on the Gettier problem*. Oxford University Press.
- Bogardus, T.** 2014. Knowledge under threat. *Philosophy and Phenomenological Research* 88: 289–313.
- Bogardus, T. & Marxen, C.** 2014. Yes, safety is in danger. *Philosophia* 42: 321–34.
- Broncano-Berrocal F.** 2014a. Anti-Luck (Too Weak) Virtue Epistemology. *Erkenntnis* 79: 733–754.
- 2014b. Is Safety In Danger? *Philosophia* 42: 63–81.
- Chrisman, M.** 2012. The normative evaluation of belief and the aspectual classification of belief and knowledge attributions. *Journal of Philosophy* 109: 588–612.
- Fara, M.** 2008. Masked abilities and compatibilism. *Mind* 117: 843–65.
- Goldman, A. I.** 1976. Discrimination and perceptual knowledge. *Journal of Philosophy* 73: 771–91.
- 1979. What is justified belief? In G. Pappas (ed.), *Justification and knowledge*. Reidel: 1–23.
- Graham, P. J.** 2012. Epistemic entitlement. *Noûs* 46: 449–82.

- Greco, J.** 2000. *Putting skeptics in their place: The nature of skeptical arguments and their role in philosophical inquiry*. Cambridge University Press.
- 2010. *Achieving knowledge: A virtue-theoretic account of epistemic normativity*. Cambridge University Press.
- 2012. A (different) virtue epistemology. *Philosophy and Phenomenological Research* 85: 1-26.
- Hawley, K.** 2003. Success and Knowledge-How. *American Philosophical Quarterly* 40: 19-31.
- Haddock, A. Millar, A. & Pritchard, D.** 2010. *The nature and value of Knowledge: Three investigations*. Oxford University Press.
- Heil, J.** 2005. Dispositions. *Synthese* 144: 343-56.
- Kahneman, D.** 2011. *Thinking, fast and slow*. Farrar, Straus and Giroux.
- Kallestrup, J. & Pritchard, D.** 2014. Virtue epistemology and epistemic twin earth. *European Journal of Philosophy* 22: 335-57.
- Kelp, C.** 2013. Knowledge: The safe-apt view. *Australasian Journal of Philosophy* 91: 265-78.
- Lackey, J.** 2007. Why we don't deserve credit for everything we know. *Synthese* 158: 345-61.
- 2009. Knowledge and credit. *Philosophical Studies* 142: 27-42.
- Maier, J.** 2013. The agentive modalities. *Philosophy and Phenomenological Research* 87: 113-34.
- McKinnon, R.** 2014. You make your own luck. *Metaphilosophy* 45: 558-77.
- Millikan, R.** 1984. *Language, thought, and other biological categories*. MIT Press.
- Palermos, S. O.** 2013. Could reliability naturally imply safety? *European Journal of Philosophy* doi: 10.1111/ejop.12046.
- Pritchard, D.** 2006. A defence of quasi-reductionism in the epistemology of testimony. *Philosophica* 78: 13-28.
- 2010. Cognitive ability and the extended cognition thesis. *Synthese* 175: 133-51.
- 2012. Anti-luck virtue epistemology. *Journal of Philosophy* 109: 247-79.
- Proust, J.** 2013. *The philosophy of metacognition: Mental agency and self-awareness*. Oxford University Press.
- Sosa, E.** 1991. *Knowledge in perspective*. Cambridge University Press.
- 2007. *A virtue epistemology: Apt belief and reflective knowledge (Vol. I)*. Oxford University Press.

- 2009. *Reflective knowledge: Apt belief and reflective knowledge (Vol. II)*. Oxford University Press.
- 2011. *Knowing full well*. Princeton University Press.
- 2015. *Judgment and agency*. Oxford University Press.
- Turri, J.** 2011. Manifest failure: The Gettier problem solved. *Philosophers' Imprint* 11: 1-11.
- (*forthcoming*). Knowledge as achievement, more or less. In M. A. Fernández (ed.), *Performance epistemology*. Oxford University Press.
- Vihvelin, K.** 2004. Free will demystified: A dispositional account. *Philosophical Topics* 32: 427–50.
- Zagzebski, L.** 1994. The inescapability of Gettier problems. *Philosophical Quarterly* 44: 65-73.