



**University of
Zurich**^{UZH}

**Zurich Open Repository and
Archive**

University of Zurich
University Library
Strickhofstrasse 39
CH-8057 Zurich
www.zora.uzh.ch

Year: 2018

Virtue-Reliabilism and the Value of Knowledge: Classical and New Problems

Meylan, Anne

DOI: <https://doi.org/10.4324/9781315712550>

Posted at the Zurich Open Repository and Archive, University of Zurich

ZORA URL: <https://doi.org/10.5167/uzh-167733>

Book Section

Accepted Version

Originally published at:

Meylan, Anne (2018). Virtue-Reliabilism and the Value of Knowledge: Classical and New Problems. In: Battaly, Heather. *The Routledge Handbook of Virtue Epistemology*. New York: Taylor Francis, 317-329.

DOI: <https://doi.org/10.4324/9781315712550>

Virtue-Reliabilism and the Value of Knowledge: Classical and New Problems

Anne Meylan, University of Basel

Word count: 5820 (including references)

Abstract

This first part of this chapter presents the virtue-reliabilist's answer to the classical value problems of knowledge. According to this solution, knowledge is a better cognitive state than what falls short of it—mere true belief or true+Gettierized belief—because when a subject knows, she deserves credit for her true belief. The second part of this chapter is devoted to showing that this solution cannot be extended to solve the “new” value problem, that is to say, the problem of explaining why some higher form of knowledge—what Sosa calls *full knowledge*—is better than some lower form of knowledge—Sosa's animal knowledge. The basic problem for Sosa is that, when a subject fully knows, it is not necessarily the case that she deserves *more credit* than when she has animal knowledge.

Virtue-Reliabilism and the Classical Value Problems

Three value problems

The problem of the value of knowledge results from two solid but unfortunately opposing axiological intuitions. On one hand, as observed by Plato in the *Meno* (97a-d), as far as the achievement of our practical goals is concerned, knowledge does not seem more valuable than mere true belief. Animals do not need to know that a predator is coming in order to increase their chances of survival. A true belief is sufficient to make them run away. Or, to re-use Plato's famous example, having a true belief about the way to Larissa is sufficient to reach Larissa. Knowledge seems superfluous. On the other hand, we do value knowledge more than true belief. Suppose someone gives you the choice between:

- a. getting knowledge that this plant is toxic or;

b. getting a mere true belief that this plant is toxic

while making clear that there are no special practical costs attached to one or the other of these choices. The vast majority of us will prefer knowledge to mere true belief (Pritchard 2009a).¹

At the root of the problem of the value of knowledge lie these two opposite intuitions.

1. Knowledge does not seem more practically valuable than mere true belief;
2. But knowledge seems to be more valuable than mere true belief.

Solving the problem of the value of knowledge requires explaining why knowledge that *p* is a better cognitive state than merely believing truly that *p*, even though knowing that *p* and merely believing truly that *p* are on a par as far as their practical value is concerned. More precisely, this is required to solve the *primary* value problem of knowledge (Pritchard 2007, 2009a).

There is, indeed, a *secondary* value problem (and even a *tertiary* value problem). The secondary value problem arises from the observation that knowledge is not only better than mere true belief. Knowledge is more valuable than *any* cognitive state that falls short of it. Specifically, knowing is also better than holding a Gettierized+true belief, that is to say, than holding a justified true belief that falls short of knowledge because it is subject to Gettier problems (Pritchard 2007, 2009a). Here as well, we cannot rely on any additional *practical* value to explain this difference. Knowledge that *p* is not, presumably,² more practically valuable than the Gettierized+true belief that *p*. For instance, suppose that I bet that the person who will be hired has ten coins in his pocket (Gettier 1963). My holding the Gettierized+true belief that the person who will be hired has ten coins in his pocket allows me to win my bet as efficiently as my knowing that the person who will be hired has ten coins in his pocket would. Solving the *secondary* value problem requires explaining *why* it is better to know than to be in a state that falls short of knowledge— a Gettierized+true belief. And this must be done even though an instance of knowledge and a Gettierized+true belief have the same practical value.

Making matters even more complicated, knowledge does not simply possess a higher

degree of value than, mere true belief and Gettierized+true belief. Knowledge is also valuable in a different way. That is, knowledge and true belief do not simply occupy different positions on a single evaluative continuum, with knowledge at the top and true belief further down.³ The value of knowledge also belongs to a different *kind* of value. In Pritchard's words:

Accounting for the value of knowledge requires us to offer an explanation of why knowledge has not just a greater degree, but also a different kind of value than whatever falls short of knowledge. Call this the *tertiary value problem* (Pritchard, 2009a: 4).

Solving the *tertiary* value problem requires explaining why knowledge possesses a different *kind* of value. One elegant way of meeting the challenge raised by the tertiary value problem would be showing that knowledge is, in contrast to mere true belief, *finally* valuable.⁴ Final values are traditionally opposed to mere instrumental values. For an entity to be *finally* valuable is for this entity to be valuable for its own sake. For an entity to be *instrumentally* valuable is for this entity to be valuable for the sake of what it leads to (Rabinowicz and Rønnow-Rasmussen 1999). Morally good actions are generally taken to be finally valuable because they seem good for their own sake, they seem to draw their goodness at least partly from what they are and not solely from what they lead to. In contrast, drugs, warm baths, and punctual trains are merely instrumentally valuable. They draw their value solely from the goodness of their effects. It is also possible for certain entities to be both finally and instrumentally valuable. Beautiful paintings, for instance, might be said to be both finally valuable in virtue of their aesthetic qualities and instrumentally valuable in virtue of the aesthetic pleasure they induce. Solving the tertiary value problem by claiming that knowledge is, in contrast to what falls short of it, *finally* valuable requires showing that knowledge is good for its own sake and not merely good in virtue of its good outcomes. As we will see below, this is a requirement that virtue-reliabilism fulfills especially well.

Alongside the aforementioned value problems lies the problem of the analysis of knowledge. Which property (or properties) turns a mere true belief into an instance of knowledge? The value problems and the problem of the analysis of knowledge are intertwined because it is a condition on a successful analysis of knowledge that it solves the value problem(s).⁵ Suppose that I defend an analysis of knowledge in which the exemplification of some property *I* is sufficient to turn a true belief into an instance of knowledge. This means that my solution to the value problems will have to be that the exemplification of *I* by a true belief makes this true belief more valuable than anything that falls short of knowledge. Given that knowledge is a cognitive state that is essentially characterized by its superior value, if the exemplification of *I* by a true belief turns out not to make the true belief more valuable than anything that falls short of knowledge, then exemplifying *I* cannot be sufficient to turn a true belief into an instance of knowledge. Indeed, if exemplifying *I* does not make the true belief more valuable than anything that falls short of knowledge, I am forced to say that the additional value of knowledge results from some *other property*. And this, of course, is to admit that exemplifying *I* is not sufficient to turn a true belief into an instance of knowledge.

This difficulty threatens classical process-reliabilism. According to the classical process-reliabilist analysis of knowledge (Goldman 1979), knowledge is true belief that causally results from a reliable process (a process that generally produces true beliefs). Accordingly, the process-reliabilist solution to the primary value problem must be an explanation of why the exemplification, by a true belief, of the property of resulting from a reliable process makes this true belief even more valuable. Arguably, this explanation fails. This is because the instrumental value that the property of resulting from a reliable process should add is *swamped* by the fact that the belief that results from this reliable process is already true (Kvanvig 2003; Zagzebski 2003).⁶ To explicate, consider two cups of coffee that are exactly on a par as far as their tastefulness is concerned but differ as regards the reliability of

the processes that produced them. A reliable coffee machine —one that generally produces tasty coffee—produced the first cup. The machine that produced the second cup is, in contrast, unreliable. However, on this occasion, it produced a tasty coffee, a coffee that is as tasty as the first cup, the one produced by the reliable coffee machine. The fact that a reliable coffee machine produced the first cup of coffee does not make the first cup better than the second one. Any alleged instrumental value that one might be tempted to assign to the first cup of coffee because it results from a reliable coffee machine is *swamped* by the fact that this first cup is already tasty. The same reasoning holds with respect to true beliefs. A true belief that is the outcome of a reliable process —viz. an instance of knowledge according to process-reliabilism— does not draw any additional value from its being the result of such a process. Any alleged ‘additional’ value is swamped by the fact that the belief is already true.

The virtue-reliabilist’s solution to the value problems

One important advantage of the virtue-reliabilist’s⁷ analysis of knowledge — defended by John Greco (2003, 2010, 2012) and Ernest Sosa (2007, 2009a, 2011, 2015)— over process-reliabilism is its ability to provide an elegant answer to the three value problems above. Before presenting the virtue-reliabilist’s solution to the value problems, let me briefly recall what knowledge is according to virtue-reliabilism.

In cases of knowledge, S believes the truth because S’s belief is produced by ability (Greco 2012: 1).

A subject *S* knows whether *p* only if *S* believes the truth regarding *p* because *S* believes from intellectual virtue (or excellence, or ability) (Greco and Reibsamens forthcoming: 1).

According to Sosa (2007, 2009a, 2015), animal knowledge is *apt* belief, that is, belief whose truth manifests a skill or competence seated in the subject. When a subject’s true belief manifests a competence, I will say that this true belief *results* from her exercising

her competence. The virtue-reliabilist's idea is that this specific relationship is *not* instantiated in the classical Gettier cases (Gettier 1963). In the classical Gettier cases, the true belief does *not* result from the exercise of the subject's competence. The subject arrives at a true belief due to luck. This explains why the subject does not know.

Let us now consider virtue-reliabilism from the perspective of the value problems. Given what they take knowledge to be, the task for virtue-reliabilists is to explain why a true belief *that results from the exercise of a competence* is better than one that does not instantiate this property. The virtue-reliabilist's explanation usually starts with the comparative evaluation of two performances. Here is an example of such a comparative evaluation.

The lucky physician

A physician examines a patient who suffers from severe stomach pains. Even though the physician is not able to give a diagnosis on the basis of this examination, she prescribes her patient a medicine that she picks at random from her compendium. Now, luckily, this is exactly the drug that her patient needs and the patient is healed.

The competent physician

A physician examines a patient who suffers from severe stomach pains. After having examined her, she arrives at the diagnosis that she suffers from diverticulitis.

She prescribes her patient the appropriate medicine and the patient is healed.

Compare the performances of the lucky and the competent physician. Their medical performances—the lucky and the competent prescription—are both *successful*: they both attain their constitutive purposes.⁸ Both cause the patient's recovery. Their being successful makes them both valuable performances. But, virtue-reliabilists insist, we value the competent performance even *more* than the lucky performance. Causing the recovery of a patient is a better performance when this recovery is not a lucky outcome but results from the exercise of some

medical competence. The competent performance possesses an additional value that the lucky performance does not possess. Moreover, what is true of physicians' successful performances is true of successful performances in general. And, so:

A subject S's successful performance is even more valuable when it is not due to luck but results from the exercise of S's competence.

The idea is sometimes expressed in terms of what is *creditable* to the subject (assuming that a successful performance that results from the exercise of a subject S's competence is creditable to S). A subject S's successful performance is even more valuable when it is not due to luck but is creditable to S.⁹

When a successful performance—e.g. a medical cure, an accurate shot in archery, a winning chess move—is creditable to a subject, this performance displays two values. The first value is simply the value that a performance possesses when it achieves its purpose, that is, when it is successful. Its additional value is what I have called its *credit value* (Meylan 2013). It is the value that a successful performance possesses when it is creditable to a subject and is not a lucky outcome.

I have said that the upholders of virtue-reliabilism owe us an explanation. Why is a true belief *that results from the exercise of a competence* better than a true belief that does not instantiate this property? Such an explanation follows directly.

Regarding the value of knowledge, we may note that, in general, we value success from ability over mere lucky success. Our preference for knowledge over mere true belief may now be understood as an instance of this more general valuing. Again: in cases of knowledge, true belief is no mere lucky success; rather, S's believing the truth is attributable to the exercise of ability (Greco 2012: 2).

The virtue-reliabilist's answer to the value problems is thus as follows. True beliefs that result from the exercise of a subject's competence—that is, instances of knowledge—

are more valuable than true beliefs that do not result from such an exercise —i.e., mere true beliefs or true+Gettierized beliefs— because the former are successful performances that are creditable to the subject while the latter are lucky outcomes. Or, to say it more briefly, true beliefs that result from the exercise of a subject's competence possess an additional credit value that is not possessed by true or Gettierized+true beliefs that are luckily true.

Besides offering a solution to the primary and the secondary value problems, virtue-reliabilism also solves the tertiary value problem. Recall that answering the tertiary value problem requires explaining why knowledge possesses a different *kind* of value than that possessed by anything that falls short of knowledge. Answering the tertiary value problem requires explaining why the value of knowledge differs from the value of mere true or Gettierized+true belief, not only as matter of degree but also as a matter of kind. This is a challenge that the aforementioned virtue-reliabilist answer meets very well. Indeed, credit value—the additional value that, according to virtue-reliabilism, only knowledge possesses — is *not* a value that knowledge possesses in virtue of its good effects. As virtue-reliabilists emphasize, credit value is not an instrumental but a *final* value of knowledge (recall the previously presented distinction). In contrast, the value possessed by mere true and Gettierized+true belief is merely *instrumental*. Mere true and Gettierized+true beliefs draw their value solely from their good practical effects, e.g. they help us not to eat toxic plants, they allow us to reach Larissa, etc. Mere true and Gettierized+true beliefs are merely instrumentally practically valuable.

To recap, according to virtue-reliabilism, the value of knowledge differs from the value of mere true and Gettierized+true beliefs as a matter of kind because the additional value of knowledge is a *final* value while true and Gettierized+true beliefs are merely *instrumentally* valuable.¹⁰

The Jokester case

Virtue-reliabilism seems to fare well with the value problems.¹¹ Famously though, the most serious threat for virtue-reliabilism comes from its apparent inability to deal with cases involving environmental luck, cases like the *Jokester*:

You 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).¹²

Your belief that the surface is red results from the exercise of your visual competence. But, even if the Jokester provides you with the good white-light+red surface combination, you arguably do *not* know that the surface is red. In such a case, you might too easily have been wrong, and when you might too easily have been wrong you don't have knowledge. The risks that the Jokester has presented you with the bad combination are too high. So the problem for virtue-reliabilism is that, in cases involving environmental luck, the virtue-reliabilist's conditions are satisfied while the subject does not know (Kvanvig 2003; Pritchard 2009a; Pritchard 2009b; Sosa 2007; Sosa 2015). The virtue epistemological literature contains various replies to this objection.¹³ Sosa's answer is especially interesting; indeed, Sosa's reply to the Jokester case leads to the emergence of a *new value problem*. Let us see why.

Sosa's Virtue-Reliabilism and the New Value Problem

Sosa's solution to the Jokester objection relies on a distinction between kinds of knowledge. Famously, Sosa supports this distinction between kinds of knowledge by differentiating between various standards along which performances can be evaluated (Sosa 2007; Sosa 2009a; Sosa 2009b; Sosa 2011; Sosa 2015). I briefly introduce these standards.

Success, aptness and meta-aptness

First, for Sosa, a performance is valuable when it is *successful*, when it achieves its constitutive

purpose. The act of prescribing medication is better when it succeeds in curing the patient than when it fails to do so. Second, as seen above, a successful performance is even more valuable when it results from the exercise of a subject's competence than when it is due to luck. A successful hunter's shot is even better when it results from the exercise of the hunter's competence and is not a lucky outcome. Sosa calls a successful performance that results from the exercise of a subject's competence an *apt* performance. Thus Sosa's *standard of aptness* can be formulated in the following terms:

A subject S's successful performance P is apt when P results from the exercise by S of one of her competences.

Third, a successful performance is also sometimes, according to Sosa, *meta-apt*.

A shot is meta-apt iff it is well-selected: i.e., iff it takes appropriate risk, and its doing so manifests the agent's competence for target and shot selection (Sosa 2009b: 12).

For instance, a hunter's shot is meta-apt when it is performed while the hunter truly believes that the risks of missing her target are not too high (Sosa 2015: 68). Paraphrasing, Sosa's *standard of meta-aptness* entails (at least) the following:

A subject's performance P is *meta-apt* when P is accompanied by the second-order true¹⁴ belief that the risks of P being unsuccessful are not too high.¹⁵

Finally, the meta-aptness of a performance should be distinguished from its *environmental safety* where:

A subject's successful performance P is environmentally safe when P takes place in an environment in which the risks of P being unsuccessful are not too high.

Consider the case of the skillful hunter once again. Suppose that she goes hunting on a sunny, non-windy day and that she woke up in normal good form. The risks of missing her target today are not especially high. She shoots and hits her target as a result of her hunting

competence. Her successful performance is apt and environmentally safe. What about the meta-aptness of the performance? As we just saw, the latter depends on whether she, furthermore, believes that the risks of missing her target are not too high. In order for the successful performance of the hunter to be meta-apt, the environment must not only be such that her performance could not easily have been unsuccessful, but the hunter also has to believe that the environment is such. Put briefly, Sosa's standard of meta-aptness combines:

1. A *doxastic condition* that requires that the subject believes, on a second-order level, that the environment is safe, that is, that the risks of failing are not too high; with
2. A *modal condition* that requires that the environment really be safe.

That meta-aptness requires the satisfaction of the second, modal, condition is made clear by Sosa's repeated insistence that the second-order belief has to be apt and thus true (Sosa 2007: 24; Sosa 2015: 69). As for the first, doxastic, condition, the fact that it has to be fulfilled is made clear by Sosa's basketball player case (Sosa 2015: 71-72). Sosa asks us to consider:

a shooter as she approaches a distance to the basket near her relevant threshold of reliability....Suppose her to be above the threshold, but indiscernibly so to her. A statistician-coach-observer might know perfectly well that the player is now barely above the threshold. Suppose he has studied her success rate extensively, aided by a device that measures with exactitude her distance from the basket....He can tell that she is reliable enough at that distance.... But she herself is very far from knowing any such thing (Sosa 2015: 71-72).

Sosa continues:

The player may still attain her basic aim: namely, to sink that shot in the basket. In that respect her shot may... be apt... What then is she missing? Anything? ...

Unlike the statistician, she is unable to tell that her shot is still reliable enough at

that distance. If she shoots anyhow, and her shot turns out to be reliable enough, she may aptly reach her aim of scoring that goal. What she does not aptly reach, however, is the aim of aptly scoring (Sosa 2015: 71-72).

The player's shot is *safe* since, as the statistician observes, she shoots from a distance that is above her threshold of reliability. But, as Sosa puts it, something is *missing*. Her shot is not meta-apt because the basketball player is not aware that she is reliable at that distance. This is an example in which the performance is apt and safe but not meta-apt. Note that there are also, according to Sosa, meta-apt performances that are not apt:

An archer/hunter's shot selection and risk taking may be excellent, for example, and in taking a certain shot he may manifest his competence at assessing risk, while the shot itself nevertheless fails, being unsuccessful (inaccurate) and hence inapt. The shot is hence meta-apt without being apt (Sosa 2009b: 12).

So, a shot can be apt without being meta-apt, or meta-apt without being apt. Can we say something similar of the relationship between meta-aptness and safety? As the above basketball case illustrates, there are safe performances that are not meta-apt. But, given that meta-aptness entails safety, there are no meta-apt performances that are not safe.

Animal, reflective and full knowledge

The application of the standards of aptness and meta-aptness to *true beliefs* leads to the recognition of three kinds of knowledge.

The *first*, low-level, kind of knowledge is the kind that a subject possesses when she truly believes something as a result of the exercise of her competence, that is, when she holds an *apt belief*. This first kind of knowledge is so-called *animal knowledge* (Sosa 2007; Sosa 2009a; Sosa 2015).

The *second*, higher-level, kind of knowledge is the kind that a subject possesses when she holds a second-order apt belief that the risk that her first-order belief is false is not too high.

Sosa calls this second kind of knowledge *reflective knowledge* (Sosa 2009b; Sosa 2015).

The *third*, highest-level, kind of knowledge consists in having animal knowledge *together with* reflective knowledge. Sosa calls this third kind of knowledge *full knowledge*.¹⁶

As the basketball case illustrates, a subject's performance can be apt without being meta-apt. The same is true, according to Sosa, in the realm of knowledge. A subject can possess animal knowledge (apt belief) without possessing either reflective or full knowledge (meta-apt belief). A subject can attain a first-order true belief as a result of the exercise of one of her competences without possessing the second-order *true*—as we shall see the italicization is important—belief that the risks of being wrong on the first-order level are not high.

Crucially, for Sosa, this is what happens in the Jokester case. In such a case, the subject has animal knowledge. What the subject lacks in the Jokester case—since there is undeniably something lacking that needs to be captured—is reflective and thus full knowledge. The reason why the subject, in the Jokester case, lacks reflective and thus full knowledge is not that she does not entertain any second-order belief regarding the risks of being wrong on the first-order level. This would, for sure, be one way of lacking reflective/full knowledge. But nothing in the description of the Jokester case is said regarding the presence or the absence of such a second-order belief. Nothing allows us to say that the lack of reflective/full knowledge is due to the absence of such a second-order belief regarding the risks of being wrong. What is made very clear in the description of the Jokester case, and what explains why the subject lacks reflective/full knowledge, is that her first-order belief is not safe. It is the *modal condition* (and not the doxastic condition) of meta-aptness that is not fulfilled. To put it slightly differently, even if the subject in the Jokester case entertained the second-order belief that the risks of being wrong are not high, she would still lack reflective/full knowledge since this second-order belief would be false.

The New Value Problem

Recall the virtue-reliabilist's solution to the classical value problem:

True beliefs that result from the exercise of a subject's competence—that is, instances of knowledge—are more valuable than true beliefs that do not result from such an exercise—i.e., mere true beliefs or true+Gettierized beliefs—because the former are successful performances that are creditable to the subject while the latter are lucky outcomes.

This solution explains why Sosa's animal knowledge is better than a mere true belief or a Gettierized+true belief, since an instance of animal knowledge is a true belief that results from the exercise of a subject's competence. An instance of animal knowledge is more valuable than a merely true or a Gettierized+true belief because, in the case of animal knowledge, the true belief is creditable to the subject.

What about the value of *full knowledge*? Sosa explicitly says that the complex cognitive state consisting in having animal knowledge *and* reflective knowledge is more valuable than the simpler state consisting in having only animal knowledge.

Apt belief aptly noted, reflective knowledge, is better than mere apt belief or animal knowledge, especially when the reflective knowledge helps to guide the first order belief so that it is apt (Sosa 2009b: 15).

Or, more generally, performances that are apt *and* meta-apt are, according to Sosa, better than mere apt performances (Sosa 2009b: 13; Sosa 2015: 68-74). But why would an apt belief that is also meta-apt be better than a mere apt belief? Why would full knowledge be more valuable than mere animal knowledge? This is what I suggest we call *the new value problem*. The following quotation gives us a clue as to how Sosa answers these questions.¹⁷

Diana's apt shot that kills its prey is a better shot if apt than if successful only by luck and not through competence. Moreover, it is also a better, more admirable, more creditable shot, if its success flows also from her target-selecting, shot-picking

competences. Her shot is more creditable in that case than it is when the right competence is manifest in conditions required for a successful first-order performance, but only by luck external to any such selection meta-competence on her part (Sosa 2009b: 14).

In other words, Sosa extends the aforementioned virtue-reliabilist solution to full knowledge. The reason why *full knowledge* is more valuable than mere animal knowledge is a matter of credit as well. When a subject fully knows, the true belief is even more *creditable* to the subject than when she possesses mere animal knowledge. And this explains why full knowledge is a better epistemic state than animal knowledge. To put it differently, full knowledge displays *more* credit value than animal knowledge and this is why the former is, according to Sosa, better than the latter. In the following section, I critically consider this solution to the new value problem.

Is full knowledge more creditable?

Is a subject S's true belief really *more creditable* to S when the true belief in question is not only apt but also meta-apt, that is, when the subject has full knowledge? This, I believe, is doubtful. To show why this is so, I need, first, to introduce two new cases: Risky and Secure.

Risky

Rafa believes, *on a second-order level*, that the risks of forming a false first-order belief regarding the color of the surface in front of him are not too high. His second-order belief is *false*. In fact, the environment is such that the risks of forming a false first-order belief are indeed high. Guided by this second-order belief, Rafa exercises his reliable visual competence and luckily—given the environment he is in—forms a *true* first-order belief that the surface is red as a result of the exercise of his visual competence.

Rafa's belief that the surface is red is neither safe nor meta-apt, given that his second-order

belief is false. Rafa's belief that the surface is red is, however, apt. It is only a piece of *animal knowledge*.

Secure

Sara believes, *on a second-order level*, that the risks of forming a false first-order belief regarding the color of the surface in front of her are not too high. This second-order belief is *true*. The environment really is such that the risks of Sara forming a false belief are not too high. Guided by this second-order belief, Sara exercises her reliable visual competence and truly believes that the surface is red as a result of this exercise.

Unlike Rafa's, Sara's first-order true belief is not only apt but also *meta-apt*. *It is a piece of reflective and full knowledge*.

Recall that my purpose in this section is to critically consider Sosa's answer to the new value problem. Sosa's answer consists in extending, to full knowledge, the virtue-reliabilist's solution to the 'classical' value problem. Full knowledge is better than animal knowledge because, in the case of full knowledge, the true belief is even *more creditable* to the subject. The supposition that creditability comes in degrees is not problematic. A subject can contribute more or less to the occurrence of a successful performance and this makes her successful performance more or less creditable to her. More precisely, the attribution of creditability is, on my view, very probably governed by the following principle.

The degrees of creditability principle

A successful performance P1 is more creditable to a subject S than a successful performance P2 only if S has, in one way or another (for instance, by exercising one of her competences) contributed more to the occurrence of P1 than to the occurrence of P2.

Bearing this principle in mind, let us reconsider the two cases above. Does the subject deserve

more credit for her first-order true belief in “Secure” than in “Risky”, as Sosa’s answer to the new value problem forces him to say? It does not seem so. The contribution of a subject to the possession of a true belief is *not* higher when she —like Sara in “Secure”— has formed a *true* second-order belief about the risk of forming a false first-order belief than when she —like Rafa in “Risky”— has formed a *false* second-order belief about this risk. Indeed, we can assume that Sara in “Secure” does not *do* anything to increase her chances of forming a true belief that Rafa does not do in “Risky”. Sara and Rafa both take time to assess the risks of being wrong, take these risks to be acceptable, and exercise their visual competence. They are exactly on a par as regards their contribution to the adoption of a true belief, and thus—given the principle above— exactly on a par as regards their creditability.

Contrary to what Sosa’s answer to the new value problem claims, Sara does not deserve more credit than Rafa for her true belief. And, once again, this is because Sara does not contribute more in coming to possess a true belief than Rafa. They are on the par as far as their respective endeavours to obtain a true belief are concerned. Both Sara and Rafa try as much as they can to tell whether the circumstances are safe. The only difference is that Rafa is, despite his best efforts, unfortunately wrong about them. This unfortunate fact of the matter does not, however, make Sara more creditable than Rafa for her true belief.

Conclusion

To recap, the recognition of higher forms of knowledge—reflective and full knowledge offers a way of capturing what is going wrong for the subject in the Jokester case. In such a case, the subject has animal knowledge but lacks reflective and, thus, full knowledge. The subject’s belief is apt but it is not meta-apt. As I have emphasized, the reason why it is only apt is *not* that the doxastic condition of meta-aptness is violated. In the Jokester case, we simply do not know whether the subject possesses a doxastic attitude regarding the risks of believing something falsely on the first-order level. The reason why the subject’s belief is only apt is that

the modal condition of meta-aptness is violated. The belief is not environmentally safe. This, as just said, explains why the cognitive situation of the subject in the Jokester case seems to be lacking. His cognitive situation is not ideal *because* the subject does not have reflective or full knowledge. He possesses only animal knowledge. Most importantly for my purpose here, such an explanation presupposes that possessing full knowledge is somehow *better* than possessing animal knowledge. But why would full knowledge be better than mere animal knowledge? It is at this point that the new value problem makes its appearance. The upshot of this contribution is that the virtue-reliabilist's solution to the 'classical' value problem cannot be extended to solve this new problem. The superiority of full knowledge over animal knowledge cannot be captured in terms of credit. As I have argued, the subject does not deserve *more credit* for her first-order true belief when she has full knowledge. Sara's first-order true belief in "Secure" is not more creditable to her than Rafa's belief in "Risky".

¹ Some have doubted whether knowledge that p is a better cognitive state than true belief that p. See e.g. Dutant 2012.

² See Olsson 2007, Olsson 2011 for the view that knowledge, understood as reliable true belief, is in fact practically more valuable than any subset of its parts.

³ The reason why knowledge cannot be situated on the same evaluative continuum as a mere true belief and a Gettierized+true belief is, according to Pritchard (2009a: 4), that it would fail "to explain why the long history of epistemological discussion has focused specifically on the stage in this continuum of epistemic value that knowledge marks rather than some other stage (such as a stage just before the one marked out by knowledge, or just after)."

⁴ This is not the only possible solution. Rather than relying on the distinction between instrumental and final value, another way of solving the tertiary value problem is to make use of the distinction between various spectra of evaluation. Consider a pluralist about values who thinks that there are at least three spectra of evaluation that apply to beliefs: the epistemic

spectrum (that brings the various epistemic values together), the practical spectrum (that gathers the different practical values), and the moral spectrum. Now, suppose the pluralist in question manages to show that knowledge possesses some moral value—whether final or instrumental—that the states that fall short of knowledge do not, in contrast, possess. This amounts to establishing that knowledge is characterized by a different *kind* of value. This is another way of solving the tertiary value problem.

⁵ The reverse is also true. One cannot solve the value problems without, thereby, imposing an answer on the analysis of knowledge. Usually, however, epistemologists address the problem of analysis first. They then try to find an explanation of the superior value of knowledge that is compatible with their analysis.

⁶ For interesting replies on behalf of reliabilists, see e.g. Goldman and Olsson 2009; Olsson 2007, Olsson 2011.

⁷ Virtue-reliabilism is to be distinguished from virtue-responsibilism according to which virtues are not reliable competences like memory or vision but acquired character-traits. See Battaly 2008; Greco and Reibsamens forthcoming. Sosa casts doubt on the distinction between virtue-reliabilism and virtue-responsibilism in Sosa 2015 (chapter 2).

⁸ According to Sosa, performances are characterized by constitutive aims. See Sosa 2015: 65; Sosa 2015: 88, note 5.

⁹ Still another way of saying the same thing is in terms “achievement”. Successful performances that result from the exercise of a competence are *achievements* and achievements are better than successful performances that are due to luck. Pritchard (2009b: 409) suggests this formulation. It is now very common; see Greco (2010); Greco and Reibsamens forthcoming: 9.

¹⁰ Note that even if virtue-reliabilists thought that mere true belief was in some sense finally valuable, they could still try to respond to the tertiary value problem as follows. They could contend that although mere true belief and knowledge are both finally valuable, the final value

that characterizes knowledge is credit value and thus belongs to a different kind of final value than the kind of final value characterizing true beliefs.

¹¹ The virtue-reliabilist's answer to the value problems is not, however, free from difficulties. See Kelp and Simion forthcoming; Meylan 2013.

¹² Sosa's Jokester case is analogous to Goldman's famous "Barns case" initially presented in Goldman 1976.

¹³ See Greco's pragmatic attribution theory (Greco 2012; Greco and Reibsamens forthcoming) and Pritchard's anti-luck virtue epistemology (Pritchard forthcoming). Only Sosa's solution to the Jokester case appeals to the differentiation of various levels of knowledge and, thereby, leads to the emergence of what I call the *new value problem*.

¹⁴ Sosa often emphasizes that the second-order belief that regards the risks must itself be apt, that is, *true* in order for the first-order performance to be meta-apt (Sosa 2007: 24; Sosa 2015: 69).

¹⁵ In addition to aptness and meta-aptness, Sosa distinguishes a third normative standard that applies to performances: full aptness. As aptness is the standard of animal knowledge, meta-aptness the standard of reflective knowledge, full aptness is the normative standard satisfied by *full knowledge*. A performance P is fully apt if and only if "it is guided to aptness through the agent's reflectively apt risk assessment." (Sosa 2015: 9) Thus, full aptness consists in a specific combination of aptness and meta-aptness that occurs—to say it briefly—when a performance is apt *because* it is meta-apt. For sake of simplicity, I shall leave full aptness aside in this contribution.

¹⁶ More precisely, a subject S fully knows, according to Sosa, when S's animal knowledge *results* from S having aptly assessed the risk of being wrong, that is, from S having reflective knowledge. Her animal knowledge, in some sense, has to be due to her reflective knowledge. See Sosa 2015: 66-74.

¹⁷ See also e.g. Sosa 2009b: 14 ; Sosa 2015: 85.

References

- Dutant, J. (2012) “The Value and Expected Value of Knowledge”, *Dialogue* 51: 1, 141-62.
- Gettier E. (1963) “Is Justified True Belief Knowledge?”, *Analysis* 23: 6, 121-123.
- Goldman, A. (1976) “Discrimination and Perceptual Knowledge”, *Journal of Philosophy* 73, 771-91.
- (1979) “What Is Justified Belief?”, in G.S. Pappas (ed.) *Justification and Knowledge*, Dordrecht: Reidel, 1-23.
- Goldman, A. and Olsson, E.J. (2009), “Reliabilism and the Value of Knowledge”, in A. Haddock, A., Millar, A., and Pritchard, D. (eds.) *Epistemic Value*, Oxford: Oxford University Press, 19-41.
- Greco, J. (2003) “Knowledge as Credit for True Belief”, in M. DePaul, and L. Zagzebski (eds.) *Intellectual Virtue: Perspectives from Ethics and Epistemology*, New York: Oxford University Press, 111-134.
- (2010) *Achieving Knowledge: a Virtue-Theoretic Account of Epistemic Normativity*. Cambridge: Cambridge University Press.
- (2012) “A (Different) Virtue Epistemology”, *Philosophy and Phenomenological Research* 85, 1-26.
- (forthcoming). “Knowledge, Virtue and Safety”, in M. A. Fernández (ed.) *Performance Epistemology*, New York: Oxford University Press.
- Kelp, C. and Simion, M. (forthcoming) “The Tertiary Value Problem and the Superiority of Knowledge”, *American Philosophical Quarterly*.
- Kvanvig, J. (2003) *The Value of Knowledge and the Pursuit of Understanding*, Cambridge: Cambridge University Press.

Meylan, A. (2013) “The Value Problem of Knowledge: an Axiological Diagnosis of the Credit Solution”, *Res Philosophica* 90:2, 261-75.

Olsson, E. (2007) “Reliabilism, Stability, and the Value of Knowledge”, *American Philosophical Quarterly*, 44: 4, 343-55.

----- (2011) “The Value of Knowledge”, *Philosophy Compass* 6: 12, 874-83.

Pritchard, D. (2007) “Recent Work on Epistemic Value”, *American Philosophical Quarterly* 44: 2, 85-110.

----- (2009a) “The Value of Knowledge”, *Harvard Review of Philosophy* 16, 2-20. ^[1]_[SEP]

----- (2009b) “Apt Performance and Epistemic Value”, *Philosophical Studies* 143, 407-16.

----- (forthcoming) “Knowledge, Luck and Virtue: Resolving the Gettier Problem”, in C. de Almeida, P. Klein and R. Borges (eds.) *The Gettier Problem*, Oxford: Oxford University Press.

Rabinowicz, W. and Rønnow-Rasmussen, T. (1999) “A Distinction in Value: Intrinsic and For Its Own Sake”, *Proceedings of the Aristotelian Society* 100, 33-52.

Sosa, E. (2007) *A Virtue Epistemology*, Oxford: Oxford University Press.

----- (2009a) *Reflective Knowledge*, Oxford: Oxford University Press.

----- (2009b) “Knowing Full Well: The Normativity of Beliefs as Performances”, *Philosophical Studies* 142 :1, 5-15

----- (2011) *Knowing Full Well*. Princeton: Princeton University Press.

----- (2015) *Judgment and Agency*. New York: Oxford University Press.

Zagzebski, L. (2003) “The Search for the Source of the Epistemic Good”, *Metaphilosophy* 34, 12-28; and reprinted in M. Brady and D. Pritchard (eds.) *Moral and Epistemic Virtues* Oxford: Blackwell, 13-28.