

Epistemic Dependence and Cognitive Ability

Fernando Broncano-Berrocal

(published in *Synthese*)

Abstract

In a series of papers, Jesper Kallestrup and Duncan Pritchard argue that the thesis that knowledge is a cognitive success because of cognitive ability (*robust virtue epistemology*) is incompatible with the idea that whether or not an agent's true belief amounts to knowledge can significantly depend upon factors beyond her cognitive agency (*epistemic dependence*). In particular, certain purely modal facts seem to preclude knowledge, while the contribution of other agents' cognitive abilities seems to enable it. Kallestrup and Pritchard's arguments are targeted against views that hold that all it takes to manifest one's cognitive agency is to properly exercise one's belief-forming abilities. I offer an account of the notion of cognitive ability according to which our epistemic resources are not exhausted by abilities to produce true beliefs as outputs, but also include dispositions to stop belief-formation when actual or modal circumstances are not suitable for it (*precautionary cognitive abilities*). Knowledge, I argue, can be accordingly conceived as a cognitive success that is also due to the latter. The resulting version of robust virtue epistemology helps explain how purely modal facts as well as other agents' cognitive abilities may have a bearing on the manifestation of one's cognitive agency, which shows in turn that robust virtue epistemology and epistemic dependence are not incompatible after all.

The problem is the following. On the one hand, there is a well-established view, *robust virtue epistemology*, which understands knowledge as a sort of cognitive success (i.e., a true belief) that is because of a set of properly exercised cognitive abilities.¹ The view has proved useful in solving important problems

¹ For two influential versions of robust virtue epistemology, see Greco (2010; 2012) and Sosa (2007; 2015). Greco opts for a causal explanatory reading of the 'because of' relation and accordingly understands knowledge as a cognitive success that is *explained* by one's cognitive abilities. Sosa prefers a metaphysical reading according to which knowledge is a cognitive success that *manifests* one's cognitive abilities. The version of robust virtue epistemology that I will give in section 3 endorses the metaphysical reading.

such as the Gettier problem. After all, Gettier-style cases involve beliefs that are true, not in virtue of an exercise of cognitive ability, but in virtue of a lucky coincidence. More generally, the view reveals an important fact about knowledge: all it takes to know a proposition is within the limits of *cognitive agency*; in other words, there are no facts over and above the exercise of cognitive ability that determine in a significant way whether one possesses propositional knowledge. On the other hand, there is a thesis, *epistemic dependence*, according to which knowledge can be dependent upon factors which are completely external to one's cognitive agency. More specifically:

Epistemic Dependence Thesis

Whether or not an agent's undefeated epistemic support for her true belief that p counts as knowledge that p can significantly depend upon factors outwith her cognitive agency. (Pritchard 2015: 306).

The problem is that, by the looks of it, epistemic dependence and robust virtue epistemology are incompatible: as long as one deems the former true, one has to reject the latter, and the other way around. After all, robust virtue epistemology claims that whether or not we possess knowledge exclusively depends on how we exercise our cognitive abilities, and namely on whether or not the fact that our beliefs are true manifests or is explained by such an exercise, whereas the epistemic dependence thesis entails that factors beyond cognitive agency—i.e., factors that are not related to the exercise of cognitive ability—can indeed explain why we possess knowledge.

In a series of recent papers, Jesper Kallestrup and Duncan Pritchard (henceforth K&P) have insistently defended the incompatibility claim. In particular, their main point is that the epistemic dependence thesis is incompatible with robust virtue epistemology in a way that proves the theory both too strong and too weak for knowledge. In order to prove the former claim, they appeal to cases in which factors that have nothing to do with the exercise of cognitive ability enable the acquisition of knowledge. In order to prove the latter, they propose cases in which modal facts that are independent of the exercise of cognitive ability preclude the possession of knowledge. Either way, the upshot, according to K&P, is that there is no plausible conception of cognitive agency

on which such knowledge-precluding modal facts and knowledge-enabling factors may have a bearing on the manifestation of one's cognitive agency.²

Robust virtue epistemologists could simply react to K&P's challenge by rejecting the epistemic dependence thesis. The problem with this move is that the examples of epistemic dependence that K&P use to justify the incompatibility claim are plausible. In this way, robust virtue epistemologists have no other option but to explain K&P's cases away.

In this paper, I will argue that while K&P's arguments are targeted against virtue accounts based on a simple understanding of cognitive abilities as belief-forming dispositions, our epistemic resources also comprise cognitive abilities that are not aimed at producing true beliefs as outputs but at stopping belief-formation when actual or modal circumstances are not suitable for it (*precautionary cognitive abilities*). I will accordingly propose an account of knowledge in terms of all our epistemic resources (i.e., in terms of both kinds of cognitive abilities) and argue that it gives rise to a conception of cognitive agency on which purely modal facts as well as other agents' cognitive abilities may have a bearing on the exercise and manifestation of one's cognitive agency. This account, I contend, preserves the main tenet of robust virtue epistemology—that knowledge is a cognitive success because of cognitive ability—and is compatible with the epistemic dependence thesis.

1 Epistemic Dependence: Positive and Negative

Some initial caveats about epistemic dependence are in order. Firstly, it is not K&P's claim that all knowledge is dependent in the stipulated way, but that it *may* be so dependent. Accordingly, the burden of proof is on robust virtue epis-

² The idea of epistemic dependence is discussed at length in Pritchard (2015). In their 2014 paper, Kallestrup and Pritchard introduce their key case of knowledge-precluding epistemic dependence: the epistemic twin earth case (see below) and try to block possible replies by robust virtue epistemologists. Their 2012 paper is mainly concerned with knowledge-enabling epistemic dependence in testimony cases. Their 2013a and 2013b papers represent good summaries of their objections to robust virtue epistemology. Finally, in Kallestrup and Pritchard (2016) they focus on versions of robust virtue epistemology that account for knowledge in terms of the notion of manifestation of cognitive ability. All objections, arguments and cases by Kallestrup and Pritchard that I will discuss in this paper are spread over these works. To avoid constant repetition of these references, I omit them in what follows (for specific points, however, I do refer the reader to the papers where they are discussed in-depth).

temology, which claims that there are no significant factors beyond cognitive agency fixing whether or not a given true belief amounts to knowledge. Secondly, as K&P point out, the epistemic dependence thesis must be understood neither as the thesis that knowledge of p depends on there being a truth-maker for p , nor as the thesis that the mental states constituting that knowledge are individuated by physical or sociolinguistic external factors.³ The epistemic dependence thesis is properly epistemic, in that it holds that factors having nothing to do with the exercise of one's cognitive abilities bear on whether or not one knows.

K&P explain that epistemic dependence comes in two guises, depending on whether such factors enable or prevent the possession of knowledge. Knowledge-precluding epistemic dependence (or *negative epistemic dependence*) seems to prove robust virtue epistemology too weak. For suppose that you exercise your cognitive abilities in the right way but your cognitive success does not amount to knowledge because of being dependent on knowledge-precluding factors beyond your cognitive agency. Robust virtue epistemology would incorrectly predict that you know. By contrast, knowledge-enabling epistemic dependence (or *positive epistemic dependence*) seems to prove robust virtue epistemology too strong. For suppose that your cognitive success is partly down to knowledge-enabling factors. It seems that however you exercise your cognitive abilities this wouldn't be enough to justify the claim that your cognitive success is solely due to them. In this way, the view would incorrectly predict that you don't know.

In order to show that this is effectively the case, K&P need to provide plausible examples of both negative and positive epistemic dependence. They illustrate positive epistemic dependence with certain cases of *testimonial knowledge*. For example, when confronted with a sincere and a reliable speaker, a standard hearer need not exhibit more than a minimal level of cognitive ability to gain knowledge—which is different from saying that she knows by blindly trusting her interlocutor. This means, according to K&P, that factors beyond her cognitive agency (namely, the informant's cognitive abilities) enable to a large degree her cognitive success. The same point can be also made with pro-

³ See especially Kallestrup and Pritchard (2013a) for this point.

tective social environments. K&P appeal to Sanford Goldberg's idea that some communities or third party epistemic agents monitor and police testimonial exchanges in such a way that they ensure that hearers are mostly exposed to reliable speakers (primary schools would be an example of this).⁴ Again, when hearers acquire knowledge in this way little cognitive ability is manifested on their behalf. Instead, their true beliefs are largely due to factors beyond their cognitive agency (e.g., the friendliness of the environment).

K&P's example of negative epistemic dependence is more complicated and requires the introduction of some terminology.⁵ They distinguish between local, regional and global environmental conditions. Suppose that an agent *S* forms a belief at time *t*. *Local environmental conditions* are the conditions of the environment where *S* is when forming the belief at *t*—roughly: the actual world. *Regional environmental conditions* are the conditions of the environment where *S* might easily have been when forming the belief at *t* or at times close to *t*—roughly: nearby possible worlds. *Global environmental conditions* are the type of conditions in which *S* typically is when forming the same kind of beliefs (e.g., visual beliefs) about the same kind of propositions (e.g., propositions about colors).

With these distinctions in place, K&P modify Putnam's classic twin earth thought experiment to give a plausible example of negative epistemic dependence.⁶ In the original twin earth case, there is an agent, *S*, who forms the belief that water is wet, say, when washing her hands. *S*'s intrinsic physical duplicate, *S**, inhabits twin earth, which is exactly like earth except for the fact that watery stuff is not H₂O but XYZ. Water and twin-water are perceptually indistinguishable and both agents lack knowledge of chemistry. The point of the case is that when *S** forms a belief in the same proposition (that water is wet), her belief is different from *S*'s insofar as their respective beliefs are not exclusively individuated by their intrinsic physical properties, but by their contents, which respectively refer to XYZ and H₂O. This shows, according to Putnam's famous slogan, that 'meaning ain't in the head'.

⁴ See Goldberg (2011). For relevant discussion of third-party epistemic dependence, see especially Kallestrup and Pritchard (2012).

⁵ See especially Kallestrup and Pritchard (2014) for this case.

⁶ See Putnam (1973) for the original example.

K&P offer an epistemic version of the twin earth thought experiment. In the modified case, there are the same two agents, *S* and her intrinsic physical duplicate *S**, who are perceptually and conceptually competent but chemically ignorant. Relative to their respective global environmental conditions on earth and epistemic twin earth, *S* and *S** are equally highly reliable in forming true water-beliefs, i.e., true beliefs about the presence of H₂O.

On earth there is only H₂O, which means that all watery stuff *S* encounters in her local, regional and global environmental conditions is H₂O. *Epistemic twin earth*, however, is a trickier place. On its surface, there is *typically* H₂O (i.e., *S**'s global environmental conditions include water), which explains why *S**'s perceptual capacities are reliably tuned to its presence. When forming the belief that water is wet at time *t*, *S**'s local (i.e., actual) environment also includes water. But *S** is in a region of epistemic twin earth where she could easily have encountered twin-water (i.e., XYZ) at *t* or at times close to *t*.

K&P claim—and most commentators will agree—that, while *S*'s belief that water is wet amounts to knowledge, *S**'s belief does not. The explanation is simple: it is a matter of luck that *S** gets thing right. More generally, the moral that K&P draw is, firstly, that robust virtue epistemology is too weak, because it incorrectly predicts that *S** knows; secondly, that it fails to offer a conception of cognitive agency on which purely modal facts (such as the presence of XYZ in the regional environment) bear on the manifestation of one's cognitive agency.

My point is that K&P can draw their conclusions only on the assumption that all it takes to manifest one's cognitive agency is to exercise one's *belief-forming* abilities in the right way. Indeed, K&P's arguments are targeted against virtue accounts that understand cognitive abilities simply as reliable dispositions to form true beliefs (e.g., Greco 2010), so the assumption is understandable. In the next section, however, I will offer an alternative conception of cognitive agency—or of proper exercise of cognitive ability—according to which our epistemic resources are not exhausted by our belief-forming abilities: we also count with what I call *precautionary abilities*.

2 Precautionary Abilities

The main tenet of robust virtue epistemology is that knowledge is a sort of cognitive success that is because of a set of properly exercised cognitive abilities, but it does not say anything on how to understand the notion of cognitive ability itself. In particular, it does not say that knowledge should be exclusively due to our belief-forming abilities.

In order to get a better understanding of the notion of cognitive ability, we need to consider an alternative way to individuate abilities. Abilities and cognitive abilities are typically classified in terms of the domains in which they are used: cooking abilities, musical abilities, reasoning abilities, visual abilities, and so on. But a different way to individuate abilities is by the *functional role* the play.⁷

Accordingly, some abilities play a *task-completion role*, in that they are aimed at completing tasks, such as the ability to dice an onion or the cognitive ability to form true beliefs about propositions pertaining to some domain. Other abilities play a *precautionary role*: they do not directly accomplish tasks (dicing onions, forming true beliefs) by producing outputs (onion dices, true beliefs), but prevent the production of such outputs when the conditions are not suitable for it. In other words, precautionary abilities prevent the exercise of corresponding task-completion abilities when their conditions are not appropriate.

By way of illustration, a reliable precautionary ability corresponding to the ability to dice onions should prevent one from dicing onions in circumstances that are inappropriate for it, e.g., when one's knife is not sharp enough or when one's cutting board is, say, balancing on the tip of a cone. Another example is the ability of a sniper to refrain from shooting when the wind is such that it would deviate bullets way off target. In the cognitive case, we are typically disposed to avoid forming visual beliefs about objects being such-and-such when the light conditions are inadequate. This is a *precautionary cognitive ability*.

Insofar as the reliability of abilities is relativized to a certain set of conditions, abilities might be broad or narrow depending on the kind of condi-

⁷ In Broncano-Berrocal (2017), I discuss this way to individuate abilities at length.

tions they are relativized to. Most abilities (precautionary or task-completion abilities) are *broad*, because they are reliable relative to global conditions, i.e., relative to the internal and external conditions in which agents normally or typically use them or relative to the type of conditions in which they acquired or developed them in the first place. For example, a white surface illuminated with red light to make it look red is not what we would describe as 'normal conditions' for the ability to determine the color of objects, because neither they are the typical circumstances where this ability is employed, nor they correspond to the sort of conditions where such a visual ability was acquired or developed in the first place. Interestingly, there is a key difference between the sort of conditions that respectively count as 'normal' for task-completion and precautionary abilities: namely, for the former normal conditions are adequate circumstances for completing tasks, whereas for the latter they are inappropriate conditions for completing those tasks.

Abilities may be also *narrow*, in the sense that they may not be reliable relative to global conditions but to local and regional conditions instead (roughly: reliable only in the actual and in nearby possible worlds). A cyclist, for example, might win a race only because she has used drugs that temporarily boost her endurance. Her improved ability to pedal is not reliable relative to global conditions, where those drugs are typically unavailable, but relative to her actual circumstances and very similar ones. Analogously, the ability to visually form true beliefs might be locally improved by binoculars, microscopes, pills increasing one's visual acuity, and so on.

In the same way, a precautionary ability may be as narrow as to be reliable only relative to very specific circumstances. If a trusted source informs you that many works of the museum that you are about to visit have been randomly replaced with forgeries, you won't probably believe that they are authentic or that they have been made or painted by such-or-such an artist. Your locally acquired disposition not to form such beliefs counts as a precautionary cognitive ability insofar as it reliably prevents you from using your belief-forming abilities in unfriendly circumstances. By the same token, if you come to know that a prankster has set your watch one hour back, you are disposed to avoid forming false beliefs about the time when consulting it. In view of this, it doesn't matter whether precautionary abilities are grounded on an

innate physical basis, an arduous *learning process*, the competent use of a *device* or reliable *information* from a trustworthy informant. What matters is that one comes to be reliably disposed to stop the exercise of a corresponding task-completion ability when the circumstances are not suitable for it, however this disposition is grounded.⁸

It is important to keep in mind that the sense in which precautionary abilities are reliable is different from the sense in which task-completion abilities are reliable. To see this, let's examine the notion of reliability in more detail. In its most general sense, reliability can be defined as a tendency to produce outputs of one type such that the ratio of produced outputs of that type to produced outputs of a different type is sufficiently high. In addition, the notion is typically linked to a standard of assessment that categorizes the different output types either as good or bad, successful or unsuccessful, accurate or inaccurate, true or false, correct or incorrect, and so on.⁹

Both task-completion and precautionary abilities are reliable in this *general sense*. The former are reliable in this way because their ratio of successful outcomes (i.e., of completed tasks) to unsuccessful outcomes (i.e., to uncompleted tasks) is sufficiently high.¹⁰ For example, one does not have the ability to ride a bike unless one exhibits a tendency to successfully move by bike (e.g., from point A to point B without falling) such that the ratio of successful to unsuccessful journeys is sufficiently high. What is important to keep in mind is that, while precautionary abilities are also reliable in this general sense (i.e., in the sense that they also exhibit a high ratio of successful to unsuccessful outputs), their outputs are different in kind from the outputs produced by task-completion abilities—basically because precautionary and task-completions abilities serve different functions—, and this marks a difference in the way the two types of abilities are *specifically reliable* (viz., reliable in com-

⁸ Of course, this disposition must be adequately integrated within one's cognitive system. See Breyer and Greco (2008) for relevant discussion on the notion of cognitive integration.

⁹ This way to understand reliability is a generalization from the way the term 'reliability' is typically used in epistemology. In quantitative research (e.g., in experimental psychology), reliability is understood more simply as consistency or repeatability of a certain type of outcome (e.g., in an experiment). It is the term 'validity' that is used instead to refer to the tendency to produce outcomes of a certain type such that they meet the requirements of a standard of assessment (e.g., being successful, accurate, correct, and so on).

¹⁰ What counts as 'sufficiently high' depends on the domain of the relevant ability.

pleting the specific functions they serve to complete). Let's analyze these two different specific notions of reliability in turn.

The successful and unsuccessful outputs of task-completion abilities are respectively completed and uncompleted tasks. This means that their specific reliability is *reliability in completing tasks* (e.g., in dicing onions in the right way, in forming true beliefs) and that the way they fall short of it is by not completing those tasks (e.g., by dicing onions in the wrong way, by forming false beliefs).

By contrast, the function of precautionary abilities is, as we have seen, to prevent task-completion endeavors in inappropriate circumstances. In other words, their specific reliability is *reliability in preventing task-completion abilities from producing their own outputs in inappropriate circumstances for it*. This means that a successful output of a precautionary ability consists in preventing the manifestation of a corresponding task-completion ability when the circumstances (i.e., the circumstances for completing tasks) are inappropriate. Interestingly, unsuccessful outputs may be of two types: *either* if the relevant precautionary ability fails to prevent task-completion abilities from producing outputs when the circumstances are inappropriate for completing tasks, *or* if it prevents task-completion when the circumstances are appropriate for it. In the former case, we can describe the failed performance of the precautionary ability as *reckless*. In the latter, we can describe it as *overprotective*. These are the two main ways in which precautionary abilities may fall short of its specific reliability.

By way of illustration, suppose that the light conditions are not good enough to determine whether the surface in front of you is orange or red (i.e., the circumstances are inappropriate for the reliable exercise of your visual belief-forming abilities). If you nevertheless form a belief about the color of the surface, say, by mere guessing, the performance of your corresponding precautionary cognitive abilities can be deemed unreliable in the *reckless* sense (independently of whether or not you end up getting things right). By contrast, suppose this time that the light conditions are excellent but you encounter a misleading defeater, namely misleading but compelling evidence that you are facing a white surface illuminated with red light. This disposes you to avoid forming a true visual belief (that the surface is red) under excellent light con-

ditions. Such a disposition, if manifested, fails to fulfill an adequate precautionary role in the *overprotective* sense (the circumstances are appropriate but belief formation is prevented). In this way, to be reliable precautionary abilities must stop the activation of corresponding task-completion abilities *only* in circumstances that are inappropriate for it. In appropriate circumstances, they are reliable as long as they remain inactive.¹¹

That precautionary abilities are reliable does not mean of course that they are infallible. Precautionary abilities are *fallible* precisely because there might be appropriate conditions for task-completion in which they end up preventing the manifestation of the corresponding task-completion abilities, or else there might be inappropriate conditions for task-completion in which they nevertheless let the corresponding task-completion abilities produce an output. Indeed, as we will see, extraordinary inadequate conditions call for extraordinary precautionary abilities that humans typically lack, which reveals the limits of our cognitive agency.

If our resources when it comes to completing tasks in general comprise both kinds of abilities, it is natural to think that our *epistemic resources* also comprise both types of cognitive abilities. Indeed, one not only manifests one's agency by skillfully acting in ways that serve to accomplish tasks when the circumstances are appropriate for it, but also by competently refraining from carrying them out when they are not. Therefore, it is also natural to think that one manifests one's *cognitive agency* by exercising one's reliable belief-forming dispositions in appropriate circumstances as well as by activating one's cognitive precautionary abilities to withhold belief in unsuitable ones.

Furthermore, as we have seen, when an agent exercises her task-completion abilities under appropriate circumstances and succeeds in some endeavor, whatever protective abilities the agent has also competently succeed in not being

¹¹ Another important remark about the specific reliability of precautionary abilities: if a precautionary ability is broad or narrow so will be its reliability. A broad disposition to stop visual belief formation when the light conditions are bad is broad enough to be reliable in any situation in which the light conditions are bad. A narrow disposition to stop belief formation on one occasion based on trustworthy testimonial information that some sort of deception is going on is reliable only relative to that situation. Suppose that one erroneously believes such information to be true of a large number of situations (e.g., think of believers in conspiracy theories). Stopping belief formation in those situations is reckless, because they are beyond the range of reliability of the locally acquired disposition to stop belief formation in the specific circumstances the information is about.

overprotective by not stopping it. Therefore, an agent's success may also manifest her precautionary abilities in appropriate circumstances. The same idea applies to precautionary cognitive abilities. If the light conditions are good, one's visual precautionary abilities—which would stop belief formation under poor light conditions—competently succeed in not being overprotective by not stopping belief formation when they shouldn't.

3 A Robust Enough Virtue Epistemology

A more comprehensive *version of robust virtue epistemology* takes into account the foregoing ideas and acknowledges the epistemic role played by precautionary cognitive abilities. Recall that the main tenet of robust virtue epistemology is that knowledge is down to the exercise of cognitive ability, which is different from saying that it must be exclusively due to one's belief-forming dispositions. Since our epistemic resources also include precautionary cognitive abilities, one way to understand knowledge has it that an agent's true belief amounts to knowledge just in case her cognitive success is because of (in the sense that it manifests) all epistemic resources, i.e., not only her reliable belief-forming dispositions but also her reliable precautionary cognitive abilities.

This view assumes a conception of abilities as dispositions of agents and accordingly understands cognitive abilities as the kind of dispositions that are relevant to knowledge and other epistemic standings such as epistemic justification or understanding. Dispositions, in general, are adequately triggered under certain normal or appropriate conditions, which means that the resulting states manifest them only under such conditions. For example, salt manifests its solubility when stirred into water (appropriate conditions) but not when stirred into gasoline (inappropriate conditions).

Analogously, an agent's cognitive success (i.e., a true belief) manifests her cognitive abilities and hence qualifies as knowledge *only under certain appropriate conditions*. After all, those conditions are the conditions under which the agent's cognitive abilities are reliable, which in the case of belief-forming abilities are, as we have seen, appropriate conditions for belief formation and,

in the case of precautionary abilities, appropriate conditions for preventing belief formation or, in other words, inappropriate conditions for belief formation. Since the view says that knowledge is a cognitive success that manifests both kind of abilities, this means that the circumstances for the relevant belief-forming and precautionary abilities must be appropriate (in the two senses of appropriateness just distinguished).

In the case of precautionary abilities, circumstances must be appropriate both in the local and the regional environment. Indeed, it is not surprising that the kind of precautionary cognitive abilities that give rise to knowledge must be sensitive not only to the inappropriateness of the agent's actual environment but also of the environment in which the agent could easily be (i.e., her regional environment). After all, knowledge is commonly thought to have a modal dimension.

To see this, consider one of the examples of the previous section. Suppose that you enter a museum where, unbeknownst to you, many works have been randomly replaced with forgeries. Suppose that you look at a real Kandinsky painting. Intuitively, your true belief does not amount to knowledge because you could easily have formed that belief by looking at any of the many forgeries in the room. But suppose that this time you count with the relevant precautionary abilities. For example, you might wear computerized glasses that alert you every time you look at a forgery, or you might be an art expert who is able to detect signs of forgery. The competent use of the device or your forgery detection skills dispose you to stop belief formation when the circumstances are inappropriate for it—forgeries in general are not part of the conditions relative to which standard art-related cognitive abilities are reliable. That being the case, you can know that the painting in front of you is a Kandinsky painting when you are in fact before one.¹²

Some final caveats are in order. The previous case shows that a disposition to stop belief formation when needed might be grounded on the use of a device. In keeping with foregoing discussion, it does not matter whether precautionary cognitive abilities are *innate*, *learned* or *temporarily possessed* in virtue of relevant information from a trustworthy speaker or a reliable device.

¹² This case is analogous to the much-discussed fake barn case (Goldman 1979).

From an epistemic point of view, all that matters is that the agent counts with a cognitive disposition fulfilling a precautionary role in the right way.

In addition, note that the model of cognitive agency defended here is *two-level* in that it comprises one type of cognitive abilities that are *first-order* (viz., belief-forming abilities) and one type of cognitive abilities that are *second-order* (viz., precautionary abilities).¹³ However, this doesn't mean that the latter need to be *reflective* (in the sense that agents need not be aware of the effect they have on belief formation) or *metacognitive* (in the sense that the kind of control exerted by precautionary abilities over first-order belief-forming processes need not be exclusively based on internal monitoring of the latter). Many precautionary abilities *are* reflective or metacognitive, but this is not necessarily the case.

Firstly, requiring precautionary abilities to be necessarily reflective might run into the risk of over-intellectualization. It seems implausible to suppose that in all cases in which localized information is put into action to prevent belief-formation in a specific situation a reflective process is always involved. For example, if someone replaced some of the fresh eggs in your fridge with boiled eggs, you wouldn't notice the difference by sight—such conditions are inappropriate for your visual belief-forming abilities. But suppose that you counted with a device that automatically indicated whether eggs are fresh or boiled, e.g., by showing on the screen the words 'fresh' or 'boiled'. Using this device would be a reliable method for forming true beliefs, but also for avoiding false ones. In other words, the device would dispose you not only to form true beliefs, but also not to form false beliefs in conditions that are inappropriate for your innate belief-forming abilities. The crux of the matter is that the use of this device might be a very automatic and fast (i.e., a type-1) process, as when you consult your watch or cell phone to know the time. So we shouldn't presume that precautionary abilities are necessarily reflective.

To shed more light on why precautionary abilities are not necessarily metacognitive, it will be useful to consider a complementary view to the account of

¹³ Belief-forming abilities are first-order because they are dispositions to form true beliefs. Precautionary cognitive abilities are second-order because they are dispositions to prevent the manifestation of other dispositions. However, it is important to keep in mind that the fact that the former are first-order and the latter second-order is not what distinguishes them. What distinguishes them is the fact that they play distinct functional roles.

knowledge defended here. Michaelian (2012; 2016) defends a view according to which knowledge and justified belief are produced by overall reliable belief-forming systems with the following two-level structure: a first-order level produces information that serves as content of the belief and a second-order level either endorses or rejects the information produced and thus either triggers or prevents belief formation. The overall reliability of the system is what counts for knowledge and justification and is determined by the interaction of the first-order and the second-order levels. Crucially, the first-order and second-order levels need not have the same degree of reliability in the specific cognitive tasks they perform in order for the whole system to attain a sufficient degree of reliability able to produce knowledge-level or justification-level beliefs.

Interestingly, Michaelian (2016: 154) distinguishes two ways in which a second-order endorsement mechanism might control belief formation. One is by internally (not necessarily consciously) monitoring the first-order information-producing process, the other is on the basis of outside information, e.g., information about environmental conditions.¹⁴ According to Michaelian, what precisely distinguishes a two-level metacognitive system from a two-level non-exclusively metacognitive system is that in the former but not in the latter endorsement and rejection of object-level information are determined exclusively on the basis of monitoring. Precautionary abilities, as conceived here, may very well be metacognitive—think of any of our innate abilities to prevent visual belief formation when the conditions are not appropriate. But, as we have seen, a precautionary ability may be also grounded on a simple disposition to stop belief formation in very local conditions thanks to trustworthy information about the environment. This marks a difference with Michaelian's view, which only takes into account metacognitive belief-forming systems.¹⁵

¹⁴ For extensive discussion on the notion of monitoring see Nelson and Narens (1990).

¹⁵ A note of clarification is in order. My version of robust virtue epistemology doesn't follow the simple reliabilist template according to which S knows p if and only if S's true belief is produced by a reliable belief-forming process. In this way, my view is not that knowledge is belief that is produced by a two-level belief-forming system whose overall reliability is sufficiently high. Instead, it follows the standard template of virtue accounts according to which an agent S knows p if and only if S believes p truly *because of* her reliable cognitive abilities. The main difference with standard virtue-theoretic views is that the agent must get things right because of reliably exercising not only her first-order belief-forming abilities *but also* her precautionary

As a final caveat, we have seen that a requirement on precautionary cognitive abilities is that they must be sensitive to inappropriate conditions in the local (i.e., actual) circumstances but also in the regional environment (i.e., in nearby possible worlds). This doesn't mean however that they must be sensitive to *any* condition in the regional environment. Some conditions enable the exercise of an ability (*enabling conditions*), whereas some conditions determine its failure or success (*determining conditions*). For example, the fact that you have a functioning visual system enables you to form true beliefs about Kandinsky paintings, whereas the fact that there are forgeries around determines the success or failure of your cognitive performance. The distinction between enabling and determining conditions is important inasmuch as precautionary cognitive abilities must be sensitive only to determining conditions. For instance, they must not stop belief formation if you could easily lose your visual system, but must prevent you from believing that the painting in front of you is by Kandinsky if you could easily encounter a forgery.¹⁶

cognitive abilities: in particular, the correctness of the agent's belief must *manifest* the fact that her relevant belief-forming abilities reliably produce a true belief in appropriate circumstances *and* the fact that her precautionary cognitive abilities reliably avoid stopping belief formation precisely because the circumstances are adequate. Unlike simple reliabilist views, the 'because of' clause (which applies to both types of abilities) serves, as in standard virtue accounts, as a measure against Gettierization. In particular, Gettierized beliefs are not only true but typically reliably formed; however, their correctness *does not manifest* (i.e., it is not because of) the exercise of ability (namely, of belief-forming *and* precautionary ability). In Broncano-Berrocal (2017), I develop this account in more detail. See Broncano-Berrocal (*forthcoming*) for how this view fares better than rivals. A more sophisticated but significantly different version of this view would define knowledge as true belief that manifests overall reliability, i.e., as belief that is true because of the overall reliable belief-forming process that results from the interaction of the agent's first-order belief-forming abilities and second-order precautionary abilities. As Michaelian insightfully explains, the reliability of an overall reliable process allows for different degrees of reliability of the former and the latter in such a way that knowledgeable beliefs can be attained, e.g., with a reckless precautionary policy if the relevant first-order abilities enjoy a greater degree of reliability than usual. For present purposes, I will stick with the first (and arguably simpler) version of the view, keeping in mind that it could be enriched with Michaelian's careful examination of how the reliability of first-order and second-order cognitive processes interact in two-level belief-forming systems.

¹⁶ See Mackie (1974) for a similar distinction between causes and background conditions.

4 Epistemic Dependence and the Limits of Cognitive Agency

The epistemic dependence thesis says that whether or not an agent's undefeated epistemic support for her true belief that p counts as knowledge that p can significantly depend upon factors beyond her cognitive agency. The question that I will address in what follows is whether this thesis is compatible with the previous version of robust virtue epistemology. As we will see, not only is it compatible, but this compatibility also reveals some insights into the limits of our cognitive agency.

4.1 Negative Epistemic Dependence

Let's start with negative epistemic dependence and the epistemic twin earth case. Consider an agent, A , who is infallible in the following way: her belief-forming abilities are such that under a set of appropriate conditions she only forms true beliefs. On epistemic twin earth, A 's local conditions are appropriate, so she gets things right. K&P's point is that no matter how infallible A 's cognitive abilities are, she fails to know because modal facts in her regional environment (namely, the presence of twin water) undermine the actual epistemic support she has for her true belief. On the face of it, they conclude not only that robust virtue epistemology is too weak, but also that there is no plausible conception of cognitive agency on which purely modal facts may have a bearing on the manifestation of one's cognitive agency. My point is that K&P make these claims on the assumption that the only kind of cognitive abilities that are relevant to knowledge are belief-forming abilities.

While it is undeniable that the previous case is a genuine case of negative epistemic dependence, K&P's conclusions don't follow. To see this, let's start by comparing A with another agent B , who, like A , is endowed with infallible belief-forming abilities. The difference with A is that B also counts with infallible precautionary cognitive abilities, in the sense that she always withholds belief in inappropriate conditions for belief formation. This sensitivity to inappropriate conditions also includes unfriendly circumstances in the regional environment (roughly, in nearby possible worlds). Finally, B 's precautionary

abilities are never overprotective, in that they never stop belief formation under appropriate conditions.

Things are different for *B* on epistemic twin earth. Whenever she encounters twin water, *B*'s infallible precautionary abilities make her withhold belief—e.g., such abilities could be based on an innate disposition to stop belief formation when there is twin water around, could involve the competent use of a twin water tracking device or, more simply, knowledge of chemistry. Recall now the museum case. In the same way as you would know that the painting in front of you is by Kandinsky if you had the relevant precautionary abilities (e.g., if you counted with a pair of computerized glasses alerting you of forgeries), *B* plausibly knows that water is wet in her local conditions thanks to her infallible precautionary abilities, i.e., no matter whether there is twin water in her regional environment. After all, if that environment became actual, she would withhold belief.

Therefore, K&P cannot claim this time that no matter how infallible *B*'s cognitive abilities are, her true belief that water is wet doesn't amount to knowledge because of modal factors beyond her cognitive agency. It is in fact the other way around: thanks to her superb epistemic resources, and in particular thanks to her infallible precautionary abilities, *B* would never form such a true belief if modal factors beyond her cognitive agency undermined the epistemic support provided by her also infallible belief-forming abilities (this is the protective role of *B*'s abilities).

In addition, insofar as *B*'s precautionary abilities are sensitive to such modal factors, they are not beyond her cognitive agency. This reveals an interesting fact about cognitive agency and its relationship with negative epistemic dependence: purely modal factors beyond one's cognitive agency undermine knowledge only insofar as one's epistemic resources (and in particular one's precautionary abilities) are not sufficiently tuned to them. Sometimes, our cognitive agency is too limited to play the kind of non-standard precautionary role that unusual unfriendly circumstances such as epistemic twin earth demand.

The result is that the fact that there are plausible cases of negative epistemic dependence (*A*'s case) doesn't exclude that there is also a conception of cognitive agency on which purely modal factors can have a bearing on one's manifestation of cognitive agency. On the alternative conception of cognitive

agency introduced here, an agent fully manifests her cognitive agency just in case she manifests her belief-forming dispositions as well as her precautionary cognitive abilities.

Given that the previous version of virtue epistemology assumes this conception and accordingly accounts for knowledge in terms of the manifestation of both types of abilities, it is able to offer the right prediction about epistemic twin earth. Unlike versions of robust virtue epistemology that understand cognitive abilities simply as belief-forming dispositions, it predicts that *A*, the agent with infallible belief-forming abilities, lacks knowledge. Let's see why.

Unlike *B*'s infallible precautionary abilities, *A*'s standard precautionary abilities are insensitive to her regional conditions. The reason is that circumstances in which twin water is present are inappropriate for the manifestation of *A*'s standard precautionary visual abilities. In particular, the presence of a visually indistinguishable liquid (a fake for that matter) does not belong to the kind of appropriate conditions relative to which standard precautionary visual abilities are reliable in withholding belief (e.g., circumstances in which the light conditions are bad). Since, as we have seen, a cognitive success manifests cognitive ability only under appropriate conditions (in the case of precautionary cognitive abilities, under specific inappropriate conditions for belief formation), *A*'s cognitive success does not amount to knowledge. In this way, cases of negative epistemic dependence neither prove this account of knowledge too weak, nor the main tenet of robust virtue epistemology wrong.

One might still wonder why in K&P's original case, *S* (the agent on earth) knows whereas *S** (the agent on epistemic twin earth) does not. After all, one could argue, *S* and *S** are intrinsic physical duplicates. Here is the explanation: precisely because they are intrinsic physical duplicates (and none of them counts with privileged information about the environment or uses a device that the other lacks), they are equals in terms of precautionary cognitive ability. In this sense, both are endowed with standard precautionary visual abilities that would prevent them from forming beliefs in standard inappropriate conditions for vision (e.g., in the dark). These abilities are reliable on earth, but unreliable on twin earth. In particular, they are unreliable in stopping belief formation when there is twin water around, which explains why *S** doesn't know.

4.2 Positive Epistemic Dependence

Let's consider now cases of positive epistemic dependence, i.e., cases in which one exercises one's cognitive abilities in the right way but acquires knowledge not only in virtue of them but also in virtue of factors beyond one's cognitive agency. In particular, let's focus on testimonial knowledge, which is K&P's key example of positive epistemic dependence.

K&P agree that testimonial knowledge cannot arise out of complete gullibility but hearers need to exhibit a minimal level of cognitive ability. They don't think, however, that it is as significant as to lend support to the claim that in standard testimonial exchanges the hearer's cognitive success is because of her cognitive abilities (as robust virtue epistemologists hold). After all, the speaker's cognitive abilities or the epistemically favorable nature of the circumstances (in cases where hearers are guaranteed to be exposed only to reliable speakers) significantly contribute to it. In this way, K&P conclude, positive epistemic dependence proves robust virtue epistemology too strong.

This conclusion, however, does not generalize to all robust virtue epistemological views, as we will see next. In particular, views that understand knowledge as a cognitive success that *manifests* the agents' cognitive abilities steer clear of it. The first thing to note is that K&P's reasoning is based on the following thesis:

Hindrance Thesis

In a testimonial exchange between a speaker *S* and a hearer *H*, where *S*'s cognitive abilities contribute to *H*'s cognitive success, the manifestation of *H*'s cognitive agency is epistemically less significant than if *H* exercised her cognitive abilities alone.

The hindrance thesis sees the intervention of other agents' cognitive abilities as a hindrance to the manifestation of one's cognitive agency. Suppose that you want to know the location of the train station and ask someone for directions. This person happens to be a sincere and a reliable informant so you form a true belief that plausibly amounts to knowledge. The hindrance thesis says that insofar as your cognitive success manifests not only your own cognitive abilities but also the other person's, the manifestation of your cognitive agency

is epistemically less significant than if you found the station on your own. This thesis, however, is dubious. Let's see why.

Belief-forming abilities are dispositions of agents to reliably produce true beliefs, and true beliefs produced by belief-forming abilities manifest such dispositions insofar as they are exercised in the right way under appropriate conditions. A quite popular idea in the literature on dispositions is that the resulting state of a disposition that has been triggered in normal conditions is not only the manifestation of the disposition but also of at least one dispositional partner.¹⁷ By way of illustration, the result of stirring salt into water (the state of salt in solution in water) is the manifestation of the solubility of salt as well as of the disposition of water to dissolve salt. Such a state manifests both dispositions insofar as it is a common product of both. Moreover, these reciprocal dispositions need each other to produce it.

Testimonial exchanges (at least those that typically give rise to knowledge) can be plausibly understood along the same lines. Instead of viewing the contribution of the speaker's cognitive abilities as a hindrance to the manifestation of the hearer's cognitive agency, a more plausible view has it that they are part of what makes such a manifestation possible.

To see this, consider the following analogy. When salt is in solution in water, the solubility of salt is *no less manifested* because water also manifests its disposition to dissolve salt. On the contrary, the solubility of salt can be manifested because water also manifests its reciprocal disposition to dissolve salt. In the same way, when a hearer forms a true belief after a testimonial exchange, her cognitive abilities are no less manifested because the speaker has also manifested hers. Indeed, the hearer's testimonial cognitive success arises because the cognitive abilities of both parties mutually cooperate. In this way, contrary to what the hindrance thesis says, the speaker's cognitive abilities do not hinder the manifestation of the hearer's cognitive agency: they make it possible instead.

Consequently, robust virtue epistemologists should not fall into the error of understanding knowledge as a cognitive success that *exclusively* manifests the agent's cognitive abilities. That would be too strong. All robust virtue

¹⁷ See Heil (2005), Martin (2008) and Mumford & Anjum (2011) for some representative defenses of this idea.

epistemologists need to be committed to is the idea that knowledge is a cognitive success that non-exclusively manifests them. This opens the door to there being knowledge that arises in reciprocal cooperation with other agents' cognitive abilities. In this way, while it is undeniable that standard testimony-based knowledge represents a paradigmatic example of positive epistemic dependence, robust virtue epistemology and, in particular, manifestation views are not too strong.

It is still puzzling, however, how robust virtue epistemology can explain more complicated cases of testimonial knowledge. In their 2016 paper, K&P give one such a case that is analogous to the epistemic twin earth thought experiment. In the example, the same two agents *S* and her intrinsic physical duplicate *S** encounter honest informants in the same local and global environments. K&P claim that both thus manifest the very same cognitive abilities. The difference between *S* and *S** is in their respective regional environments. Whereas *S* would only encounter honest interlocutors, *S**'s regional environment is populated by dishonest informants. Here is K&P's diagnosis:

The problem should now be manifest. For while the agent who has an epistemically favourable regional environment gains testimonial knowledge, her internal duplicate who has an epistemically unfavourable environment does not gain testimonial knowledge. And yet, as we have seen, there can be no difference in these agents' possession or manifestation of cognitive abilities. Whatever the reason for why they differ in terms of what they know, then, it is not a difference which is a function purely of their manifestation of their cognitive abilities. (Kallestrup and Pritchard 2016: 44).

But there is a difference in the manifestation of their cognitive abilities. In particular, the difference is that only *S* manifests precautionary cognitive ability. This doesn't mean that *S* possesses better precautionary abilities. *S* and *S** are equals in that respect, at least as far as testimonial exchanges are concerned. What happens is that *S*'s regional circumstances are appropriate for her precautionary abilities, whereas *S**'s are not.

To see this, note that the kind of circumstances relative to which their precautionary abilities would reliably prevent them from believing dishonest state-

ments include circumstances in which the transmitted information is remarkably outlandish or in which informants show ostensible signs of dishonesty or unreliability. As a matter of fact, *S*'s local and regional circumstances don't feature any of these conditions. In addition, *S*'s precautionary abilities competently remain inactive in her local environment. Consequently, her cognitive success not only manifests her belief-forming abilities but also her standard precautionary abilities. This explains why *S* knows.

Things are different for *S**. *S**'s Truman-show-like regional environment, in which everyone aims to deceive each other, is certainly not the kind of environment that is appropriate for the standard precautionary abilities *S** is endowed with. Remember that an agent's cognitive success manifests her cognitive abilities (precautionary or non) only under appropriate conditions. Accordingly, the reason why *S** fails to manifest her precautionary abilities relative to her regional circumstances is that they are inappropriate. This explains why *S** doesn't know.

A subject might of course acquire testimonial knowledge in a Truman-show-like situation. But such an epistemically harsh environment, like epistemic twin earth or the museum full of forgeries, calls for exceptional precautionary cognitive abilities that humans typically lack. In standard testimonial exchanges, by contrast, even if the relevant capacities for monitoring dishonesty or deception might not be very reliable, hearers can still acquire testimonial knowledge, firstly, because the base rate of honest speakers is high and, secondly, because they tend to be truth-biased (they are more likely to judge received testimony as true than as deceptive). This means that the precautionary abilities of standard hearers are reliable in *not* stopping belief formation when they shouldn't, which is most of the times.¹⁸

This and the foregoing discussion reveal an important fact about our cognitive agency and its relationship with positive epistemic dependence: factors beyond our cognitive agency may enable the possession of knowledge insofar as the manifestation of our cognitive abilities sometimes needs the reciprocal

¹⁸ See Michaelian (2010) for relevant discussion of empirical findings concerning the claims that our capacities for monitoring deception are not reliable, that most people tend to be honest, and that we are truth-biased, as well as of the implications that these have for the epistemology of testimony.

cooperation of other agents' cognitive abilities. This epistemic dependency on such knowledge-enabling factors often requires no more than a permissive precautionary policy not to call into question the reliability of reciprocal cognitive abilities which contribute to one's cognitive success.¹⁹ Knowledge is readily available in such circumstances. But applying that policy in a less friendly environment would make one trust unreliable or dishonest sources. If anything, unfriendly circumstances call for epistemic caution. In this sense, they require rock-solid precautionary abilities that we often lack. It is then when we fall prey to the perils of negative epistemic dependence. But robust virtue epistemologists should not worry about this. What this paper has shown is that their main tenet—that knowledge is a cognitive success because of cognitive ability—is not jeopardized by any of these two forms of epistemic dependence.

Acknowledgements

I am grateful to an anonymous reviewer for *Synthese* for helpful suggestions.

References

- Breyer, D. & Greco, J. 2008. Cognitive Integration and the Ownership of Belief: Response to Bernecker. *Philosophy and Phenomenological Research* 76: 173–184.
- Broncano-Berrocal, F. (2017). A Robust Enough Virtue Epistemology. *Synthese* 194: 2147–2174.
- Broncano-Berrocal, F. (forthcoming). Purifying Impure Virtue Epistemology. *Philosophical Studies*, DOI: <https://doi.org/10.1007/s11098-017-0873->.
- Goldberg, S. 2011. The Division of Epistemic Labour. *Episteme* 8: 112–25.
- Goldman, A. I. 1979. What is Justified Belief? In Ernest Sosa & Jaegwon Kim (eds.), *Epistemology. An Anthology*. Blackwell: 340–353.
- Greco, J. 2010. *Achieving Knowledge*. Cambridge University Press.

¹⁹ This is particularly true of the kind of cooperative environments described by Goldberg (2011) in which one's community or a third party epistemic agent monitors and polices testimonial exchanges in such a way that one is mostly exposed to reliable speakers (see section 1).

- Greco, J.** 2012. A (Different) Virtue Epistemology. *Philosophy and Phenomenological Research*, 85: 1–26.
- Kallestrup, J., & Pritchard, D. H.** 2012. Robust Virtue Epistemology and Epistemic Anti-Individualism. *Pacific Philosophical Quarterly* 93: 84–103.
- Kallestrup, J., & Pritchard, D. H.** 2013a. Robust Virtue Epistemology and Epistemic Dependence. *Knowledge, Virtue and Action: Putting Epistemic Virtues to Work*, T. Henning & D. Schweikard (eds.), ch. 11, Routledge.
- Kallestrup, J., & Pritchard, D. H.** 2013b. The Power, and Limitations, of Virtue Epistemology. *Powers and Capacities in Philosophy: The New Aristotelianism*, (eds.) J. Greco & R. Groff, ch. 13, Routledge.
- Kallestrup, J., & Pritchard, D. H.** 2014. Virtue Epistemology and Epistemic Twin Earth. *European Journal of Philosophy* 22: 335–357.
- Kallestrup, J., & Pritchard, D. H.** 2016. Dispositional Robust Virtue Epistemology versus Anti-Luck Virtue Epistemology. *Performance Epistemology: Foundations and Applications*, M. Fernández (ed.) ch. 2, Oxford University Press.
- Heil, J.** 2005. Dispositions. *Synthese* 144: 343–356.
- Mackie, J. L.** 1974. *The Cement of the Universe*. Clarendon Press.
- Martin, C. B.** 2008. *The Mind in Nature*. Oxford University Press.
- Michaelian, K.** 2010. In Defence of Gullibility: The Epistemology of Testimony and the Psychology of Deception Detection. *Synthese* 176: 399–427.
- Michaelian, K.** 2012. Metacognition and Endorsement. *Mind & Language* 27: 284–307.
- Michaelian, K.** 2016. *Mental Time Travel: Episodic Memory and Our Knowledge of the Personal Past*. MIT Press.
- Mumford, S. & Anjum, R.** 2011. *Getting Causes from Powers*. Oxford University Press.
- Nelson, T. O., & Narens, L.** 1990. Metamemory: A Theoretical Framework and New Findings. *The Psychology of Learning and Motivation: Advances in Research and Theory*, vol. 26, G. Bower (ed.), Academic Press: 125–173.
- Pritchard, D. H.** 2015. Epistemic Dependence. *Philosophical Perspectives* 29: 305–324.
- Putnam, H.** 1973. Meaning and Reference. *Journal of Philosophy* 70: 699–711.
- Sosa, E.** 2007. *A Virtue Epistemology: Apt Belief and Reflective Knowledge*. Oxford University Press.
- Sosa, E.** 2015. *Judgment and Agency*. Oxford University Press.