



Kelp, C. (2013) How to motivate anti-luck virtue epistemology. *Grazer Philosophische Studien: International Journal for Analytic Philosophy*, 88, pp. 211-225.

There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

<http://eprints.gla.ac.uk/140950/>

Deposited on: 15 May 2017

Enlighten – Research publications by members of the University of Glasgow
<http://eprints.gla.ac.uk>

How To Motivate Anti-Luck Virtue Epistemology

Christoph Kelp*

Abstract

Duncan Pritchard has recently defended an account of knowledge that combines a safety condition with an ability condition on knowledge. In order to explain this bipartite structure of knowledge he appeals to Edward Craig's work on the concept of knowledge. This paper argues that Pritchard's envisaged explanation fails and offers a better alternative.

1 Introduction

In a number of recent pieces Duncan Pritchard [e.g. 2010, 2012] has defended a view he calls anti-luck virtue epistemology (ALVE). According to ALVE,

S knows that *p* if and only if *S*'s safe true belief that *p* is the product of her relevant cognitive abilities (such that her safe cognitive success is to a significant degree creditable to her cognitive agency).

[Pritchard 2012: 273]

While there are a number of subtleties in Pritchard's version of ALVE, what's most important for present purposes is that the view countenances two independent epistemic conditions on knowledge, to wit, a safety-from-error condition and a truth-from-ability condition. In other words, according to ALVE, knowledge has bipartite epistemic structure.

*Centre for Logic and Analytic Philosophy, KU Leuven, Kardinaal Mercierplein 2, BE-3000 Leuven, Email: christoph.kelp@hiw.kuleuven.be

The question that sets the agenda for this paper (henceforth also “the crucial question”) is why knowledge has the bipartite structure ALVE claims it has.¹

In the following passage Pritchard also asks this question and immediately goes on divide it into two subquestions:

One question that we might ask about anti-luck virtue epistemology is why knowledge has this two-part structure in the first place. Put another way, why do we have these two master intuitions about knowledge [subquestion 1], and have them in such a fashion that they make distinct demands on our theory of knowledge [subquestion 2]?

[Pritchard 2012: 274]

In order to see why Pritchard thinks that the crucial question can be broken down into these two subquestions, it is important to realise that Pritchard takes ALVE ultimately to be motivated by what he refers to as “two master intuitions about knowledge”. The first “anti-luck” intuition has it that knowledge is incompatible with luck in the sense that if one knows a proposition, p , then it cannot be the case that one’s underlying belief that p is merely luckily true. In contrast, according to the second “ability” intuition, knowledge requires that one hit the mark of truth thanks to intellectual ability. In other words, if one knows a proposition, p , then that one believes the truth about p must result from the exercise of an intellectual ability [2012: 247-8]. Given that this is how ALVE’s bipartite structure is to be motivated, it is not hard to see that the question why knowledge should have this

¹ One might worry that, ever since Gettier’s famous 1963 paper, there is excellent reason for thinking that knowledge requires a further epistemic condition besides justification. By the same token, there is excellent reason to think knowledge has bipartite epistemic structure and, as a result, Pritchard’s attempt at answering the crucial question is simply superfluous. I would like to make two points by way of response: First, the bipartite structure of knowledge ALVE postulates does not map onto the distinction between the justification condition on the one hand and the condition that deals with Gettier cases on the other. Rather, Pritchard thinks that ALVE’s two conditions both take on anti-Gettier duties: while some Gettier cases are handled by ALVE’s ability condition, others are handled by the safety condition (as becomes clear in Pritchard’s discussion of the Temp and Barney cases [2012: 260-1, 267-9]). Thus, even if Gettier cases serve to motivate a kind of bipartitism, it is not the one that Pritchard needs. Second, Pritchard’s concern with this question appears to be motivated at least in part by a worry, raised for instance by John Greco [2010b: 227], that an account of knowledge that is motivated only by its ability to handle various cases is *ad hoc* in a problematic way. It is not hard to see that in order to address this worry Pritchard needs a motivation that is entirely different than the one Gettier’s argument could supply.

structure should boil down to the questions of why we have these two master intuitions and why we have them in such a way that they make distinct demands on our theory of knowledge.

It may be worth noting just why it matters that Pritchard can offer a good answer to the crucial question and especially to subquestion 2. First, success from ability very plausibly contrasts with lucky success. By way of illustration, consider a competent archer, *A*, who fires a shot at a target. Suppose that, unbeknownst to *A*, (i) strong winds are blowing around the target and (ii) *A* is fortunate enough to have a helper with a wind machine who ensures that his shot will hit the target anyway. In this case, *A*'s hitting the target is lucky for *A* rather than the result of his ability. Lucky success contrasts with success from ability here. Given, additionally, that this point generalises and that, in the case of knowledge, the relevant kind of success is true belief, we get the result that believing the truth from ability contrasts with believing the truth as a matter of luck. It may now appear that the anti-luck intuition can be derived from the ability intuition and that, in consequence, the two intuitions do not impose independent demands on knowledge.² Even if we are willing to grant that we have the two master intuitions, their mere existence is thus not sufficient to motivate ALVE in the way envisaged by Pritchard. In addition, Pritchard needs some account of why we have the two intuitions in such a way that they make distinct demands on our theory of knowledge. In other words, he needs an answer to the crucial question.

This paper discusses how, if at all, the crucial question can be answered. More specifically, I will outline Pritchard's own proposal (section 2), show that it remains unsatisfactory (section 3), and provide an alternative that will do the job for Pritchard (section 4). As I see it, then, the ultimate aim of this paper is to lend Pritchard a helping hand.

² Pritchard [2012: 248-9]. This is a line often pursued by virtue epistemologists such as Greco [2010a, 2012], Wayne Riggs [2002, 2009] and Ernest Sosa [2007, 2011]. It may also be worth noting that Pritchard does not mean to deny that the ability intuition entails the absence of a certain type of luck (sometimes referred to by Pritchard as 'intervening luck' [e.g. 2008]). On the contrary, Pritchard embraces this result (see also §2.2 below). However, he contends that, in addition, knowledge entails the absence of yet another kind of luck ('environmental luck'). It is this requirement that is captured only by the anti-luck intuition.

2 Pritchard's proposal

In order to answer the crucial question, Pritchard draws on Edward Craig's [1990] work on the function of the concept of knowledge.³ (I will henceforth also use '...' in smallcaps as a shorthand for 'the concept of ...'. Thus, for instance, 'KNOWLEDGE' means 'the concept of knowledge'.) In order to be able to assess Pritchard's answer, it will be necessary to take a brief look at Craig's relevant work.

2.1 Craig's genealogical account of the concept of knowledge

Craig's core thesis is that the function of KNOWLEDGE is to flag good informants. In support of this thesis, he first imagines a linguistic community, *L*, which resembles our present linguistic community in that its members need true beliefs about its environment and can get them via testimony from other members. At the same time, *L* also differs from us in that it does not possess KNOWLEDGE. Craig observes that *L* would need a concept to flag good informants and goes on to offer a detailed account of what this concept would have looked like (Stage 1, I will henceforth also refer to the concept at this stage as 'the concept of protoknowledge') and how further needs could have caused it to evolve into the concept of knowledge we are familiar with today (Stage 2).

Concerning Stage 1, Craig argues that we can get at the application conditions of PROTOKNOWLEDGE by considering what general properties an inquirer of *L*, who does not yet have a belief on the question at hand, would want a prospective informant to have. By way of illustration, suppose Elmo and Zoe are members of *L* and Elmo does not yet have a belief on whether the paintings they are looking at are Gonzos but wants to get one from Zoe. We can get at the application conditions of PROTOKNOWLEDGE whether the paint-

³ I'd like to flag a couple of issues I will not discuss within the scope of this paper, *viz.* (a) questions concerning Craig's methodology and (b) whether the crucial question is best answered by appeal to a thesis about the function of the concept of knowledge. For the purposes of this paper, this should not be a cause for concern. Regarding (a), notice that Pritchard also helps himself to Craig's work without further ado. Since the aim of the paper is to lend Pritchard a helping hand, I can leave it to Pritchard to remedy any defect that should arise from so doing. Regarding (b), notice that in order to get his sought after motivation for ALVE, what matters to Pritchard is that there be *some* answer to the crucial question. As a result, the question whether Pritchard has identified the *best* answer is not immediately pressing for Pritchard. Should it turn out that there is a better, entirely different way of achieving this, Pritchard can happily accept it.

ings are Gonzos by asking ourselves what general properties Elmo would want Zoe to have. More specifically Craig suggests that the inquirer (Elmo) would want the informant (Zoe) to tell her the truth on the question (i); he would want her to be as likely to tell the truth on question as his concerns require (ii); he would want her to be detectable as sufficiently likely to tell the truth (iii); he would want the channels of communication between them to be open (iv); and he would want her to be accessible here and now (v) [1990: 85]. According to Craig, then, *PROTOKNOWLEDGE* features a truth condition (i), a reliability condition that is relative to inquirer concerns (ii), and a number of indicator conditions (iii – v) that are also inquirer-relative in various ways.

Concerning Stage 2, Craig argues that, where *L* has attained a sufficient degree of cognitive sophistication, it will be extremely useful for members of *L* to have a concept that enables its members to *recommend* informants on certain questions to others, even when the recommender does not himself know the answer to the question or the needs of the recommendee. Suppose Elmo, Zoe and Oscar are members of *L*. Suppose Elmo does not have a belief on whether the paintings exhibited are Gonzos and has asked Oscar. Oscar himself does not know the answer to Elmo's question and does not know why Elmo asked, but he wants to recommend Zoe as an informant. In this situation, it may be useful for *L* to have a concept of good informant that Oscar can properly apply to Zoe. As Craig argues, since *PROTOKNOWLEDGE* is inquirer-relative in a number of ways, it is unsuitable for this purpose. In our toy case, Zoe may not satisfy the indicator conditions relative to Elmo. Moreover, since it is not known to Oscar (the recommender) what exactly Elmo's concerns require, it is not clear to Oscar that Zoe meets the reliability condition relative to these concerns. What is needed instead is an objective, inquirer-neutral concept of good informant. In response to this need, *PROTOKNOWLEDGE* evolves into exactly such an objective, inquirer-neutral concept (a process Craig's refers to as 'objectivisation').

Exactly how does *PROTOKNOWLEDGE* evolve here? Craig argues that objectivisation will relax the indicator conditions. This is to allow the concept to be properly applicable to Zoe, even when she does not meet these conditions relative to Elmo. At the same time, objectivisation will tighten the reliability condition. Given that the concept is to be applicable even when the concerns of the inquirer are not known, claims Craig, one had better prepare for the worst. Count-

ing only informants that are reliable enough to suit anyone's needs ensures this. Oscar will be able to recommend Zoe as an informant to Elmo without further qualms, even when he doesn't know how demanding Elmo's concerns are. The concept objectivisation outputs (relaxed indicator conditions, stringent, interest-invariable reliability condition) resembles our present concept of knowledge sufficiently closely to confirm Craig's initial hypothesis [1990: 86-91].

2.2 Pritchard's answer to the crucial question

Let's move on to Pritchard's answer to the crucial question. Pritchard invites us to imagine that we have completed Stage 1 of Craig's project and are now at Stage 2. He states his first central idea in the following passage:

[T]here is an important ambiguity in the very notion of a reliable (and hence good) informant. In one sense, it means an informant who possesses a reliable cognitive ability with regard to the target subject matter (and who is willing to sincerely communicate what she believes, something that we will take for granted in what follows). In another sense, it means an informant whom one can rely on (that is, whose information will not lead you astray).⁴

[Pritchard 2012: 276]

More specifically, Pritchard's thought is that as PROTOKNOWLEDGE evolves into KNOWLEDGE in Craig's story, the distinction between informants with reliable cognitive abilities on the one hand and informants that we can rely on will open up [2012: 278]. It will not come as a surprise that Pritchard's ultimate aim is to use this distinction to explain the existence and distinctness of the two master intuitions

⁴ There is an obvious problem with the thought that there could be an ambiguity in a *notion*. Ambiguity is the property of having more than one meaning. Presumably, however, notions, like concepts, are individuated by meanings in such a way that there cannot be ambiguous notions. *Words* (or phrases) can have more than one meaning (*viz.* when they express more than one concept or notion) and can thus be ambiguous. One way in which Pritchard may venture to circumvent this problem is by distinguishing between the word 'knowledge' and the notion (or concept) of knowledge. He may then assume that when the concept with the function of flagging good informants emerged, 'knowledge' was introduced to express this concept. The concept expressed by 'knowledge' at this time is PROTOKNOWLEDGE, which subsequently evolves into KNOWLEDGE so that 'knowledge' comes to express KNOWLEDGE. That said, for simplicity's sake, I will allow myself to go along with Pritchard's talk of ambiguity and multiple senses of notions or concepts.

about knowledge. In order to get there, Pritchard goes on to argue, first, that the two senses of the notion are independent of one another and, second, that the demands they impose correspond to the demands of the two master intuitions.

In order to show that the two senses of the notion of reliable informant that give rise to them are independent of one another, Pritchard adduces two cases:

Misleading Evidence. *S* is a reliable barn spotter but his audience or hearer, *H*, has misleading evidence that *S* is in Fake Barn County (an environment in which nearly all the structures that look like barns are mere façades designed to fool unwitting agents like *S*) [2012: 277].⁵

Here, *S* is an informant who possesses a reliable cognitive ability but not an informant *H* can rely on. The two senses of the notion of a reliable informant can thus be dissociated.

Helpful Information. *S* is a normally unreliable predictor of the outcomes of horse races, say, but *H* has additional evidence that the upcoming race's outcome will be fixed in accordance with *S*'s prediction [2012: 277].

Here, *S* is an informant *H* can rely on but not an informant with a reliable cognitive ability. The two senses of the notion of a reliable informant can thus be doubly dissociated.

Since the two senses of reliable informant are independent of one another, they are fit to impose distinct demands on PROTOKNOWLEDGE at this stage of its evolution into KNOWLEDGE. According to Pritchard, they will do so, since there is every reason to expect that once the distinction has opened up and it is clear that its two parts are suitably independent, PROTOKNOWLEDGE will henceforth respect both sides of the distinction: neither agents we cannot rely upon nor agents who have failed to fail to manifest a reliable cognitive ability will be counted as knowers [2012: 278].

⁵ Recall that when Pritchard first introduces the notion of an informant that one can rely on, he characterises an informant that one can rely on as an informant whose information will not lead one astray. However, given how *Misleading Evidence* is intended to work, this characterisation cannot be what Pritchard really has in mind. After all, *S* actually is a reliable barn spotter in a suitably friendly environment (not Fake Barn County, etc.). As a result it is not as if the information *S* provides will lead *H* astray. Rather, the problem is that, in view of the additional information *H* possesses, *H* cannot trust *S*'s word. Accordingly, I will assume that what Pritchard really means by 'informant that one can rely on' is 'informant whose word one can trust'.

Even if we are prepared to follow Pritchard this far, we still have to see that the two demands correspond to the demands of the two master intuitions. In order to see how Pritchard argues that they do, it will be necessary to take a closer look at the role Pritchard takes luck to play in the cases that establish the independence of the two senses of the notion of a reliable informant. In cases in which an agent has a reliable cognitive ability but cannot be relied on (such as in *Misleading Evidence*), good epistemic luck (*S* has the relevant ability⁶) is cancelled out by bad epistemic luck (*H* has additional misleading evidence) with the result that *H* cannot rely on *S*. At the same time, in cases in which the agent can be relied on despite not having a reliable cognitive ability (such as *Helpful Information*), bad epistemic luck (*S* lacks the ability) is cancelled out by good epistemic luck (*H* has additional helpful information) [2012: 277-8]. The thought here seems to be that what explains why in these cases the two senses of reliable informant come apart is that there are (different kinds of) luck operative, which correspond to the kinds of luck at issue in the two master intuitions [2012: 278].⁷

⁶ One might ask just why we should think that *S*'s having the relevant ability counts as an instance of good epistemic luck. Unfortunately, Pritchard does not offer much by way of an answer to this question. Part of the idea here appears to be that it is lucky for *H* that *S* has the ability. However, on reflection, it is not really clear why even this should be so (or why it should be relevant to whether *S* knows, see §3.2). That prospective informants have mundane cognitive abilities, such as the ability to spot barns, does not seem to be lucky for relevant inquirers. While this aspect of Pritchard's explanation thus remains somewhat puzzling, I will go along with it for the time being.

⁷ For the record, here is how Pritchard himself states his answer to the crucial question, after having introduced the distinction between the two senses of the notion of a reliable informant, having argued their independence by means of *Misleading Evidence* and *Helpful Information*, and having made the point about the role of luck in these cases:

With this ... in mind, it ought to be clear why this ambiguity in the idea of a reliable informant explains why the concept of knowledge that evolves from the proto-concept will generate *both* the anti-luck and the ability intuition. For as the range of cases which the concept of knowledge is meant to apply to widens [i.e. as PROTOKNOWLEDGE evolves into KNOWLEDGE], so the distinction will open up between good informants who are reliable and good informants that we can rely on, and we would expect the concept of knowledge that results to respect both sides of this distinction. In particular, examples where an agent possesses the relevant reliable cognitive abilities but where the presence of epistemic luck means that we would not be able to rely on this agent *qua* informant would not be counted as cases of knowledge. Similarly, those cases in which an agent forms a true

In a nutshell, then, Pritchard’s answer to the crucial question is that, as PROTOKNOWLEDGE evolves into KNOWLEDGE, (i) a distinction between two senses of the notion of a reliable informant opens up, (ii) the two senses impose different demands on the concept of protoknowledge (and hence on the concept of knowledge into which it evolves), and (iii) that these demands correspond to the demands of the two master intuitions about knowledge.

3 Problems with Pritchard’s explanation

3.1 Problem 1: Knowers who cannot be relied on

Let’s now ask how convincing Pritchard’s answer really is. Notice that Pritchard’s story relies on the claim that the concept of knowledge features a condition that rules against agents we cannot rely on. (Recall that, according to Pritchard, “examples where an agent possesses the relevant reliable cognitive abilities but where the presence of epistemic luck means that we would not be able to rely on this agent *qua* informant would not be counted as cases of knowledge” [2012: 278].) One problem for Pritchard is that the concept of knowledge simply does not feature such a condition. Perhaps somewhat surprisingly, Craig himself provides the relevant evidence for this:

[S]omeone who both knows the truth and is keen to reveal it may be useless to others because he has no credibility with them: the boy who cried ‘Wolf!’ so often that no-one would believe him when the wolf really came is a cautionary example, as is Matilda of Hillaire Belloc’s *Cautionary Tales*.

[Craig 1990: 17]

belief in an epistemically friendly environment—such that any true belief so formed would not be subject to epistemic luck—would not be counted as cases of knowledge so long as the agent concerned failed to exhibit the relevant reliable cognitive abilities (even though we could rely on this agent *qua* informant). In short, the concept of knowledge that results will both (i) disallow cases of true belief as knowledge where the belief isn’t appropriately due to the relevant cognitive abilities on the part of the agent, and (ii) disallow cases of true belief as knowledge where the truth of the belief is substantively due to luck and hence unsafe.

[Pritchard 2012: 278]

What Craig's cases illustrate is that an agent can know a proposition even though others cannot rely on that agent qua informant. The boy knows that the wolf came and Belloc's Matilda knows that the house is on fire even though one not could rely on them because they have a track record of making the corresponding false claims that undermines their credibility.

One might be inclined to object on behalf of Pritchard that in Craig's cases it is not the "presence of epistemic luck" that prevents us from being able to rely on the relevant agent qua informant. Recall that, in Pritchard's *Misleading Evidence*, *S* has a reliable ability to detect barns but cannot be relied on by *H* because *H* has misleading evidence that *S* is in Fake Barn County. Here, *H*'s misleading evidence constitutes an element of bad epistemic luck that prevents her from being able to rely on *S*. However, the same is not true in Craig's cases. There is no misleading evidence here. Hence, it might be thought that, even though in Craig's cases we cannot rely on the agent qua informant, by Pritchard's lights, this does not disqualify him from being counted as knowing.

To see why this move fails, notice that we can easily amend the cases in such a way that epistemic luck is involved in the same way as in Pritchard's own case: we just construe the case in such a way that the hearers falsely believe on the basis of misleading evidence that the agent has relevant track record. Even so, if the boy believes that the wolf came because he has seen that he did and if Matilda believes that the house is on fire because she has seen that it is, they count as knowing what they have seen. Notice, furthermore, that the same holds true in Pritchard's *Misleading Evidence* in which *H* has misleading evidence that *S* is in Fake Barn County: If *S* believes that he is facing a barn, say, because he has seen that he is facing a barn, he counts as knowing that he is facing a barn, no matter whether *H* can rely on his word that he is facing a barn.

There is thus excellent reason to believe, *pace* Pritchard, that the concept of knowledge does not feature a condition ruling against agents we cannot rely on. As a result, Pritchard's answer to the crucial question remains unsatisfactory. There is after all little reason to think that the sense of the notion of reliable informant according to which a reliable informant is one that can be relied on had much of a role to play in shaping the structure of the concept of knowledge.

3.2 Problem 2: Hearer vs. speaker luck

Pritchard might venture to respond that while the concept of knowledge we are familiar with today does not feature a condition ruling against agents we cannot rely on, an ancestor of it did. At that point, the anti-luck intuition came into existence. While the condition ruling against agents we cannot rely on was subsequently filtered out by objectivisation, the anti-luck intuition remained.

As a first observation, notice that on the present story the condition ruling against agents we cannot rely on makes a brief appearance, leaves Pritchard exactly what he needs for his project, only to then conveniently withdraw from the scene again. I would not be surprised if there were many ears to which this sounds just too good to be true.

Unfortunately, even this story won't do the trick for Pritchard. The reason for this is that there is a further problem concerning his step from the condition ruling against agents we cannot rely on qua informants to the thesis that this condition would give rise to the anti-luck intuition. To begin with, Pritchard's argument leaves a lot implicit here. Pritchard offers a diagnosis of cases in which the agent has a reliable cognitive ability but cannot be relied on in terms of luck. For instance, in *Misleading Evidence*, what explains why *H* cannot rely on *S* is that *H* suffers from bad epistemic luck—he has misleading evidence. Presumably, the thought then is (a) that the relevant kind of bad epistemic luck is incompatible with knowing and (b) that this incompatibility generates the anti-luck intuition.

Concerning (a), we have already seen that there is excellent reason to think that the relevant kind of epistemic luck is not incompatible with knowing. In Craig's cases as well as in Pritchard's own *Misleading Evidence*, the agents may know the relevant propositions even though they cannot be relied on qua informants—for instance, if they believe the propositions because they have seen that they are true. In fact, a closer look at Pritchard's diagnosis provides further support for this point. What a careful look at the diagnosis makes clear is that the person who is affected by luck here may be the hearer only. To see this, consider once more Pritchard's own case, *Misleading Evidence*. Here it is only *H* who is suffering from epistemic bad luck. *H* has the misfortune of having misleading evidence about *S* and as a result cannot rely on *S*. As opposed to *H*, *S* is not suffering from epistemic bad luck at all: *S* exercised a reliable cognitive ability in favourable circumstances. It is hard to see, however, why *H*'s suf-

fering from bad epistemic luck should undermine *S*'s knowing and hence why the relevant kind of bad luck at issue here should be incompatible with *S*'s knowing.

Furthermore, contrary to (b), there is little reason to think that a condition that disallows the kind of bad luck at issue here would generate the anti-luck intuition. After all, as we have just seen, it may well be that the only one affected by bad luck in the relevant cases is the hearer. It is hard to see, however, why a condition that rules against this kind of bad luck should generate an intuition according to which lucky truth of the *speaker's* belief is incompatible with knowledge. Thus, even if we accepted that the condition ruling against agents we cannot rely on entered and existed the story in just the way Pritchard would need it to, there is still little hope that his attempt at generating the anti-luck intuition would succeed.

3.3 Problem 3: The wrong hypothesis

A final worry about Pritchard argument arises from a challenge to the very hypothesis he adopts from Craig. As I have argued in recent work [Kelp 2011], Craig's hypothesis is incorrect as there is a better alternative in the offing. According to this alternative, the function of the concept of knowledge consists in marking when an agent is entitled to inquire no further into a given question. In order to bring this point home, I offer a direct comparison of how the two hypotheses fare given Craig's own methodology. While I grant that both hypotheses generate a concept that is very similar to the concept of knowledge we are familiar with, I show that the alternative hypothesis produces a better match. There are a couple of types of case in which only the concept generated by the alternative hypothesis passes the intuitively correct verdicts. In the first type of case, agents have beliefs they won't assert because they are bound by professional secrecy. I argue that at least some such agents (such as priests under the seal of confession) do not fall under the concept generated by Craig's hypothesis—they are not good informants even in the objective sense—while, intuitively, they know. In the other type of case, the agents' grounds for asserting differ from their grounds for believing. Crucially, the agent's grounds for asserting are epistemically unexceptionable, while the grounds for believing are epistemically highly problematic. Here I show that the agents fall under the concept generated by Craig's hypothesis—they are good informants even in the objective sense—while, intuitively, their ill-based beliefs don't

qualify as knowledge. As opposed to that, the concept generated by the alternative hypothesis—the agents have an entitlement to inquire no further—passes the right verdicts in both cases: priests under the seal of confession have an entitlement to inquire no further into the relevant questions concerning the sins of penitents, while agents who base their beliefs on epistemically problematic grounds have no such entitlement. Given that the alternative hypothesis produces a better match with our contemporary concept of knowledge, Craig-style epistemology confirms the alternative hypothesis whilst disconfirming Craig's own contender.

The fact that the alternative hypothesis is preferable to Craig's need not pose a problem for Pritchard. It won't if Pritchard can recover his argument given the alternative hypothesis. Unfortunately for Pritchard, the prospects for such recovery are rather dim. After all, Pritchard's crucial distinction between agents with a reliable cognitive ability and agents we can rely on qua informants must not be expected to have any relevance to a concept with the function of marking when an agent is entitled to inquire no further into a given question. In particular, I take it that no one would be tempted by the idea that there is a sense of being entitled to inquire no further into a given question according to which an agent possesses an entitlement in this sense only if he can be relied on qua informant. That is to say, once the problematic Craigian hypothesis is replaced by the alternative, already Pritchard's first move loses any plausibility it may have initially appeared to have had.

4 An alternative proposal

There is thus bad news for Pritchard: his answer to the crucial question remains unsatisfactory. Fortunately, there is also good news. As I am about to argue, a better answer can after all be given. This answer improves on Pritchard's own proposal in a number of respects. First, it does not rely on there being a sense of reliable informant according to which a reliable informant is one whom we can rely on qua informant. Second, it does not presuppose an arguably incorrect hypothesis about the function of the concept of knowledge but works on the more plausible alternative. Third, unlike Pritchard's answer, it avoids the detour via the two master intuitions. Instead, the explanation of why knowledge features both a safety and an ability condition

is direct.⁸

So how can the alternative hypothesis be used to answer the crucial question? I will begin with the following observation about the range of cases in which we would want a concept with this function to apply: we sometimes want to say that an agent is entitled to inquire no further into some question when we ourselves do not know the answer to it. This may happen in cases like the following: Ernie is interested in the question whether the surface Bert is currently looking at is red but does not himself know the answer. In this situation Ernie may want point out that Bert is entitled to inquire no further into the matter. By the present hypothesis, KNOWLEDGE will allow him to achieve this. And so it does: Ernie can attribute to Bert knowledge whether the surface is red.

In a similar way as Craig did, we can now ask ourselves what sorts of conditions would govern a concept that serves this function and is applicable when we ourselves don't know the truth about the question. For one, we would expect it to respect a modal condition: Were we to find out that the attributee might so easily have been mistaken, we could not attribute to him an entitlement to inquire no further into the issue. Were Ernie to be informed that nearly all of the surfaces in Bert's environment that appear red are in fact non-red, Ernie would have to withdraw his attribution of entitlement to pursue the issue no further. What's more, if Ernie knew that Bert's answer to the question of the colour of the surface is based only on Bert's taking appearances at face value, Ernie would have to deny that Bert is entitled to pursue the question no further. Ernie would have to say that Bert doesn't know whether the surface is red.

At the same time, we would also expect a concept with the envisaged function to respect an ability condition: Were we to find out that the attributee does not have the relevant cognitive ability, we could not attribute to him an entitlement to inquire no further. Were Ernie to be informed that Bert is red-green colour blind, Ernie would have to withdraw his attribution of entitlement to pursue the issue no further. What's more, if Ernie knew that Bert's answer to the question of the colour of the surface is based only on exercising his ability to dis-

⁸ To see that this is an advantage, notice that Pritchard will need further assumptions to get from the master intuitions to ALVE. In particular, he will need his modal account of luck [Pritchard 2005: ch.5] (or something very close to it) to get from the anti-luck intuition to ALVE's safety condition. As a result, there is yet another way in which he makes his argument vulnerable to failure: the modal account of luck may prove to be wrong.

cern colours, Ernie would have to deny that Bert is entitled to pursue the question no further. Ernie would have to say that Bert doesn't know whether the surface is red.

It is not hard to see that these two conditions are independent of one another. In the first case Ernie would have to withdraw his attribution of knowledge/deny Bert knowledge even when Bert exercises a highly reliable ability to tell colours apart. In the second case, Ernie would have to withdraw his attribution of knowledge/deny Bert knowledge even when the surface Bert is looking at could not have been any colour other than red.

These considerations suggest that a concept with the function of marking when an agent is entitled to inquire no further into some question will feature both a safety and an ability condition. Given the alternative hypothesis, according to which the concept of knowledge is exactly this concept, we thus have an answer to the crucial question. Furthermore, the present answer can be expected to avoid the problems Pritchard encountered since, unlike Pritchard's alternative, it does not venture to generate the bipartite structure via a distinction between the notion of an informant with a reliable cognitive ability and the notion of an informant we can rely on.⁹ Rather, the bipartite structure is generated by appeal to conditions that would force a withdrawal of an attribution of the concept or replacement by denial. These are exactly the kinds of consideration that would serve to explain why the concept features the corresponding necessary conditions. Thus, unlike Pritchard's account, the present account does offer the right kind of explanation of the bipartite structure of knowledge.

References

Craig, E. 1990. *Knowledge and the State of Nature. An Essay in Conceptual Synthesis*. Clarendon Press, Oxford.

⁹ Correlatively, the proposed account need not appeal to *Helpful Information* and *Misleading Evidence* to achieve its goal. In fact, there is excellent reason to think that it will give the right predictions in both of these cases. In *Helpful Information*, the agent does not earn an entitlement to inquire no further as he does not have the relevant cognitive ability. In contrast, in *Misleading Evidence*, the agent may earn such an entitlement. He believes the truth from ability and is in a suitably hospitable environment, which renders his belief safe. In contrast with Pritchard's account, the fact that his audience has misleading evidence concerning his trustworthiness is, on the present account, entirely orthogonal to whether he may earn this entitlement. All of this is, of course, entirely as it should be.

- Gettier, E. 1963. Is justified true belief knowledge? *Analysis*, 23:121–23.
- Greco, J. 2010a. *Achieving Knowledge*. Cambridge University Press, Cambridge.
- Greco, J. 2010b. The value problem. In Pritchard, D. and Bernecker, S., editors, *The Routledge Companion to Epistemology*. Routledge, London.
- Greco, J. 2012. A (different) virtue epistemology. *Philosophy and Phenomenological Research*, forthcoming.
- Kelp, C. 2011. What's the point of 'knowledge' anyway? *Episteme*, 8:53–66.
- Pritchard, D. 2005. *Epistemic Luck*. Oxford University Press, Oxford.
- Pritchard, D. 2008. Greco on knowledge: Virtues, contexts, achievements. *The Philosophical Quarterly*, 58:437–47.
- Pritchard, D. 2012. Anti-luck virtue epistemology. *The Journal of Philosophy*, 109:247–79.
- Pritchard, D., Millar, A., and Haddock, A. 2010. *The Nature and Value of Knowledge*. Oxford University Press, Oxford.
- Riggs, W. 2002. Reliability and the value of knowledge. *Philosophy and Phenomenological Research*, 64:79–96.
- Riggs, W. 2009. Luck, knowledge, and control. In Haddock, A., Millar, A., and Pritchard, D., editors, *Epistemic Value*. Oxford University Press, Oxford.
- Sosa, E. 2007. *A Virtue Epistemology. Apt Belief and Reflective Knowledge*, volume 1. Oxford University Press, Oxford.
- Sosa, E. 2011. *Knowing Full-Well*. Princeton University Press, Princeton/NJ.