CrossMark

Episteme, 11, 4 (2014) 369–389 © Cambridge University Press
doi:10.1017/epi.2014.24

FAKE BARNS AND FALSE DILEMMAS

CLAYTON LITTLEJOHN
cmlittlejohn@yahoo.com

ABSTRACT

The central thesis of robust virtue epistemology (RVE) is that the difference between knowledge and mere true belief is that knowledge involves success that is attributable to a subject's abilities. An influential objection to this approach is that RVE delivers the wrong verdicts in cases of environmental luck. Critics of RVE argue that the view needs to be supplemented with modal anti-luck condition. This particular criticism rests on a number of mistakes about the nature of ability that I shall try to rectify here.

1. INTRODUCTION

The central thesis of robust virtue epistemology (RVE) is that propositional knowledge is cognitive success that is attributable to a subject's abilities (Zagzebski 1996; Sosa 2007; Greco 2010). For the purposes of this discussion, cognitive success should be understood as the formation of a true belief. The proposal under consideration is that the difference between knowledge and mere true belief has to do with whether the subject's true belief is formed in such a way that this ability condition is satisfied:

> AC: One's cognitive success is properly attributable to one's cognitive abilities (i.e., one's success is because of one's abilities, manifests one's abilities, or is due to one's abilities).[1]

When one's belief is accurate and AC is met, one's belief is supposed to constitute knowledge.

Critics claim that it is possible for a subject's beliefs to satisfy AC even if that subject does not have knowledge.[2] We'll look at an anti-luck argument against RVE, one that's designed to show that there's more to knowledge than RVE suggests. We shall see that

1 There are different accounts of what it takes for accuracy to be attributable to one's abilities. Greco (2010, 2012) thinks that explanatory considerations matter when it comes to determining whether success is attributable to one's abilities. Sosa (2007, 2009) and Turri (2011) think that accuracy is attributable to ability when accuracy *manifests* one's abilities and so prefer a more metaphysically loaded account. This second approach is the approach that I prefer. I fear that the direction of fit problem that Pritchard thinks arises for robust modal accounts of knowledge will arise for virtue-accounts that do not invoke metaphysical notions like the manifestation of a disposition or power. I shall bracket this issue here. Critics of RVE seem to think that on every plausible reading of AC it is possible to satisfy AC without having knowledge.

2 Lackey (2007) raises another influential objection to this account, which is that testimony cases show that it is possible to acquire knowledge even when success should be credited primarily to the abilities of another. I hope to discuss this challenge in a later paper.

https://doi.org/10.1017/epi.2014.24 Published online by Cambridge University Press

this argument rests on a problematic conception of what it takes for success to be properly attributable to one's abilities. If my limited defense of RVE is successful, it shows that subsequent attempts to fix what's wrong with virtue-theoretic approaches to knowledge by adding additional modal conditions to vindicate anti-luck intuitions are misguided.[3] Properly understood, AC is all the anti-luck condition we need.

2. THE ANTI-LUCK ARGUMENT

Critics of robust virtue epistemology insist that however we unpack the idea of accuracy being attributable to ability there's more to knowledge than meeting AC. The argument I shall discuss here is this anti-luck argument against RVE:

> AL1. The presence of veritic luck, whether it is intervening luck or environmental luck, will prevent one's true beliefs from constituting knowledge.[4]
> AL2. The ability condition can be met in cases of environmental luck.
> ALC. Thus, there must be more to propositional knowledge than a true belief whose accuracy is properly attributable to the cognitive abilities responsible for its formation.

This argument is supposed to show that there must be more to knowledge than RVE says there is.[5] Those who find the argument convincing typically say that we need to add some modal condition to serve as an additional anti-luck condition.[6]

---

3 My defense of RVE is limited in two respects. The first is that I'm only going to address one influential objection to RVE, an argument that appeals to intuitions about environmental luck cases. I don't address Miracchi's (forthcoming) argument that RVE is too fixated on abilities to form beliefs that are true, for example, but I think that this represents a serious challenge to the viability of RVE. The second is that I do not intend to defend the claims about epistemic value or the relationship between knowledge and achievement often defended by virtue epistemologists. Nothing here should be taken as a defense of the idea that the value of knowledge is connected to the subject's performance or the idea that knowledge is an achievement. Bradford (2013) argues quite plausibly that the value of an achievement is largely determined by how difficult it is to pull the achievement off. Since so much knowledge is so easily acquired, I don't think the value of knowledge has much to do with achievement. If you're looking for reasons to be skeptical of the suggestion that knowledge is an achievement, Lackey (2007), Whiting (2012), and Hacker (2013) are good places to look.

4 For arguments for the malignancy of epistemic luck, see Engel (1992), Madison (2011), and Pritchard (2005). For a discussion of the difference between intervening and environmental luck, see Jarvis (forthcoming) and Pritchard (2009b). For skepticism about the significance of epistemic luck, see Hetherington (1998, 2012). Most of the authors in this discussion agree that something in the neighborhood of AC is needed in an account of knowledge because they think that modal accounts deliver the wrong verdicts in cases where the safety from error has too much to do with fortuitous features of the circumstances. For a discussion of such cases, see Luper (2006).

5 See Kvanvig (2003, 2010), Pritchard (2009b, 2012a, 2012c), Littlejohn (2011, 2012), Kallestrup and Pritchard (2013), and Whitcomb (Manuscript). In Littlejohn (2012), I thought that this sort of argument showed that knowledge is distinct from apt belief and that, as a consequence, we need apt belief, not knowledge, to conform to the norms governing belief, to possess reasons, and to justifiably believe something. I now see that this was a mistake. Historians take note. See Littlejohn (2013) for further discussion.

6 See Pritchard (2009b, 2012a, 2012c) for this diagnosis. He thinks that the modal condition we need to add is a safety condition. For discussion of the role of safety in an anti-luck epistemology, see Williamson (2000), Pritchard (2005), and Luper (2006). For defenses of a sensitivity condition as a

https://doi.org/10.1017/epi.2014.24 Published online by Cambridge University Press

Let's start with (AL1). The argument's first premise says that certain kinds of epistemic luck are epistemically malignant. Let's consider three cases that illustrate the two kinds of veritic luck:

> Roddy: Using his reliable perceptual faculties, Roddy non-inferentially forms a true belief that there is a sheep in the field before him. His belief is also true. Unbeknownst to Roddy, however, the truth of his belief is completely unconnected to the manner in which he acquired this belief since the object he is looking at in the field is not a sheep at all, but rather a sheep-shaped object which is obscuring from view the real sheep hidden behind. (Pritchard 2012a: 251)
> Barney: Using his reliable perceptual faculties, Barney non-inferentially forms a true belief that the object in front of him is a barn. Barney is indeed looking at a barn. Unbeknownst to Barney, however, he is in an epistemically unfriendly environment when it comes to making observations of this sort, since most objects that look like barns in these parts are in fact barn façades. (Pritchard 2012a: 251)[7]
> Chris' Clock: A demon ... wants our hero – let's call him 'Chris' – to form a belief that the time is 8:22 a.m. when he comes down the stairs first thing in the morning (the demon doesn't care whether the belief is true). Since he is a demon, with lots of special powers, he is able to ensure that Chris believes this proposition (e.g., by manipulating the clock). Now, suppose that Chris does indeed come downstairs that morning at exactly 8:22 a.m., and so forms a belief that the time is 8:22 a.m. by looking at the clock at the bottom of the stairs. Since Chris is going to form this belief anyway, the demon doesn't need to do anything to ensure that he forms the belief in the target proposition. Moreover, since Chris is forming his belief by consulting a reliable clock, one would intuitively regard this as an instance of knowledge ... . Nevertheless, the belief is clearly unsafe, since there are many near-by possible worlds in which Chris continues to form the belief that it is 8:22 a.m., and yet this belief is false (because of the interference of the demon). (Pritchard 2009a: 37)

The first case, Roddy, is a clear case of intervening luck, a case in which something gets 'in between' the exercise of the subject's relevant cognitive abilities and the conditions that determine whether the subject's belief is accurate. The intuition that Roddy's belief does not constitute knowledge is widely shared. The second and third cases, Barney and Chris' Clock, differ in an important respect.[8] There is nothing that gets 'in between' the exercise of the subject's relevant cognitive abilities and the conditions that determine whether the subject's belief is accurate. People stress that the subjects in these cases see a barn or a functioning clock and do not form their belief about the situation on the basis of any inaccurate representations. While the intuition is that subjects in these two cases will not have knowledge, critics of RVE think that RVE does not have the resources to explain why these subjects' beliefs fail to constitute knowledge.

Let's consider (AL2). In this passage, Pritchard explains why success should be attributed to Barney's abilities:

> Barney is ... really seeing a genuine barn .... In a very real sense, then, Barney's cognitive abilities are putting him in touch with the relevant fact, unlike in standard Gettier-style cases, where there is a kind of fissure between ability and fact, albeit one that does not prevent the agent from having

suitable anti-luck condition, see Nozick (1981), DeRose (1995), Black and Murphy (2007), and Roush (2007).

7 It is controversial whether Barney can know that the structure is a barn. See Gendler and Hawthorne (2005) and Sosa (2007). Sosa offers a kind of error theory to explain away the intuition that Barney does not know. As he sees it, Barney can have animal knowledge but not reflective knowledge. For critical discussion, see Battaly (2009). The Barney case first appeared in Goldman (1976).

8 Chris' Clock is from Kelp (2009).

https://doi.org/10.1017/epi.2014.24 Published online by Cambridge University Press

> a true belief regardless. [G]iven that Barney does undertake, using his cognitive abilities, a genuine perception of the barn, it seems that his cognitive success is explained by his cognitive abilities, unlike in standard Gettier-style cases. (2012a: 267)

In this passage, he explains why success should be attributed to Chris' abilities:

> While nothing intervenes between Chris's cognitive ability and his cognitive success – he really does employ his cognitive abilities in order to gain his true belief about the time – he is in a very unfriendly environment from an epistemic point of view. Nevertheless, because the demon doesn't in fact interfere in the actual case, I think we should regard Chris's true belief as a cognitive achievement – his abilities are, after all, the best explanation of why he is successful – even though his belief is only luckily true. (2009a: 40)

It looks as if Pritchard's rationale for (AL2) might be summed up as follows. In these cases, the subjects have retained the abilities that would help them to acquire knowledge in friendlier epistemic environments and there is nothing in the situation that interferes with their exercise. Since Barney and Chris' visual abilities are retained, exercised, and they subsequently relate them to their surroundings in these bad cases (i.e., non-knowledge cases) just as they would be related to their surroundings in corresponding good cases (i.e., knowledge cases), we would have to attribute accuracy to these abilities and their exercise in the same way in both cases. Since both sides assume that accuracy or success is attributable to ability in the good case, accuracy must be attributable to ability in the bad.

Sosa says something quite similar in the course of discussing his kaleidoscope case:

> Kaleidoscope: Katherine sees a surface that looks red in ostensibly normal conditions. But it is a kaleidoscope surface controlled by a jokester who also controls the ambient light, and might as easily have presented her with a red-light + white-surface combination as with the actual white-light + red-surface combination. Does she know that the surface she sees to be red is indeed red when presented with the good combination, despite the fact that, even more easily, he might have presented her with the bad combination? (2007: 31)[9]

Sosa thinks that this is a case in which success should be attributed to Katherine's relevant cognitive abilities. While factors could easily have come into play that would have interfered, they did not actually come into play and so nothing interferes with Katherine's abilities, their exercise, or the way that their exercise relates to cognitive success (2010: 76).

The rationale for (AL1) is, in effect, an appeal to widely shared intuitions. The rationale for (AL2) is a set of considerations that are supposed to help us see why success should be attributed to the subjects' cognitive abilities. There are three responses to the argument in the literature. Some find the argument compelling and conclude that there must be more to knowledge than RVE would have us believe. Pritchard thinks that we should retain AC to deal with certain kinds of trouble cases where we feel that a robust modal connection between belief and fact isn't sufficient for 'turning' a true belief into knowledge. What we shouldn't do, he thinks, is use AC to do the work of an anti-luck condition. What these environmental luck cases show is that we need an ability condition like AC and a modal condition to deal with the problem of epistemic luck. Those who defend RVE disagree

9 Modified slightly so as to introduce a character.

https://doi.org/10.1017/epi.2014.24 Published online by Cambridge University Press

about whether we should deny (AL1) or deny (AL2). Sosa, as we've seen, agrees with Pritchard that AC is met in environmental luck cases. He thinks that we should contest the intuitions about environmental luck cases that Pritchard appeals to. Greco agrees with Pritchard that environmental luck is epistemically malignant, but disagrees with Pritchard and Sosa about (AL2). As he sees it, environmental luck isn't just epistemically malignant, it shows that cognitive success cannot be attributed to ability.

Because there is a disagreement here about (AL2), I shall frame this discussion as a debate between two views, incompatibilism and compatibilism. If one thinks that the accuracy of a subject's belief in an environmental luck case is *not* attributable to the subject's relevant cognitive abilities, one is an *incompatibilist*. If, however, one thinks that the accuracy might be attributable to the subject's abilities in these cases, one is a *compatibilist*. The debate between the compatibilist and incompatibilist arises because of a disagreement about how to understand AC. We've seen why Pritchard and Sosa are compatibilists. As they see it, the features that explain why environmental luck is epistemically malignant have no bearing on what abilities one has, how they're exercised, or how they relate our subjects to their surroundings. As I see it, their rationale for compatibilism is not compelling. I shall side with Greco and offer a defense of incompatibilism.

3. ABILITIES

While the compatibilist rationale for (AL2) might strike the reader as *prima facie* plausible, it seems to me to suffer from two related defects. The first has to do with the relationship between ability and opportunity. The second has to do with the compatibilists' failure to attend to important differences between epistemic and non-epistemic abilities.

It is important to distinguish the *abilities* that are, as Sosa puts it, resident in the subject from *opportunities*, external circumstances under which these abilities might be exercised. One can remove the one without removing the other (e.g., one can drain the lake and deprive others of the opportunity to swim without taking away their ability to swim or turn off the lights to prevent someone from seeing the mess in the living room). An opportunity is not simply a situation in which an ability or a capacity might be exercised, but a situation in which an ability or capacity might be exercised in such a way that the exercise of the ability might result in success. We can deprive an interviewee of the opportunity to impress a panel by filling the room with water and preventing her from speaking. We could also deprive her of the opportunity by drugging the panel and preventing them from comprehending. Opportunity requires the absence of internal and external impediments. If some successful result is the manifestation of the subject's abilities, the subject must have had the ability, exercised it, and been given the right kind of opportunity.

Kenny (1992: 68) reminds us that there are different senses of 'can' and 'able' and cautions us against conflating them. There is the 'can' and 'able' of general ability (e.g., 'Can he read Spanish?'), a 'can' and 'able' of opportunity (e.g., 'Can the condemned have a cigarette after their last meal?'), and a 'can' and 'able' that has an overall sense, one that indicates both general ability and opportunity.[10] We can ask whether someone is able to

10 See also Maier (forthcoming) who understands S's general ability to A in terms of whether S has the option to A in a suitable range of cases.

https://doi.org/10.1017/epi.2014.24 Published online by Cambridge University Press

whistle or crack a safe and ask about general ability, opportunity, or overall ability. Since there are these different readings, we have to decide how we should understand the ability condition that's central to the robust virtue theory's approach to knowledge.

Once we see that there is a distinction to be drawn between general ability, which has to do almost exclusively with what is resident in the subject, and overall ability I would have thought that AC should be understood as having to do with overall ability, not general ability. Even if proponents of RVE have not been sufficiently clear on the matter, there is no question that there's a view in the spirit of RVE that says that overall ability, not general ability, is what we should focus on when trying to understand the difference between knowledge and mere true belief. The difference matters because there can be cases of true belief that results from the exercise of some ability that has the potential of producing knowledge when the right kind of opportunity for exercising that ability has been removed. One can remove the appropriate kind of opportunity for exercising an ability without thereby removing the ability or preventing a subject from exercising it. To modify one of Pritchard's examples, one might be an accomplished pianist and try to play a piano underwater. Being underwater does not necessarily prevent one from exercising the abilities one exercises when playing under normal conditions. It certainly does not cause one to lose those abilities. Suppose that while playing underwater the piano makes no audible noise but sets off some strange chain of events that causes a piano on land to play just the notes that one would play if one were on land. This success is not attributable to ability even though one has exercised the very abilities that we would attribute success to under normal conditions. Similarly, consider the ability that one might have to make others laugh. Suppose one has started to address people attending a funeral in the mistaken belief that the somber audience is there to hear a research talk and one starts off by telling a few very blue jokes. One might get the crowd to laugh as a result, but this would be a kind of nervous laughter that's a response to the highly inappropriate things that one is saying. This is not a case where the laughter is attributable to one's abilities. To test whether success is indeed attributable to the abilities resident in a subject, we have to do so under circumstances suitable for the exercise of those abilities.

If we read RVE as saying that knowledge is cognitive success that's properly attributable to cognitive abilities *when exercised under appropriate circumstances*, we need to ask whether the subjects in environmental luck cases have been afforded the right kind of circumstances. I shall argue that they have not and argue against (AL2) as follows:

IA1: Cognitive success is properly attributable to ability only when the subject has been afforded the right kind of opportunity for exercising a cognitive ability and the subject's exercise of the ability results in cognitive success.

IA2: In environmental luck cases, the subject has not been afforded the right kind of opportunity for exercising the relevant cognitive ability.[11]

C: The ability condition cannot be met in cases of environmental luck.

11 I disagree with Millar (2009) who thinks that the abilities involved in acquiring knowledge are not exercised when we form beliefs without being in a position to know. For the most part, I think that abilities are, as Sosa says, resident in the subject and whether one exercises a certain ability will not typically depend upon whether the situation is appropriate for exercising the relevant ability. I also disagree with Greco (2010) if, as Pritchard (2012a, 2012c) reads him, his view is that abilities themselves are indexed to appropriate circumstances. Removing opportunity does not remove ability; it only interferes with its proper exercise. It's best to think of general abilities as something one takes

https://doi.org/10.1017/epi.2014.24 Published online by Cambridge University Press

This is the incompatibilists' argument against (AL2). It seems that the key disagreement between the compatibilists and incompatibilists isn't over whether (IA1) is true, but over whether (IA2) is. If all that the critics of RVE can say in response to this is that all that matters to AC is whether cognitive success results from the exercise of an ability regardless of whether the subject has been afforded the right opportunity, we could easily add an opportunity condition and thereby undermine the anti-luck argument. In the passages above in which the incompatibilists explain why they think AC is met in environmental luck cases it certainly looks as if they think that the subjects have been afforded the right kind of opportunity, so I think we should focus on just this point.

Pritchard and Sosa seem to think that AC is met in environmental luck cases because the subject's abilities put them in touch with things in their surroundings. Let's stipulate the following. First, under normal circumstances, Barney can know by looking that some structure is a barn, Chris can tell by reading the clock that it is 8:22, and Katherine can tell by looking that the surface of the table is red. Second, the abilities that these subjects exercise in the environmental luck cases are the same. Does it follow from this that AC is met in the relevant cases? It might seem so because it follows from these two stipulated claims that (i) the subject has the same abilities in the good case and bad, (ii) the subject exercises the same abilities in the good case and bad, (iii) the result of exercising these abilities in the good case amounts to knowledge, (iv) the relationship between the exercise of the ability and the production of the relevant belief are the same in the good and bad cases, and (v) the subject's visual abilities put her in touch with the same features of her surroundings in the good case and bad. Because of (i)–(iv), it seems that the subject has the right general abilities in the good case and bad to be credited with knowledge. Because of (v), it might seem that the subject has been provided the right kind of opportunity for exercising these abilities. Thus, it might seem that the compatibilists must be right about (AL2) and the incompatibilists must be mistaken about (IA2).

This is too quick, however. To understand what the appropriate opportunities are for the exercise of an ability, it is crucial to spell out what the ability is the ability to do. Some abilities put one in touch with the things in one's surroundings. Some put one in touch with the facts. If these abilities operate differently, it wouldn't be surprising if the circumstances appropriate for their exercise differed. Let's say that *epistemic* abilities are the abilities that put one in touch with the facts and *non-epistemic* abilities are abilities that do not put one in touch with the facts. On some views, one's visual abilities will count as *cognitive* in the sense that they play an important role in the acquisition of true belief without counting as *epistemic* because these abilities never put one in touch with the facts, not even when exercised under appropriate conditions. Travis, for example, seems to defend this sort of view. Picture a piece of raw meat on what was a pristine white rug moments earlier:

> The meat is *in* the surroundings. To see it, look where it is. Look there, too, to see the condition it is in. You can watch the meat – watch it change (in condition or position), watch *for* changes. To see that the meat is on the rug, you *might* look where the *meat* is. You might also look elsewhere – in Pia's face, say (the horrified look). You cannot look 'where that the meat is on the rug'. There is no such place. You cannot *watch* that the meat is on the rug, nor watch for, nor see, changes in it.

with them when they travel and retain when they sleep and think of success manifesting ability as a matter of exercising this general ability under appropriate circumstances.

https://doi.org/10.1017/epi.2014.24 Published online by Cambridge University Press

> It is not eligible for such changes. (You can watch only what you can look for changes in.) Vision affords sensitivity to the goings on in one's surroundings, and to what undergoes them. What one is *thus* sensitive to is not that such-and-such is so … This was Frege's point in disallowing *that the meat is on the rug* as an object of visual awareness. (Travis 2013: 134)

The meat is something that one can bear a purely visual relation to, but the fact that the meat is on the rug is not itself something that one stands in a purely *visual* relation to, not if visual experience relates one to things in one's surroundings and facts about one's surroundings are not *in* one's surroundings.[12] As such, it might well be a mistake to argue that one has been afforded the right opportunity to see that or know that the meat is on the rug on the grounds that one has been afforded the right opportunity to see the bloody hunk of meat (and so see the meat on the rug, the meat that Pia had intended to cook for dinner, the last hunk of meat that butcher sold before closing his shop, what remains of the beef, the thing that Pia threw in a fit of anger, etc.).

What does it take to see the meat on the rug? Perhaps it is that one has the right kind of general abilities and the opportunity is one in which one can visually discriminate the meat from the background. The opportunity to do this might well include the one that one just imagined, a situation in which a red hunk of meat sits on a white rug in a well-lit room while Pia looks on with a horrified expression. It might also be a case in which the meat was disguised to look like a shoe before being placed on the rug. One can see a hunk of meat that looks nothing like a hunk of meat just as one can see one's father when one's father is wearing a hood and is unrecognizable when he is so covered. Thus, one can see a, which is an F, without being in any position to determine whether a is an F.

What are the abilities involved in acquiring perceptual knowledge about the objects visible to us in our surroundings? The abilities involved in visually discriminating the object from its surroundings matter, but they are not the only abilities that matter. The abilities involved in coming to know that some structure is a barn would seem to be classificatory abilities. If one has perceptual knowledge that some visible object is an F, it would seem that there would be something that vision makes available that serves as a basis for classifying the visible object as an F. Presumably, this would be done on the basis of how the object looks. The look of an object would be determined by its sensible properties. Whether an object looks like an F from a certain point of view to a particular subject would be determined by whether any difference in the sensible properties of an object and an object that is F would not make any difference to the subject's experience. To know that something is an F by looking, it seems that vision has to make something available to the subject that would function as a basis for a classificatory judgment by virtue of discriminating the Fs from the non-Fs. That thing would seem to be the way that the relevant visible object looks.[13]

12 Of course, it is controversial whether seeing that something is so involves visual awareness of a fact. McGinn (1999) and Turri (2010) defend the view that we can stand in purely visual relations to facts. For arguments that we cannot, see Brewer (2011), French (2012), and Moltmann (2013).

13 For discussions of perceptual knowledge that expand upon this idea, see Dretske (1969), Millar (2000), or Travis (2005). These approaches give us an understanding of epistemic seeing (i.e., seeing that something is so) on which one does not need to appeal to the high-level contents that figure in Siegel's (2010) account of perceptual knowledge.

https://doi.org/10.1017/epi.2014.24 Published online by Cambridge University Press

Once we have a model for how the relevant abilities involved in the acquisition of visual knowledge are supposed to work, we will have a better idea of what kind of opportunities are appropriate for their exercise. One way to remove opportunity is to prevent the subject from seeing how some object looks. One might do this by dimming the lights, hiding the object under a curtain, or moving it too far off into the distance to be seen clearly. Another way to remove opportunity is to see to it that there is no look that could serve as the visual basis for discriminating Fs from non-Fs.[14] When there isn't such a basis, there is no basis for attributing successful classification to the subject's abilities. If one sees to it that the Fs do not have a distinctive look, a look that distinguishes them from the non-Fs, the subject might occasionally classify the odd F as an F on the basis of how it looks, but it doesn't seem that we could attribute success to something that the subject was sensitive to for the simple reason that there was nothing that the subject could have been sensitive to in responding to the way that an F looked that could have been the distinctive mark of an F.

The guiding idea, then, is this. The facts that we're interested in are facts about the properties that visible objects have (e.g., the color of a table or whether a certain structure is a barn). To know that these facts obtain on the basis of one's visual experiences, one must be able to rely on something that vision provides in classifying correctly the Fs as Fs. This requires, in turn, that the subject has the opportunity to see that certain visual objects look a certain way *and* that there is something that the subject can see that is the distinctive mark of the Fs. If one sees to it that the subject does not have the opportunity to see the mark that is distinctive of Fs, either because one has seen to it that the subject cannot see how the things look or one has seen to it that the way they look is not distinctive of Fs, one has robbed the subject of the appropriate kind of opportunity for exercising her general ability to identify Fs by looking.

Let me offer an example of my own that would seem to suggest that this is the right treatment of the environmental luck cases:

> Jane
>
> Jane is a distant relative from a distant land. She writes to say that she's coming for a visit. You tell her that you'll pick her up at the airport. You don't know what she looks like and she doesn't know what you look like. You write her name on a card and stand outside of the arrivals gate holding it high. A woman sees the card, reads it, says 'Hi, I'm Jane', and you drive her home. You didn't realize it, but there were dozens of cards there that read 'Jane'. Owing to the lighting and the accidental placement of very tall people, she fixated on your card first, read the card, and judged (correctly) that you were her ride.[15]

14 To visually discriminate Fs from non-Fs requires that the Fs have a distinctive look. I take this to be a look that distinguishes the Fs from the easily encountered non-Fs, not a look that distinguishes the Fs from all possible non-Fs. If the look had to distinguish the Fs from all *possible* non-Fs, the requirement would quickly lead to skepticism. On the weaker reading, the discrimination condition does not seem to generate any untoward skeptical results. There are, of course, tricky issues to deal with here because it looks as if certain possibilities in which there are non-Fs that look the way that the actual Fs look are being deemed as somehow irrelevant for determining whether a subject's visually based judgments constitute knowledge. For helpful discussions of how to address these further issues so as to avoid skeptical worries, see Pritchard (2012b) or McKinnon (2013).

15 It helps to think about these cases from the perspective of someone trying to send a sign or signal to get a message across but struggles to do so because the features of the sign or signal are not distinctive of the kind of message that one is trying to convey. In effect, someone who sends such a signal might have the general ability to express disapproval by saying 'No', say, but has been silenced. Thinking about

https://doi.org/10.1017/epi.2014.24 Published online by Cambridge University Press

Had it not been for the other cards, the card that Jane needed would have had a distinctive look. As things stand, however, it did not. There is nothing that Jane was sensitive to that would have clued her in that the card she was looking at would have been the wrong card if she had happened to look at any of the other cards she could have easily seen. By writing her name on the card, you tried to give her the opportunity to identify you by sight and to tell you apart from the other strangers that she shouldn't take rides from. You tried, but failed in that regard. You got lucky in that she found you even though she was not afforded the right opportunity.

I think that it is clear that this case is *not* a case of success that is attributable to Jane's abilities. I also think that it is clear that the reason that this is not such a case is *not* that Jane lacked some general ability, such as the ability to read cards and determine whether the name on the card was 'Jane'. If, as it seems, this case is another case of environmental luck, one that is analogous to Barney, the compatibilist treatment of that case is problematic. What's missing from Jane, I submit, is that she lacked the opportunity she needed to classify the card she saw as the card with *her* name on it on the basis of how the card looks precisely because that card did not have a distinctive look that she was sensitive to. Had all the other cards read 'Jill', however, we would attribute success to her abilities and credit her with knowledge precisely because she had the general ability and she exercised it under appropriate circumstances.

While it's easy to see how this kind of story might work for Barney or Kaleidoscope, does it work for Jane and Chris' Clock? I think so. We tried but failed to provide Jane with the opportunity she needed to identify us by means of our sign. A barn might suffer a similar plight if a barn wanted to be recognized by passersby as such if some joker decides to disguise some nearby non-barns as barns. The barn can show us how it looks and put on a display, but when there is no way that it looks that distinguishes it from the non-barns, there is nothing it shows us in its look that would enable us to discriminate it from non-barns. Suppose the time wanted to tell Chris how things were with her. The time wanted to tell Chris that she was 8:22. Since Chris cannot literally see the time and what time it is, he has to use an instrument like a clock. The demon saw to it that the time's being 8:22 did not have a distinctive look for the demon saw to it that the clock would look the same regardless of whether it was 8:22, 8:23, 8:24, 8:25, etc.[16] The demon, we might say, silenced the time. By virtue of the decision to intervene if necessary to convince Chris that it was 8:22, the demon destroyed the conditions under which 8:22 had a look that set it apart from other times and so robbed Chris of the opportunity needed to know the time by looking at a clock.

Epistemic abilities are abilities that put one in touch with the facts. Our focus has been on epistemic abilities that involve visual abilities that put the subject in touch with things in her surroundings. The opportunities that are appropriate for these abilities are

the case from this perspective rather than focusing exclusively on the perspective of the person receiving the signal helps to make this vivid.

16 I think people are overly impressed by the fact that Chris can see the clock and tell what time the clock says it is. The ability to tell time by using a clock is not reducible to the ability to see the clock, know how to read clocks in general, and the ability to see the position of the hands. It is true that the clock is not *broken*, but it is also true that the clock in the circumstances described does not indicate the time. So, rather than focus on whether Chris can read the clock face, we need to focus on whether Chris can use the clock to tell the time. I think not because the way things look is not distinctive of 8:22.

https://doi.org/10.1017/epi.2014.24 Published online by Cambridge University Press

situations in which some target fact (e.g., the fact that some visible object is an F) are situations in which the visible object has a look that is distinctive of Fs constituted by properties that are visually available to the subject. If the subject does not have visual access to the relevant look by virtue of the fact that the subject cannot visually identify the properties or by virtue of the fact that the properties do not constitute a distinctive look, the subject does not have the right kind of opportunities for exercising the epistemic abilities responsible for perceptual knowledge in fortuitous circumstances. This is why environmental luck precludes knowledge and prevents us from saying that the odd correct perceptual judgment is one whose accuracy is attributable to the subject's ability to correctly classify things on the basis of how they look.

4. ANTI-LUCK

An interesting question to consider at this point in the discussion is whether there is anything left to the original worry, a worry that really had to do with whether AC might serve as a suitable anti-luck condition.

If one thinks that veritic luck is epistemically malignant and one wants to add a modal condition to one's account of knowledge, it might be better to think of knowledge as requiring safety than to think of it as requiring sensitivity. If one thinks that safety is indeed necessary for knowledge, we should ask whether there is any reason to add a safety condition to RVE. Should proponents of RVE say that an additional safety condition is unnecessary because there can be unsafe knowledge? They might say this, but they might not have to. They might say that the *addition* of a safety condition is unnecessary because the safety condition is redundant.[17]

The notion of ability is itself a modal notion. While critics of robust virtue epistemology have argued that cases of environmental luck are cases of unsafe belief that is nevertheless success that manifests ability, I have contested this description of the cases. If correctness manifests ability only when correctness results from the exercise of a discriminatory capacity that correctly classifies the objects one could easily encounter on the basis of some identifying mark, it looks like one must satisfy some sort of safety principle by virtue of satisfying the ability condition.

One of the reasons that I thought that it was odd to say that success was the manifestation of the subject's relevant cognitive abilities in cases of environmental luck was precisely that these seemed to be cases in which the safety condition was not met. In testing whether something was the manifestation of an individual's abilities, we do not simply ask whether it would be metaphysically possible or physically possible for the individual to produce an effect. Sinking a putt from thirty yards is something that I *can* do in

17 Carter (2013) and Jarvis (forthcoming) argue that the modal anti-luck condition is redundant on a proper understanding of what it is for success to manifest ability. I think that they are right about this point and that this is an important point. However, I think there are problems with their overall approaches to these issues. I do not see, for example, how Carter's discussion of the difference between agent- and belief-focused senses of believing truly sheds light on the issue. (The fault might be with me, not Carter's paper.) I worry that Jarvis' account delivers non-knowledge verdicts in clear cases of knowledge because he thinks that knowledge must be 'fully' attributable to the exercise of one's abilities.

https://doi.org/10.1017/epi.2014.24 Published online by Cambridge University Press

the sense that it is something that I'm physically capable of doing, but it is surely not something that's within my abilities.[18] How do we know? Well, for a start, if I did it once and tried it again, I would almost certainly fail. Surely that's a clue.

Fans of Austin will remind us that this point has to be handled with care. After saying that there is 'some plausibility' to the idea that 'I can do X' means 'I shall succeed in doing X, if I try', he remarks in a footnote:

> Plausibility, but no more. Consider the case where I miss a very short putt and kick myself because I could have holed it. It is not that I should have holed it if I had tried: I did try, and missed. It is not that I should have holed it if conditions had been different: that might of course be so, but I am talking about the conditions as they precisely were … Nor does 'I can hole it this time' mean that I shall hole it this time if I try or if anything else: for I may try and miss, and yet not be convinced that I could not have done it … But if I tried my hardest … and missed, surely there *must* have been *something* that caused me to fail, that made me unable to succeed? Well, a modern belief in science, in there being an explanation of everything, may make us assent to this argument. But such a belief is not in line with the traditional beliefs enshrined in the word *can*: according to *them*, a human ability or power or capacity is inherently liable not to produce success, on occasion, and that for no reason. (1961: 218)

Austin's example is a vivid reminder that we cannot give an analysis of ability in terms of subjunctive conditionals. It doesn't follow from the fact that he has the overall ability to sink the putt that he would sink it if he were to try or would sink it if he were to try again. Having said that, it would surely go against the spirit of Austin's remarks to take this too far and to take the possession of ability and the presence of opportunity to tell us nothing about nearby possibilities. Here, I think Kenny (1975: 142) is probably right that it should follow from the claim that someone has the overall ability to pull something off that if they have the opportunity and gave it their best shot they would *normally* succeed or would be *expected* to succeed. This suggests that in the nearby possibilities in which one tries in the kinds of circumstances one is in, the exercise of the ability will normally meet with success.[19]

Does this mean that the safety condition is otiose? It's difficult to say because there is currently a great deal of disagreement about how to best formulate the safety condition. If we opt for a version of safety according to which a belief is safe iff in most of the nearby possibilities where the belief is formed in the same way that it actually is the belief will be true, it looks like when accuracy manifests the subject's abilities, there will not be many

18 As White (1975: 23) reminds us, it's not part of the meaning of 'can,' 'could have,' or 'able' that 'Can V,' 'Could have Vd,' or 'Is able to V' are all ways of talking about an ability to V. I can pull the queen of diamonds from the deck, could have pulled that card from the deck, and was able to pull it from the deck if I did indeed do that, but I don't have the *ability* to do so.

19 Some writers (e.g., McKinnon forthcoming) think that there are cases in which someone has the ability to F where it is very unlikely that they will F even if they try provided that there is some non-zero probability that the subject does F if, say, she tries. This is surely right about whether the subject 'can' F or 'is able' to F in the relevant circumstances, but I doubt that this is right for ability. I suspect the plausibility of thinking that a non-zero probability of F-ing is sufficient for having the ability to F stems from the thought that we use 'can' and 'able' to pick out abilities, a point that White (1975) cautions against. To test whether someone really has the ability to draw a queen from a deck of cards without looking, we look to see if they can pull that off reliably. The issue isn't settled by noting that there is a non-zero probability of drawing that card out of a well-shuffled deck.

https://doi.org/10.1017/epi.2014.24 Published online by Cambridge University Press

nearby possibilities in which the ability is exercised, the circumstances are appropriate, and the belief is mistaken. Suppose one believes on the basis of how the structure looks that it is a barn and the look is indeed distinctive of barns. Could looking that way mean that the structure is a barn or count as the distinctive look of a barn if, say, there are easily encountered non-barns that have that very same look? I think not. So, perhaps built into the very conditions under which barns have a distinctive look are the conditions that ensure that the look is a safe basis for a classificatory judgment.

These points reinforce a point I made earlier about opportunity. The environmental luck cases are set up in such a way that it's guaranteed that the subject will *not* normally succeed if they give it their best shot. Thus, I think it is a mistake to think that the opportunity condition on overall ability is met. It also suggests that some version of the safety condition is satisfied whenever the ability condition is met. It might not follow that there are *no* nearby possibilities in which the ability is exercised and the subject errs, but so long as there are few enough such possibilities, it seems that success can be the manifestation of the subject's abilities. Moreover, it seems that unless these mistakes are rare enough and a sign of the inherent liability to failure that is characteristic of human abilities, powers, and capacities, we would not be willing to say that the subject had both the relevant cognitive ability and the right opportunity for their exercise. The failure of a weak safety condition is a clear indication that the subject lacked general ability or opportunity. Thus, a proper understanding of the ability condition requires us to invoke modal notions like safety.

5. KNOWLEDGE AND ABILITY

The attractions of compatibilism are due to some mistakes about the nature of abilities involved in acquiring knowledge and the nature of ability itself, specifically the relationship between ability, opportunity, and manifestation. When we think about the Barney case and ask whether Barney is able to tell that the barns are barns and that the non-barns are non-barns, the answer seems pretty clearly to be 'no'. In this section, I shall provide an independent argument against compatibilism, one that has to do with the abilities that one should have only once one knows rather than the abilities exercised in coming to know. It seems that there are independent data points that we should consider in evaluating the merits of compatibilism and I think that these data points do not support that view.

The focus thus far has been on the abilities that are involved in acquiring knowledge. These are the *abilities of acquisition*. Let's consider some of the abilities that a subject is supposed to have in coming to possess knowledge. These are the *acquired abilities*. Think about practical abilities. It's often said that an agent who acts in the belief that p will act for the reason that p iff the agent knows p.[20] Given the way the agent is in terms of her

20 See Unger (1975), Hyman (1999), Hornsby (2008) and Marcus (2012). Littlejohn (2012) and Mantel (forthcoming) say that we act for the reason that p when we act in the apt belief that p. If aptness is sufficient for knowledge, the accounts coincide. Also, see Gibbons (2001) for a discussion of the role that knowledge plays in intentional action. For criticism of these approaches, see Hofmann (forthcoming) and Hughes (forthcoming). Hughes uses environmental luck cases to attack these views. In arguing that success shouldn't be attributed to a subject's abilities in these cases, I hope to show that the subject cannot be guided by the fact *p* when her belief about *p* fails to constitute knowledge in cases of environmental luck.

https://doi.org/10.1017/epi.2014.24 Published online by Cambridge University Press

desires, wants, intentions, goals, or plans, a belief might lead her to act in a certain way or try to act in a certain way, but unless the agent knows p, her action will not be correctly described as a case for acting for the reason that p. A similar thesis has been defended about emotion. A person who is angry, sad, happy, regretful, etc. because she believes p will be angry that p, sad that p, happy that p, regretful that p, etc. iff she knows p.[21] It's only when the subject knows, say, that the neighbor ran over her cat that she could be angry that her neighbor ran over her cat or angry because of the fact that her neighbor ran over her cat. A similar thesis has also been defended about belief. A person who believes q because she believes p will only believe q for the reason that p if she knows p.[22] In light of this, it's tempting to accept Hyman's (1999: 441) suggestion that knowledge is the ability to act, feel, or believe for reasons that are facts, there are data that suggest that certain entailments hold between knowledge ascriptions and ability ascriptions.[23]

Suppose that it's false that Marsha went to the store for the reason that she's out of cigarettes if she doesn't know that she's out of cigarettes, false that Bobby is angry that Marsha has taken up smoking if he doesn't know that Marsha smokes, etc. Let's suppose that it's generally the case that the following theses hold:

Practical Ability: S can act for the reason that p only if S knows p.
Doxastic Ability: S can believe for the reason that p only if S knows p.
Affective Ability: S can emote that p only if S knows p.

The evidence that supports these theses seems pretty solid and this evidence should worry the compatibilists.[24] If these three theses are correct, the malignancy thesis tells us that environmental luck would prevent our subjects from being able to act, believe, or feel for reasons that are facts. Why would environmental luck rob us of the acquired abilities when it does nothing to undermine our abilities of acquisition?[25] I would think that the compatibilist rationale for thinking that environmental luck doesn't rob us of acquired abilities would be just as strong as the rationale for thinking that it doesn't undermine our abilities of acquisition, but we seem to have independent evidence that it *does* rob us of our acquired abilities.

If a subject in an environmental luck case who believes p without knowing that p cannot be happy that p, sad that p, angry that p, believe q for the reason that p, or act for the

21 See Unger (1975) and Gordon (1987).

22 In addition to Unger (1975), Littlejohn (2012) and Marcus (2012) defend this view.

23 For further discussion of the view that knowledge is itself a kind of ability, see also Ryle (1949), White (1982), Kenny (1992), and Hetherington (2012).

24 Readers familiar with Gordon and Unger know that the evidence that supports these claims about acquired abilities is independent from the philosophical arguments that support the claim that knowledge is acquired by means of an ability operative under suitable circumstances. Much of the data is linguistic evidence of entailments between knowledge ascriptions and ability attributions that has held up pretty well under scrutiny. The classic objection to these ability accounts is a kind of argument for error. For a helpful discussion of the objection, see Alvarez (2010).

25 One could answer the challenge as follows: environmental luck robs us of the acquired abilities without undermining our abilities of acquisition because the former require knowledge and the latter do not. Surely this is a non-answer if we're operating under the assumption that the fact that one knows is not some brute fact but a fact explained either by virtue of the abilities exercised in forming a belief, facts about the circumstances, and/or facts that determine whether a true belief could have easily been mistaken.

https://doi.org/10.1017/epi.2014.24 Published online by Cambridge University Press

reason that p, I'd think that the malignant features that undermine these abilities would threaten to undermine the subject's abilities of acquisition. If the subject settles the question whether p in such a way that the correctness of her answer is attributable to her abilities and she is subsequently made angry because she believes p, I'd think that it should be within her abilities to be made angry by the fact that p. The compatibilist had better have a good story to tell about why environmental luck threatens some of these abilities rather than others.

One reason to think that they'd have an incredibly difficult time coming up with a compelling story is that there might be implicit links between knowledge and the abilities of acquisition much in the way that there are implicit links between knowledge and the acquired abilities. In explaining why I thought that the compatibilist treatment of Barney was mistaken, I wanted to press the compatibilist to explain what precisely they thought the subjects' cognitive abilities were. With the possible exception of Capability Brown, none of us are simply able. An ability has to be an ability to do something (in a broad sense of 'do'). I focused on the ability to tell whether something was so. If the compatibilist wants to focus on some other sort of ability, that's fine, but it's hard to see how this ability could be irrelevant to the question as to whether subjects like Barney are in a position to know.

It might seem odd to think that there's an ability to tell whether something is a barn if that's taken to mean that there's something resident in the subject that they can carry with them to fake barn country and back that is the subject's ability to tell whether or know whether something is a barn. That's not the suggestion. The suggestion is that the correctness of the subject's judgment about whether p manifests her cognitive abilities only if we can say that the subject is able to tell whether p. *That* is a claim about overall ability (i.e., general ability and opportunity) and it looks like the truth of that claim depends upon whether the subject is in a position to know. Environmental luck ensures that the subject is not. Thus, critics of robust virtue epistemology who are impressed by the case for compatibilism have to tell us where this line of reasoning goes wrong: if the correctness of Barney's belief that the structure is a barn is a manifestation of his cognitive abilities, he was able to tell whether the structure is a barn. He is able to tell whether the structure was a barn only if he was in a position to know that the structure was a barn. As the malignancy thesis tells us, he was *not* in a position to know that the structure was a barn. Thus, the correctness of his belief is *not* the manifestation of his relevant cognitive abilities.[26]

26 I should note that if this line of argument helps to undermine an argument against robust virtue theory, it does point to a potential problem for the theory. If we cannot specify abilities of acquisition without focusing on the ability to tell whether something is so, it is not clear that the robust virtue theory can provide a non-circular account of what knowledge consists in. In characterizing the abilities of acquisition in terms of the ability to tell whether something is so, we've smuggled the concept of knowledge into the account of ability and then tried to use the notion of ability to explicate the concept of knowledge. If this is the best we can do, we should admit that we're part of the knowledge-first program. I confess that I fear that we cannot do any better.

https://doi.org/10.1017/epi.2014.24 Published online by Cambridge University Press

6. OBJECTIONS AND REPLIES

Readers might have the following worry:

> When we are clear that we are discussing abilities to form true beliefs (rather than attain knowledge), it seems that Barney does manifest an ability to form true beliefs about barns. We can define an ability that he doesn't manifest and someone who defends RVE could say that it's this latter ability that matters to knowledge. But then we have a disanalogy between the ordinary cases of success in performance manifesting skills (e.g., archery). So, either Barney doesn't meet the standard RVE ability condition for knowledge or Barney does meet it but this ability can no longer be motivated by the usual analogy to performances manifesting skill.[27]

The incompatibilist faces a dilemma. On the one hand, if they say that the archery case is analogous to the fake barns case, it seems that we should classify the fake barns case as a case of success that is attributable to ability. If, on the other hand, the incompatibilist were to say that the cases weren't analogous, the incompatibilist couldn't appeal to (purported) analogies between performances that manifest skill and accurate beliefs that manifest skill. What should the incompatibilist say in response?

Incompatibilists can offer a two-part response. The first part of the response involves grasping the first horn of the dilemma. Although the cases involving Barney and Robin are analogous, the objection to RVE fails if neither is a case in which accuracy is attributable to the subject's abilities. To see why someone might think success isn't attributable to Robin's abilities, it's helpful to think about one further sort of case:

> Cannons
>
> Michelle used the first edition training manual when learning to fire Acme cannonballs from Acme's top of the line cannon. She would have to take account of things like the elevation of the cannon and her target, the distance between the cannon and the target, and wind speed. To hit a target, she would use the computer to adjust the cannon left or right or adjust the vertical angle of the barrel. This is an incredibly difficult skill to develop, but she was quite proficient at hitting targets. As a result she was chosen to represent her regiment in a competition. She chose a cannonball from her barrel, loaded the cannon, loaded her coordinates, fired the cannon, and the target was struck. What Michelle didn't realize is that Acme had released a second edition of their training manual, one that included new material on the proper use of Acme's cannon with new line of Acme cannonballs. Acme's cannonballs previously had a fixed weight of 20 lbs, but their new line came in different weights (18, 20, or 22 lbs), which meant that the operator would have to adjust the cannon differently to take account of the weight of the ball. Michelle knew nothing of this and was fortunate to have drawn one of the few first-generation Acme cannonballs. Had she drawn any of the other balls, she would have missed her target. Zoe had trained using the second edition of the manual and was reliable at adjusting her cannon to take account of the specific weight of the ball. She chose her cannonball from the barrel, loaded the cannon, loaded her coordinates, fired the cannon, and the target was struck.

If cannonballs came in three weights (18, 20, and 22 lbs) and the coordinates that Michelle entered would have been appropriate for only one of those weights, I don't think we would say that success would be attributable to her ability even though we should say given her training and her past successes that she has developed an impressive

27 An anonymous associate editor for this journal raised this worry.

https://doi.org/10.1017/epi.2014.24 Published online by Cambridge University Press

ability and that her impressive abilities were manifested by her performance. With Zoe, however, we can say that her success *is* attributable to her ability because she's able to adjust in response to the different weights of the cannonballs she has selected to fire.

In circumstances where Michelle couldn't have easily loaded anything but a 20 lb ball into her cannon (e.g., when she was training at her home firing range), we would expect her to hit her target and we would say that success in these circumstances is attributable to her ability. In circumstances where she easily could have loaded a lighter or heavier ball (e.g., in the competition), we would not attribute successful hits using the 20 lb ball to her ability because she would not be disposed to adjust the vertical angle of the cannon to take account of the weight of the projectiles she loaded. The settings that she punches in *could* result in a hit, of course, but they easily could result in a miss. If she uses coordinates that happen to result in a hit, that's luck. That's not so with Zoe, however. Her success is attributable to her abilities.

Michelle doesn't know how to adjust her shots to take account of the different weights of the projectile she's given so as to hit the targets, so while her actions manifest impressive abilities and result in success, we wouldn't say that the particular success is itself attributable to her abilities. This is the second part of the incompatibilist response I want to offer. Her shots do manifest an excellence at ballistics, but that's not the same thing as success being attributable to those abilities. We need to distinguish the performance that manifests an impressive skill or ability from the particular performances in which success can be attributed to ability. I don't think that a performance manifesting skill or ability is itself necessary or sufficient for knowledge because I don't think it's necessary or sufficient for success being properly attributable to ability.

The incompatibilists should say something along these lines. In determining whether a particular success is attributable to ability, we have to take account of the properties of the inputs the subject could easily be given in her situation (e.g., cannonballs, arrows, or the way that a visible object looks), the way that the subject is disposed to respond to inputs with those properties (e.g., firing at such and such an angle, classifying things that look a certain way as being an instance of a kind), and the relationship between these responses in these situations and success (e.g., whether firing as one is disposed missiles with that weight results in hit targets, whether classifying objects with that look results in correct classification). Unless success is to be expected given the subject's dispositions and the properties of the inputs the subject easily encounters, the subject's success is not properly attributable to her abilities.

Let's apply this to the Robin and Barney examples. Robin, let's say, has been trained to fire certain kinds of arrows, but her quiver contains numerous arrows unlike the ones she's competent to use. The arrows that she could easily select to fire have variable weights and we're supposed to assume that Robin's dispositions are such that while she'd be disposed to fire, say, one arrow in such a way that it would reliably strike its target, she wouldn't adjust as she'd need to if she were to reliably strike targets with them. Thus, while we might admire the skill and ability manifested in her shooting, say, because she fires in such a way that she'd hit the target if she used the arrows that hadn't been weighted, we wouldn't say that her accuracy is attributable to her abilities. That would require a disposition to fire the arrows she was given in a way that would reliably result in success, and she doesn't have those dispositions.

For his part, Barney is disposed to classify objects that look certain ways as barns. In the fake barn case, there are non-barns that have the look that would trigger Barney's

https://doi.org/10.1017/epi.2014.24 Published online by Cambridge University Press

disposition to classify an input as a barn that could easily be inputs that would lead to the mistaken classification. We don't attribute success to Barney's disposition to classify objects as barns on the basis of their looks when the look that triggers the disposition isn't a look that means that the relevant object is indeed a barn.

Does this line of response generalize to Pritchard's epistemic Twin Earth case?[28] To understand the case we have to imagine two subjects, Wilfried and Willa, and two planets, Earth and Twin Earth. Wilfried resides on Earth. Willa resides on Twin Earth. Earth, as you know, is filled with watery stuff that's composed of H_2O. On (this) Twin-Earth most of the watery stuff is H_2O, but some is XYZ. Although Willa resides on a planet that has XYZ she's never come into contact with the stuff. Let's suppose that as a result, the extension of 'water' as used by Willa and Wilfried is the same and that it applies to H_2O but not to XYZ. Now, consider three environments on Twin-Earth:

(i) The local environment is the situation Willa is in now.
(ii) The global environment is the region that Willa is in normally.
(iii) The regional environment needn't be where Willa is now or where she is normally. The regional environment is a region that contains objects and properties that Willa could easily encounter.

We can stipulate that there's XYZ in the regional environment and no XYZ in the local or global environment. Although Willa has never come across XYZ, she easily could have. If she were to do so, she'd mistakenly think that it's water.

If we assume that Wilfried knows, say, that the watery stuff he sees is water, we would have to credit success to his abilities. It seems there's nothing wrong with doing that, but now consider Willa. She easily could have come across samples of watery stuff that she'd falsely believe to be water because she easily could have come across some watery stuff that was XYZ. She couldn't tell such samples apart from samples of water (i.e., the liquid, solid, or gas composed out of H_2O), so we don't credit her with knowledge when she judges that some sample is water. Kallestrup and Pritchard maintain that this is a case in which success *is* attributable to Willa's abilities. If they're right, success can be properly attributable to a subject's abilities even if the subject's belief fails to constitute knowledge.

Incompatibilists should challenge their claim that success in this case is attributable to Willa's abilities. When Willa is classifying chemicals, she does so on the basis of the superficial properties of samples and water. The fact that some sample has the properties that triggers her disposition to classify something as water doesn't mean that the sample is water because these properties are properties of liquids that aren't water that Willa could have easily come across. If, as I've suggested above, we shouldn't credit success in classification to abilities if the properties that constitute the basis of the classification aren't properties that are distinctive of the relevant kind, we shouldn't say that correctly classifying water as water is attributable to Willa's classificatory abilities. (For these properties to be distinctive properties of the Fs, there cannot be easily encountered non-Fs that have these properties.) Notice how natural it is to say that her successful classification isn't primarily down to the fact that she's sensitive to the properties of water that distinguish water from non-water. Notice how unnatural it is to acknowledge this but then insist that her

28 An anonymous referee wanted to know.

https://doi.org/10.1017/epi.2014.24 Published online by Cambridge University Press

successful classification should be attributed to her abilities and the way that she goes about classifying liquids as water or non-water.

The problem with the example is the assumption that the regions or environments that matter when determining whether success can be attributed to an individual's abilities are regions or environments characterized in non-modal terms (e.g., regions where the subject is now or is typically found). If the regions are characterized in modal terms (e.g., regions or environments that the subject can easily find herself in), the Twin Earth example doesn't threaten RVE. The examples above suggest that regions or environments should be understood in terms of easy possibilities. Consider Cannons. Because the subject could easily be in a situation where success isn't something we'd expect, the mere fact that the variations in her circumstance that would have frustrated her efforts weren't present doesn't seem to matter because it's clear that she cannot adjust to deal with variations in the circumstance that could have easily been present. Once we see that success is attributable to ability only when the subject can deal with the variations that could have easily come to pass, we can see why opportunities appropriate for the exercise of ability should be understood in modal terms and why the Twin Earth example isn't a threat.

7. CONCLUSION

The environmental luck cases pose no threat to robust virtue epistemology. Arguments to the contrary rest on some mistaken assumptions about the nature of ability and about the abilities involved in the acquisition of perceptual knowledge. Critics of RVE appeal to environmental luck cases because they think (quite rightly) that when you acquire knowledge you're in an environment in which we'd expect that you'd succeed. They also think (quite mistakenly) that this observation causes trouble for RVE. It doesn't. The point is perfectly compatible with RVE because satisfying the AC condition is possible only when the subject's abilities are exercised in appropriate circumstances and these are circumstances in which success is something we'd expect. By focusing on general ability and its exercise and ignoring whether the subject has been afforded the right kind of opportunity, critics of RVE have created a kind of false dilemma, one in which proponents of RVE are told they have to either bite the bullet and ascribe knowledge to subjects like Barney or supplement their view with a further safety condition. As we've seen, however, this objection reflects a mistake about what it takes for cognitive success to be the manifestation of the subject's cognitive abilities. I haven't tried to show that RVE is a fully adequate theory of knowledge, but I hope to have shown that RVE has the resources to handle cases of environmental luck. Doing more than this is beyond my abilities.

REFERENCES

Alvarez, M. 2010. *Kinds of Reasons*. Oxford: Oxford University Press.
Austin, J. L. 1961. *Philosophical Papers*. Oxford: Oxford University Press.
Battaly, H. 2009. Review of Ernest Sosa, *A Virtue Epistemology: Apt Belief and Reflective Knowledge*. *Analysis*, 69: 382–5.
Black, T. and Murphy, P. 2007. 'In Defense of Sensitivity.' *Synthese*, 154: 53–71.
Bradford, G. 2013. 'The Value of Achievements.' *Pacific Philosophical Quarterly*, 94: 204–24.

https://doi.org/10.1017/epi.2014.24 Published online by Cambridge University Press

Brewer, B. 2011. *Perception and its Objects*. Oxford: Oxford University Press.
Carter, A. 2013. 'A Problem for Pritchard's Anti-Luck Virtue Epistemology.' *Erkenntnis*, 78: 253–75.
DeRose, K. 1995. 'Solving the Skeptical Problem.' *Philosophical Review*, 104: 1–52.
Dretske, F. 1969. *Seeing and Knowing*. London: Routledge and Kegan Paul.
Engel, M. 1992. 'Is Epistemic Luck Compatible with Knowledge?' *Southern Journal of Philosophy*, 30: 59–75.
French, C. 2012. 'Does Propositional Seeing Entail Propositional Knowledge?' *Theoria*, 78: 115–27.
Gendler, T. and Hawthorne, J. 2005. 'The Real Guide to Fake Barns: A Catalogue of Gifts for Your Epistemic Enemies.' *Philosophical Studies*, 124: 331–52.
Gibbons, J. 2001. 'Knowledge in Action.' *Philosophy and Phenomenological Research*, 62: 579–600.
Goldman, A. 1976. 'Discrimination and Perceptual Knowledge.' *Journal of Philosophy*, 73: 771–91.
Gordon, R. 1987. *The Structure of Emotion*. Cambridge: Cambridge University Press.
Greco, J. 2010. *Achieving Knowledge*. Cambridge: Cambridge University Press.
—— 2012. 'A (Different) Virtue Epistemology.' *Philosophy and Phenomenological Research*, 85: 1–26.
Hacker, P. 2013. *The Intellectual Powers*. Chichester: Wiley-Blackwell.
Hetherington, S. 1998. 'Actually Knowing.' *Philosophical Quarterly*, 48: 453–69.
—— 2012. *How to Know: A Practicalist Conception of Knowledge*. Chichester: Wiley-Blackwell.
Hofmann, F. Forthcoming. Gettier for Justification. *Episteme*.
Hornsby, J. 2008. 'A Disjunctive Conception of Acting for Reasons.' In F. MacPherson and A. Haddock (eds), *Disjunctivism: Perception, Action, and Knowledge*, pp. 244–61. Oxford: Oxford University Press.
Hughes, N. Forthcoming. 'Is Knowledge the Ability to ϕ for the reason that p?' *Episteme*.
Hyman, J. 1999. 'How Knowledge Works.' *Philosophical Quarterly*, 49: 433–51.
Jarvis, B. Forthcoming. 'Knowledge, Cognitive Achievement, and Environmental Luck.' *Pacific Philosophical Quarterly*.
Kallestrup, J. and Pritchard, D. 2013. 'The Power, and Limitations, of Virtue Epistemology.' In R. Groff and J. Greco (eds), *Powers and Capacities in Philosophy*, pp. 248–70. London: Routledge.
Kelp, C. 2009. 'Knowledge and Safety.' *Journal of Philosophical Research*, 34: 21–31.
Kenny, A. 1975. *Will, Freedom, and Power*. Oxford: Blackwell.
—— 1992. *The Metaphysics of Mind*. Oxford: Oxford University Press.
Kvanvig, J. 2003. *The Value of Knowledge and the Pursuit of Understanding*. Cambridge: Cambridge University Press.
—— 2010. 'Sosa's Virtue Epistemology.' *Critica*, 42: 47–62.
Lackey, J. 2007. 'Why We Don't Deserve Credit for Everything We Know.' *Synthese*, 158: 345–61.
Littlejohn, C. 2011. 'Evidence and Knowledge.' *Erkenntnis*, 74: 241–62.
—— 2012. *Justification and the Truth-Connection*. Cambridge: Cambridge University Press.
—— 2013. 'The Russellian Retreat.' *Proceedings of the Aristotelian Society*, 113: 293–320.
Luper, S. 2006. 'Restorative Rigging and the Safe Indication Account.' *Synthese*, 153: 161–70.
Madison, B. 2011. 'Combating Anti-Anti-Luck Epistemology.' *Australasian Journal of Philosophy*, 89: 47–58.
Maier, J. Forthcoming. 'The Agentive Modalities.' *Philosophy and Phenomenological Research*.
Mantel, S. Forthcoming. 'Acting for Reasons, Apt Action, and Knowledge.' *Synthese*.
Marcus, E. 2012. *Rational Causation*. Cambridge, MA: Harvard University Press.
McGinn, C. 1999. 'The Concept of Knowledge.' In *Knowledge and Reality: Selected Essays*, pp. 7–36. Oxford: Oxford University Press.
McKinnon, R. 2013. 'Lotteries, Knowledge, and Irrelevant Alternatives.' *Dialogue*, 52: 523–49.
—— Forthcoming. 'You Make Your Own Luck.' *Metaphilosophy*.
Millar, A. 2000. 'The Scope of Perceptual Knowledge.' *Philosophy*, 75: 73–88.
—— 2009. 'What is it that Cognitive Abilities are Abilities to Do?' *Acta Analytica*, 24: 223–36.
Miracchi, L. Forthcoming. 'Competence to Know.' *Philosophical Studies*.
Moltmann, F. 2013. *Abstract Objects and the Semantics of Natural Language*. Oxford: Oxford University Press.
Nozick, R. 1981. *Philosophical Explanations*. Cambridge, MA: Harvard University Press.
Pritchard, D. 2005. *Epistemic Luck*. Oxford: Oxford University Press.

https://doi.org/10.1017/epi.2014.24 Published online by Cambridge University Press

—— 2009a. 'Safety-Based Epistemology: Whither Now?' *Journal of Philosophical Research*, 34: 33–45.
—— 2009b. *Knowledge*. London: Palgrave MacMillan.
—— 2012a. 'Anti-Luck Virtue Epistemology.' *Journal of Philosophy*, 109: 247–79.
—— 2012b. *Epistemological Disjunctivism*. Oxford: Oxford University Press.
—— 2012c. 'In Defense of Modest Anti-Luck Epistemology.' In T. Black (ed.), *The Sensitivity Principle in Epistemology*, pp. 173–92. Cambridge: Cambridge University Press.
Roush, S. 2007. *Tracking Truth: Knowledge, Evidence, and Science*. Oxford: Oxford University Press.
Ryle, G. 1949. *The Concept of Mind*. London: Hutchinson.
Siegel, S. 2010. *The Contents of Visual Experience*. Oxford: Oxford University Press.
Sosa, E. 2007. *A Virtue Epistemology: Apt Belief and Reflective Knowledge*, Volume I. Oxford: Oxford University Press.
—— 2009. *A Virtue Epistemology: Apt Belief and Reflective Knowledge*, Volume II. Oxford: Oxford University Press.
—— 2010. 'How Competence Matters in Epistemology.' *Philosophical Perspectives*, 24: 465–75.
Travis, C. 2005. 'A Sense of Occasion.' *Philosophical Quarterly*, 219: 286–314.
—— 2013. *Perception: Essays after Frege*. Oxford: Oxford University Press.
Turri, J. 2010. 'Does Perceiving Entail Knowing?' *Theoria*, 76: 197–206.
—— 2011. 'Manifest Failure: The Gettier Problem Solved.' *Philosopher's Imprint*, 11: 1–11.
Unger, P. 1975. *Ignorance*. Oxford: Oxford University Press.
Whitcomb, D. Manuscript. 'Knowledge, Virtue, and Truth.'
White, A. 1975. *Modal Thinking*. Oxford: Blackwell.
—— 1982. *The Nature of Knowledge*. Totowa, NJ: Rowman & Littlefield.
Whiting, D. 2012. 'Epistemic Value and Achievement.' *Ratio*, 25: 216–30.
Williamson, T. 2000. *Knowledge and its Limits*. Oxford: Oxford University Press.
Zagzebski, L. 1996. *Virtues of the Mind*. Cambridge: Cambridge University Press.

CLAYTON LITTLEJOHN is a Lecturer in the Department of Philosophy at King's College London. His first book, *Justification and the Truth-Connection* (Cambridge University Press), defended an externalist approach to justification. More recently, he and John Turri edited *Epistemic Norms: New Essays on Action, Belief, and Assertion* (Oxford University Press). His current research concerns excuses in epistemology, the nature of epistemic rationality, and the normative significance of ignorance.

https://doi.org/10.1017/epi.2014.24 Published online by Cambridge University Press