

I am a little concerned today regarding the topic of my paper, for at least two distinct reasons. The first is that the conference is on “epistemic value,” and I will hardly mention such value throughout my talk. This may be thought by some to be somewhat inappropriate, and perhaps even obnoxious. The second reason I’m a little worried about the topic of my paper today is that what I will be spending a lot of time on is the notion of luck, and in particular, epistemic luck. Now, in one way, this is a particularly appropriate place to deliver a paper on that topic. Duncan Pritchard, our illustrious host, has recently done a lot of interesting and important work on the subject. His book titled *Epistemic Luck* is an admirably clear and unprecedentedly thorough examination of the role of luck, specifically epistemic luck, in our various epistemic evaluations. This is obviously a place where an essay on epistemic luck will fall on knowledgeable ears. But, of course, I’d have little to say about the subject if I thought that Duncan had gotten the story about luck exactly *right*. Indeed, a large portion of my paper is dedicated to convincing you all (with perhaps Duncan himself excepted) that his account of luck in general, and in particular his account of epistemic luck, are fundamentally misguided. So, not only is my choice of topic inappropriate and possibly obnoxious, I fear it evinces a deep ingratitude on my part for being invited here.

In these opening remarks, I do hope to address the first of these concerns, and explain why I think the topic of epistemic luck is deeply important to at least some important issues involving epistemic value. As for my apparent ingratitude, I will have to rely on Duncan’s implicit trust that I am engaged with him in the spirit of honest inquiry to mutually come to the best understanding of epistemic luck that we can. Failing that, I suppose he can always withhold my reimbursement check.

So why is epistemic luck relevant to a discussion of epistemic value? On my view, epistemic luck is the central notion in the best account available of the value of knowledge. Indeed, I think it is a central notion in the best account of the *nature* of knowledge as well, but that is a topic for another day. The account I have in mind is the so-called “credit theory” of knowledge. It posits that a condition of adequacy on any theory of knowledge is that knowledge turn out to be something for which the knowing agent deserves credit. I have defended this theory in the past, though not in the usual fashion prevalent in contemporary epistemology. I do

not have a version of the theory worked out in sufficient detail to test it against our intuitions in specific real and imagined cases. Various other epistemologists have defended more detailed theories of knowledge that, I would argue, fall under the category of “credit-theories” of knowledge. I have in mind the theories of John Greco, Ernest Sosa and Linda Zagzebski, for instance. One measure, then, of how plausible the “credit theory” of knowledge is would be how well these more detailed theories fare by the usual methods. While I think each of these theories is as plausible and powerful an account of knowledge as most, I am not content to rest the plausibility of the credit-theory on their specific successes or failures. Indeed, I have defended the credit theory elsewhere on the basis of its explanatory power in explaining (1) why both internalist and externalist conditions on justification are necessary for an adequate account of knowledge, (2) why it seems as though internalists and externalists are rival accounts of the same phenomenon, even though they are not, and (3) why knowledge is more epistemically valuable than mere true belief.

So now it all becomes clear. The credit theory of knowledge best accounts for the value of knowledge, and epistemic luck is a crucial notion in the development of that theory. Thus, I can now spend the rest of my time talking about luck with a clear conscience, with the connection between epistemic luck and epistemic value apparent to all. Well, perhaps, but the tie between the credit theory and epistemic value is tighter than this rather shallow discussion would have it. After all, “credit” is itself a notion deeply steeped in value. Credit is closely related to notions like praiseworthiness (and blameworthiness if we adopt a usage that does not imply positive value in every case), responsibility, attribution, and so on. Accepting the view that knowing is essentially a matter of credit-worthy true believing puts the phenomenon of knowledge in the same family of notions as responsible action. It suggests that as knowers we are evaluable in many of the same ways that we are evaluable as actors. This goes against recent trends in externalist epistemology that seem to treat believers as mere cognitive machines, that either perform well or not at their assigned task of believing only truths. The credit theory also makes explicit some deep parallels between epistemology and ethics. In short, the credit theory brings to the fore dimensions of value that might be overlooked, and even denied, on other theories.

I have not yet explained how luck is central to the credit theory, and as I want to get on to talking about luck itself, I will have to briefly lay out the notion of credit I have in mind and how

luck is involved in it. In a nutshell, deserving credit for some state of affairs is largely, if not entirely, a matter of being sufficiently responsible for that state of affair's obtaining. (There are all kinds of tricky issues here that I am simply going to have to bypass for present purposes.) One is responsible for a state of affairs if its obtaining is due to (a) your abilities, and (b) your intentions. In other words, for you to be deserving of credit for some state of affairs, that state of affairs has to have come about in a way that involves you as an agent. It has to be something that you can be said to have *brought about*, at least in a sufficiently robust sense. That requires that at least two kinds of circumstances be ruled out. First, it cannot be that the state of affairs is simply a fortuitous accident as far as you are concerned. That is the point of insisting that the state of affairs comes about as the result of something that you are able to *do*—as a consequence of your abilities or skills, not simply as a consequence of your happening to be standing at a particular place, say.

Second, your contribution to bringing the state of affairs about cannot be inadvertent. Whether by virtue of an explicit intention, or by way of a standing desire or more-or-less unconscious internal monitoring system, the state of affairs must have come about in part because of your *meaning* for it to. If an agent's contribution to bringing about a state of affairs is subject either to luck or to inadvertence, then the agent is not responsible for the state of affairs, and so it cannot be credited to her. These represent two different kinds of luck, each of which can undermine credit.

I take these to be perfectly general features of our notions of credit, responsibility, and luck. I think they hold whether we are talking about luck in action or epistemic luck. Thus, a general account of luck should be possible that gives a good account of lucky events, and that can be put to use in an anti-luck account of knowledge as well. I think that such an account is possible, and it is that account that I want to advance in this paper. This account of luck is one in which agent control, or its lack, is the prime determinant of whether some event is lucky for some given agent.

What follows will be a two-part defense of this account of luck. The first part will consist of a critique of the best-developed rival account of luck around—Pritchard's "safety account" of luck. I will argue that his modal definition of luck falls to decisive counterexamples. I will also argue that the specific instantiation of this definition that he develops to account for epistemic luck fails as well, as does his subsequent "safety" theory of knowledge which is presented in

terms of such luck. The second part of the defense will be to respond on behalf of the “control account” of luck to the most important objections raised to it, and to show that not only do the objections fail, their failures offer us some interesting insights into the nature of luck that other accounts do not.

### **Pritchard’s Safety Account of Luck**

My goal in this section is to show that both Pritchard’s account of luck in general, and his account of epistemic luck in particular, are mistaken. Pritchard’s general account of luck consists of two principles. They are as follows:

L1: If event E is lucky, then it is an event that occurs in the actual world but which does not occur in a wide class of the nearest possible worlds where the relevant initial conditions are the same.

L2: If an event is lucky, then it is an event that is significant to the agent concerned (or would be significant, were the agent to be availed of the relevant facts).

While L2 raises some interesting issues of its own, I will not have anything to say about it for the moment. I will concentrate, as other recent commentators have, on L1.

The usual strategy one follows when attacking someone’s analysis of something or other is to show that the conditions offered are neither necessary nor sufficient, as evidenced by various counterexamples. I will be following this method here. I think the case against the necessity of Pritchard’s two principles is quite powerful, and it has been made by, among others, Jennifer Lackey. I will briefly recount her argument for this in a moment. As for the sufficiency of Pritchard’s principles, I am skeptical here too, though the arguments for this position are not nearly as potent. Nonetheless, I will make the best case I can against the sufficiency of Pritchard’s account, because even if one does not find the argument decisive, it still adds some weight to the overall argument against the safety analysis of luck. [I’m calling it the “safety” analysis because of the similarity of its modal structure to safety accounts of knowledge, one of which he goes on to develop himself.]

In a very penetrating review<sup>1</sup> of Pritchard's book, *Epistemic Luck*, Jennifer Lackey makes a very strong case against the necessity of Pritchard's two conditions on luck. To do this, one needs to show that an event can be (intuitively) lucky even though it occurs in all or nearly all nearby possible worlds. She does this by way of a counterexample, which she calls BURIED TREASURE. I shall recount it in full so as to see the force of her argument.

BURIED TREASURE: Sophie, knowing that she had very little time left to live, wanted to bury a chest filled with all her earthly treasures on an island she inhabited. A suitable site had to be located on the northwest corner of the island—where she had spent many of her fondest moments in life—and it had to be a spot where rose bushes could flourish—since these were her favorite flowers. As it happens, there was only one spot on the island that met both these criteria. Sophie, being excellent at detecting rose-friendly soil, immediately located this spot and buried her treasure there. One month later, Vincent, a distant neighbor of Sophie's, was driving on the NW corner of the island—which was also his most beloved place to visit—and was looking for a place to plant a rose bush in memory of his mother who had died ten years previously—since these were her favorite flowers. Being excellent at spotting rose-friendly soil, he immediately located the very spot that Sophie had buried her treasure one month earlier. As he began digging a hole to plant the rosebush in, he was astonished to find buried treasure there. [Lackey, p??]

Lackey points out that Vincent's discovery is both clearly lucky, and yet also clearly happens in all nearby possible worlds. This is because all the features of the case that led both to Sophie's burying the treasure where she did as well to Vincent's digging where he did are quite central features of the actual world. Thus, one has to go to a quite distant world in order to find one where either Sophie failed to bury her treasure in that spot or where Vincent failed to dig in that spot. Thus, only in a far distant world does Vincent fail to discover buried treasure, despite his doing so being a paradigm instance of luck.

Lackey claims that any number of such counterexamples can be constructed against Pritchard's view. She even offers us a formula for constructing them.

---

<sup>1</sup> Lackey's review of *Epistemic Luck*

To do so, first choose a paradigm instance of luck, such as winning *Jeopardy* through a purely lucky guess, ... Second, construct a case involving such an event in which both central aspects of the event are counterfactually robust, though there is no deliberate or otherwise relevant connection between them... Voilà, you have a counterexample to [Pritchard's account of luck]. [Lackey, p ??].

Counterexamples are useful, of course, in showing that there is something wrong with a view, but they are unhelpful in making further progress unless they also give us some idea of how to diagnose the failure. Lackey's commentary on her own counterexample points us in the direction of such a diagnosis.

The fundamental problem with such modal accounts is that counterfactual robustness can be ensured through a combination of features that is entirely fortuitous. For instance, an event that appears in both the actual world and all of the relevant nearby possible worlds can nonetheless be lucky because the relevant counterfactual robustness is achieved purely through a lucky combination of external events. [Lackey, p ??]

Lackey notes that the problem lies in the fact that the counterfactual robustness of the lucky event is itself fortuitous in the problem cases she describes. There is no "deliberate or otherwise relevant connection" between the two elements of the case whose conjunction yields the luck (e.g., Sophie burying the treasure in spot X and Vincent's digging at spot X). So, presumably, for an account of luck to get these Lackey-type cases right, it must differentiate between situations in which there is such a relevant connection between the elements of luck in the case, and situations in which there is not. Pritchard's safety account does not do this.

Let me turn now to a consideration of the sufficiency of Pritchard's conditions on luck. To show that these conditions are not sufficient to account for luck, one must describe an event which (1) obtains in the actual world, (2) fails to obtain in a wide class of nearby possible worlds, and yet (3) is nevertheless not lucky. It is hard (perhaps impossible!) to find cases like this where the non-luckiness of the event is intuitively very clear. But let me try anyway. Consider the following case:

Smarty is the valedictorian of her high school class who is about to take her computer-delivered college entrance exams. Despite her formidable intelligence, she decides to prepare for the upcoming exam by studying diligently and taking many practice exams. The night before the exam, she gets a good night's sleep, and awakens fresh, sharp, and ready to excel. She takes the exam and scores very highly. Unbeknownst to Smarty, however, a fiendishly clever hacker with debilitating test anxiety had decided prior to the exam to wreak vengeance on all the clever students about to take it. Driven mad by his inability to get into a good college because of his poor test scores, he has vowed that all those smarty-pants test-takers will suffer just as he has had to suffer. He compiles a list of all the high school valedictorians for that year, and hacks into the exam program. For the valedictorians, he replaces the usual questions with questions from an advanced college physics exam. As it happens, he accidentally skips Smarty's name on the list (despite being very careful not to miss anybody), and so she gets the usual questions. As a result, she is the only valedictorian who did well on the exam.

The question to be settled is this: was Smarty's high score on the exam a matter of luck? For my own part, I find that I am pulled in two directions here. I am willing to say that Smarty is lucky that her questions were not changed, but that seems to me not to amount to saying that she is lucky to have done well on the exam.

To bring out this point more clearly, suppose that she had gotten the advanced physics questions. Forced to simply guess in every case, she nevertheless chooses the correct answer in virtually every case. Here I think it is clear that her doing well on the exam is due to luck. And yet, there seems to be no obvious difference between the two cases in terms of how they fare according to Pritchard's safety condition. In each case there is a large class of nearby possible worlds in which Smarty does not get many correct answers on the exam. In the original case, this is due to the fact that in most nearby possible worlds, the hacker does not make the mistake of missing Smarty's name. In the revised case, it is due to the fact that Smarty guessed differently in many of the nearby possible worlds. So even if it is not clear that in the original case Smarty is not lucky to have done well on the exam, there still seems to be a difference in degree between the luck involved there and the luck involved in the revised case, even though there is no obvious

difference even of degree in how “unsafe” they are in the two situations. What this suggests to me is that the safety account of luck is not looking in the right place for the appropriate conditions for luck.

## **Epistemic Luck & Knowledge**

The first thing to point out about Pritchard’s theory of epistemic luck is that he introduces a distinction that does not arise in his more general discussions of luck. He distinguishes between what he calls *veritic* epistemic luck and *reflective* epistemic luck.

*Veritic epistemic luck*: it is a matter of luck that the agent’s belief is true.

*Reflective epistemic luck*: given only what the agent is able to know by reflection alone, it is a matter of luck that her belief is true.

Veritic luck is simply the application of his modal definition of luck to the specific event of coming to hold a true belief. Reflective luck, by contrast, is a slightly more convoluted application of the general principle. While there are some interesting issues surrounding Pritchard’s understanding of reflective luck, like most of his commentators to date, I will focus primarily on veritic luck. The main reason for this is that the definition of knowledge that Pritchard endorses takes account only of veritic luck.

I hope to have cast doubt on his general modal approach to luck, but I want to show how this approach fails specifically in an account of knowledge. Pritchard defines knowledge as, essentially, non-*veritically-lucky* true belief. So, putting his account of luck in place as the element that turns true belief into knowledge, we should get the following: S knows that p iff S believes p, p is true, and there is no wide class of nearby possible worlds in which S believes p when p is false. This is, in fact, more or less the penultimate version of Pritchard’s safety account of knowledge. However, he is forced to revise it to make the constraints on knowledge much stricter by requiring that there be *no* nearby possible worlds *at all* in which S believes p when p is false. He does this under pressure from what he takes to be the correct analysis of lottery-type cases. Here is a paraphrase of Pritchard’s definition of luck in terms of modal safety.

S knows that p (where p is a contingent proposition) iff S believes p in the actual world, p is true in the actual world, and in all nearby possible worlds in which S forms her belief about p in the same way as she does in the actual world, S believes p only when p is true.

(I have reformulated this principle to make certain things more obvious, but I have taken pains not to change its substance. The only possible exception to this is that I have rendered the safety condition as both necessary and sufficient for knowledge, where Pritchard makes it only necessary. However, it seems clear from the context that he means for it to be a full account of knowledge, so I will make this explicit by treating it as both necessary and sufficient.)

Because this is simply a special application of his more general account of luck, one would assume that Lackey's formula for generating counterexamples would serve to undermine this account of knowledge as well. Pritchard is claiming with this principle that knowledge implies a non-veritically-lucky true belief, and that veritically lucky true belief implies a lack of knowledge. But, if Lackey is right, we should be able to construct cases wherein an agent comes to hold a belief in a counterfactually robust way, and the belief is true in a counterfactually robust way, but where the connection between the way the agent came to hold the belief and the truth of that belief are not connected in the right sort of way, and so their having come to hold a true belief is lucky, despite meeting the modal conditions of the safety principle.

Indeed, these cases are constructible. What they amount to are counterfactually robust Gettier cases. Consider:

Suppose the North Korean government undertook a publicity campaign to inform its citizens that all Americans are racists who particularly despise Asians. Further suppose that this is done very convincingly, so that most North Koreans come reasonably to believe that all Americans are racists who particularly despise Asians. Meanwhile, the U.S. government suspects that the N. Koreans have a secret nuclear facility near a village deep in the heart of the country. The CIA decides to send a covert operative to investigate this possibility. As it happens, there is only one agent who even comes close to having the training and skills necessary for this mission. He also happens to be a racist who particularly despises Asians. This prejudice is a deep and longstanding feature of his personality, produced by a series of life-shaping and counterfactually robust incidents in

his past. When the agent arrives at the village, a local villager spots him and immediately recognizes him as an American. The villager subsequently comes to believe that the person he sees is a racist who particularly despises Asians.

It looks as though the villager's belief is true, and that there are no nearby possible worlds where he comes to this belief in the same way and the belief is false. All nearby possible worlds where the villager comes to this belief in the same way he does in the actual world are worlds where the very same operative with the very same prejudices is the subject of that belief. This assumes that both the operative's skill set and his racism are deep properties that hold of him in all nearby possible worlds, and that no other American operative would have been successful in arriving at the village at all. Yet, it is still a matter of luck that his belief is true. So, Lackey's formula holds for constructing counterexamples to the safety account of epistemic luck as well. Consequently, Pritchard's definition of knowledge is faulty because the villager's belief meets the requirements of his definition though he clearly does not know. More damning still, he doesn't know *precisely because* his being right is a matter of luck, which is the consideration that Pritchard's account is supposed to be sensitive to.

But what about my rather weak argument against the sufficiency of Pritchard's conditions on luck. Do these carry over to similarly undermine his account of knowledge? The astute student of Gettierology will already have anticipated a somewhat counterintuitive commitment of my critique at this point. For what would an epistemic analogue of the Smarty case look like? Here's a proposal, chosen randomly from among the various alternatives.

Imagine that I am traveling through a picturesque countryside and I pull off into a driveway to a farmhouse because I like to look at barns. Indeed, as I approach the house I can see a barn quite clearly off in the adjacent meadow. Having accomplished my goal, and having formed the belief that "there's a barn," I drive away. As you might have surmised, I am correct in my belief. Yet, the neighbors for miles around have secretly erected fake barns to fool just such casual barnspotters as myself, and had I chosen any other driveway I would have been fooled by one of their fake barns.

So here is a case in which, I submit, it is not obvious that it is a full-blown matter of luck that I came to have a true belief. Contrast this with a case in which I am, say, riding an express train that is going so fast that I cannot distinguish one building from another as they speed past my window. I am thinking nostalgically of my upbringing on a farm, and remember quite fondly the barn where I used to milk the cows. Glancing briefly outside, and being primed by my reminiscences, I believe of the next building that I see go flashing past, that it is a barn. And, of course, it happens to be a barn. Let us construct both this case and the previous one so that the truth of my belief is equally counterfactually robust. That is, the ratio of fake barns to real ones is roughly the same as the ratio of barns to other buildings that are flashing by my train window. So in both cases the nearby possible worlds are equally populated with instances of my believing “there’s a barn” where that belief is false. I submit that, intuitively, it is not nearly so obviously a matter of luck that I am correct in the original case as it is in the modified case.

But that would imply that I think his theory should *not* discount my belief in the classic barn façade case as an instance of knowledge. But it is an item of contemporary epistemological dogma that discounting one’s true belief in the barn façade case as an instance of knowledge is a *desideratum* of a theory of knowledge, not a point of criticism. True enough. I cannot offer much of a defense of this judgment on my part, other than to point out that I am not quite alone in my judgment. Both Bill Lycan and Stephen Hetherington have recently argued that the barn façade case is not a real Gettier case because it represents a genuine case of knowledge. I also have anecdotal evidence from a number of epistemologists who say that they have always found the barn façade case to be the least persuasive of the standard menagerie of gettier cases. For now, I must simply leave this as I did my initial criticism of the sufficiency of Pritchard’s account of luck. Clearly, this criticism is not going to convince anyone on its own, but it is a significant part of the case to be made against safety accounts of luck and knowledge nonetheless.

### **‘Out of Control’ Luck**

Let me turn now to my own conception of luck that I take to be superior to the safety account. As I have indicated, the notion of luck I defend places agent control (or, rather, the lack thereof) at the forefront of the analysis. One has control over some happening to the extent that the happening is properly considered something the agent has *done*. As I have also already

indicated, this imposes two separable requirements. First, the event has to be the product of the agent's powers, abilities or skills. Second, the event has to be, at least in some attenuated sense, something the agent *meant to do*. This second requirement does not demand an actual conscious intention on the part of the agent, but it does mean that a goal or desire or intention must be guiding the exercise of one's powers, abilities or skills that brings about the event in question.

I think that even this rough description of the control account of luck is enough to see that it gives the right answer in ordinary, run-of-the-mill examples of lucky occurrences. Winning the lottery or making a lucky pool shot are prime examples of events that are not brought about by an agent's powers, abilities, or skills. They are simply the result of the operation of chance, which dictates that, given time, pretty much anything can happen, and occasionally will. Other examples, like finding buried treasure while digging in your flower garden, are doubly lucky. Not only is it not the result of your powers, abilities, or skills that you found buried treasure, your desire or intention or goal of finding buried treasure, if you even have one, played no role in your choosing to dig right there.

Furthermore, I think that this account of luck handles easily the cases discussed previously with regard to the safety account of luck. The cases I have in mind are the buried treasure case, guessing the right answer on jeopardy, and the two cases involving Smarty the valedictorian. Consider first the buried treasure case. Recall that Vincent was looking for an appropriate place to plant roses in memory of his mother. He had no intention of looking for buried treasure, so his finding it is clearly lucky on my account. Moreover, his finding buried treasure was in no way due to his abilities, etc. Guessing the right answer on jeopardy is a different matter. In that case, presumably, you desire to win the game, and that desire is very much governing your deployment of your powers, abilities and skills in coming up with an answer. But, by hypothesis, those powers, abilities and skills were not sufficient for you to know the answer. Therefore, your giving what turns out to be the correct answer is not due to any such powers, abilities, or skills.

And what of my two examples intended to show the insufficiency of the safety account of luck? In the first example, Smarty the valedictorian scored well on his college board exams because his nemesis inadvertently failed to change his questions. What does the control theory of luck say about this case? On the one hand, his score very much reflects his powers, abilities and skills. He knew the answers to the questions given because of his native intelligence and his

study. On the other hand, he would have failed miserably had the hacker not overlooked his exam and had changed his questions as he did for the others. Smarty had no control over this aspect of the situation, which seems to be crucial to his success. Notice that there is no such ambiguity in the second Smarty case, in which the hacker did change his questions, and Smarty managed to guess correctly on most of them. In that case, his high score on the test does not in any way reflect his powers, abilities or skills. This generates some ambiguity concerning whether Smarty does or does not have control over his performance on the exam. This ambiguity, I maintain, explains the difference in our intuitive responses to these two cases.

Let me now turn to two objections that Pritchard gives to the control account of luck. The first objection is both very simple and very powerful. If luck is simply a matter of some event or state of affairs being out of one's control, then all manner of mundane natural occurrences, like the rising of the sun each morning, should be counted as lucky. For that matter, many unnatural occurrences, like the continued functioning of the electricity in my house every day, are not in my control, and thus should be counted as lucky on my view. This seems counterintuitive, to say the least.

The second objection Pritchard gives is addressed more specifically to the account of epistemic luck that would be implied by the control theory of luck. It seems that beliefs, particularly perceptual beliefs, are *not* subject to our control, and thus our very having of them, never mind their being true, would be a matter of luck on my account. Thus, perceptual knowledge (as defined by the credit theory) would appear to be impossible. I will address both of these objections in turn. My response to the first requires a modification of the control view of luck, resulting in additional resources for answering the second objection.

## **Luck & Exploitation**

The first objection is to the sufficiency of the control theory of luck. Something's being out of S's control does not suffice for its being a matter of luck, or so it appears. As Pritchard says, acknowledging Andrew Latus for the point, "...the rising of the sun this morning was an event the occurrence of which was out of one's control. But would we really want to say that it was *lucky* that the sun rose this morning." Pritchard presumably expects our answer to this question to be "no." However, my answer is "it depends." One can certainly tell a story in which

the rising of the sun is lucky for someone. This would be problematic for Pritchard's account of luck, since presumably all nearby and even pretty remote possible worlds are such that the sun rises each morning, so any example of a lucky sunrise would be problematic for his view. It is less so for the control view, however. Indeed, one response on behalf of the control view is simply to appeal to his own principle (L2), which, you will recall, requires that a lucky event be of some significance to the agent. Since the rising of the sun each morning is a fairly mundane occurrence most of the time, we could explain its non-luckiness in that fashion. Principle (L2) seems perfectly detachable from the safety account of luck and perfectly compatible with the control account of luck.

While this would dodge the initial force of the objection, it will ultimately be inadequate because there are still problem cases involving events and states of affairs that are beyond one's control, are of significance to the agent, and yet we still do not want to count as lucky. To make this concrete, I will describe yet another scenario. (I apologize in advance for both the complexity and the political incorrectness of my example. It's all I could come up with in time to send the paper off!)

The scenario involves two gentleman adventurers in the golden era of such adventures. Let us call them Indiana Jones and New Jersey Smith. New Jersey Smith plans an expedition into the wilds of Africa where certain tribes of Africans with exotic customs were known to live. Smith is constrained by his schedule and finances to make this trip during a particular month of a particular year. He proposes the trip to his fellow adventurer Jones, including the specific times that he means to travel. Jones agrees to tag along. As it happens, the particular tribe that lives in the area that Smith and Jones visit has a custom of sacrificing people from outside the tribe on the equinoxes of the year. The autumnal equinox happens to fall during the time that Smith and Jones are in the area, so they are captured and held until that day so that they can be sacrificed. When the day of the autumnal equinox dawns, the tribe readies their captives for sacrifice at midday. As the tribesmen approach to kill them, Smith says to Jones, "Only a miracle could save us now!" At that precise moment, there is a total eclipse of the sun. The members of this tribe always take such exotic natural occurrences to signal the anger of their gods at them for whatever they happen to be doing at the moment. Consequently,

they set their captives free. Smith says to Jones, “that solar eclipse was an amazing stroke of luck!”

Smith certainly seems to be right—it was very lucky for him that the solar eclipse happened when it did. So it is possible for such nomically certain events to be lucky for someone. So far we have no problem for the control theory. But there is more to the story. You see, Jones has a surprising and enlightening reply to Smith’s breathless declaration of their good fortune.

Smith says to Jones, “that solar eclipse was an amazing stroke of luck!” Jones replies, “Don’t be absurd! There was nothing lucky about it. I knew all along that these people would likely try to sacrifice us on the equinox if we were captured, but I also knew that there was a total eclipse of the sun due on that very day, and that this tribe would react to that event by letting us go. Did you really think I would be stupid enough to fall into such a situation without having a plan to extricate myself?”

According to Jones, the eclipse was *not* a matter of luck for *him*. But how can this be? He was no more in control of the eclipse than was Smith. And the eclipse was equally significant to him as it was to Smith. By any criteria that we have appealed to so far, these two should be on a par with respect to the luckiness of the eclipse. Yet both seem right: Smith was lucky that eclipse happened, and Jones was not.

Figuring out what to make of the luckiness or not of the eclipse is a subtle business, and it might help to pin down some of the clearer implications of the scenario first. It seems clear, and accords with the control theory of luck, that Smith is *lucky to be alive*, while Jones is not. Jones’s *being alive* is a consequence of his deploying his powers, abilities and skills in planning against his possible capture by the tribesfolk. Smith’s being alive, on the other hand, has nothing to do with his powers, abilities or skills. So far, so good. But what of the eclipse?

It is tempting to say that, since each adventurer’s survival depends crucially on the eclipse, that each one’s survival is lucky iff the eclipse is lucky for him. This, at least, gets the facts right. Both the eclipse and his subsequent survival are lucky for Smith, and neither is lucky for Jones. But, alas, things get more complicated, as we shall see.

Suppose that Smith knew, just as well as did Jones, that an eclipse was due that day, but he did not know that such an eclipse would save his life if he were captured. Would it still make sense to say that he was lucky that the eclipse occurred? That seems a little strained. He would, of course, be lucky that the tribesfolk had the beliefs they did, which resulted in Smith being set free. And he would still be lucky to be alive. Indeed, Smith could have known everything that Jones did, but if he never put the information together and took account of it in order to make sure he would be safe, then he would still be lucky to be alive. What seems to distinguish Jones from Smith, and makes Smith lucky to be alive but not Jones, is not that Jones *knew* about the eclipse and whatnot, but that he *exploited* those facts to his own advantage. That is to say, he took them into account and planned a course of action that assumed that those things would occur. And the outcome that resulted, his survival, was a consequence of his having taken account of and exploited those facts. Thus, he was in control of his own destiny, even though he was not in control of every event that played an important role in that destiny.

Once this is pointed out, it seems obvious. After all, to the extent that we are ever in control of our actions, it is only against the backdrop of our environment. We can successfully plan and carry out actions only if the world cooperates by behaving more or less as we expect it to. But we still have not figured out how to make sense, on the control view of luck, of the fact that the eclipse itself was lucky for Smith and not for Jones. I think that the preceding discussion shows that a modification of either my definition of luck or my definition of control is required. Up until now, I have been assuming roughly the following:

E is lucky for S iff E is (too far) out of S's control, and  
S controls E iff S brought E about (where this implies both that E was the result of the application of S's powers, abilities or skills, and E was not inadvertent with respect to S).

Putting these together, we get:

E is lucky for S iff it is not the case that S brought E about (where this implies that either E was not the result of the application of S's powers, abilities or skills, or E was inadvertent with respect to S).

But the preceding discussion shows that an event can be both beyond someone's abilities to produce or even affect, and yet still fail to be lucky. That is Jones's situation with respect to the eclipse. What must be recognized here is that S's lacking control over something is not *sufficient* for it to be lucky for S. Thus the biconditional linking the two is incorrect. Instead, we must define luck this way:

E is lucky for S iff

- (a) E is (too far) out of S's control, and
- (b) S did not successfully exploit E for some purpose.

There is obviously a lot of work that needs to be done in getting clearer on what counts as "exploiting" something in this context. For now, I will have to rely on an intuitive understanding gained from the foregoing examples.

It is important to realize that this does not affect the definition of credit, which is still in terms of control, not luck in general. It would be a mistake to start giving Jones *credit* for the eclipse, just because he wasn't lucky that it happened.

We still have a problem, though, because there are all kinds of events that are out of my control and that I do not exploit, yet seem not to be a matter of luck because they do not impinge on my life in any way—the movements of the stars in a distant galaxy, for example. The response I'd like to make here is to bite the bullet and insist that the intuitions give way to the theory in this case. There is a sense of "luck" that is entwined with our concept of "fortune," as in one's fate, or what befalls one in life. It is the pressure of this sense of luck that makes us uneasy calling something that is completely irrelevant to our fortunes a matter of luck. Yet this sense of luck is completely separable from the sense of luck I am trying to get clear on, which is the sense of luck that is conceptually connected to credit and responsibility, rather than fortune. Yet I fear this may strike some as too presumptuous, especially for a view that is still in the early stages of its development and defense. The sell-out alternative, then, is to acknowledge the need for something like Pritchard's principle (L2), which added the additional requirement that a lucky event had to be one that was significant to the agent concerned. So now we have the following:

E is lucky for S iff

- (a) E is (too far) out of S's control, and
- (b) S did not successfully exploit E for some purpose, and
- (c) E is significant to the agent concerned (or would be significant, were the agent to be availed of the relevant facts).

## Epistemic Luck

Pritchard's second objection to the control theory of luck has to do more specifically with the account of veritic epistemic luck that would follow from it.

[T]he issue of control is particularly problematic when it comes to epistemic luck, because (on most views at least) belief is a component of knowledge, and it is certainly common to regard the formation of at least one's most basic perceptual beliefs as not being within one's immediate control. Nevertheless, it seems odd to argue on this basis that basic perceptual belief is inherently 'lucky'. [Pritchard, p. 127]

It is hard to know quite what to do with this objection, initially. Both Pritchard and I want ultimately to give an account of knowledge that is parasitic on our respective accounts of luck. In order to do that, the account of luck must allow for the possibility of perceptual knowledge. That requires that it be possible for someone to have a true perceptual belief that is not a matter of luck. Pritchard seems to be implying that a control theory of luck cannot allow this possibility. But he puts it in a rather puzzling way. After all, it is the luckiness or not of having a *true* perceptual belief that must be accounted for here, not the luckiness or not of our merely *having* a perceptual belief. It may well be true that we do not have immediate control over our perceptual beliefs, but that is a non sequiter in a couple of different ways. First, a control theory of luck need not be committed to our having *immediate control* over anything. We consider ourselves able to do a great many things that require cooperation from the world, and that we subsequently do not have *immediate* control over. This is a big issue, and I cannot pursue it further here.

But the second sense in which the objection seems to be a non sequiter is that the relevant question here is not whether one is lucky or not to have the belief, it is rather whether one is

lucky or not to have had a true belief in this instance. I think the case of Jones and the eclipse has shown us that it is possible for an event to fail to be a matter of luck for someone even though an immediate and important causal contributor to that event is not under the person's control. Jones was not in control of the eclipse, but he exploited it to arrange his own survival. Similarly, we may say that an agent does not control her perceptual beliefs, in the sense of bringing them about, but it can nonetheless fail to be lucky when such beliefs turn out to be true. This may sound a bit fanciful, so let me offer an analogy.

I have urged elsewhere that we think of our control over our perceptual beliefs along the same lines as we think about the relationship between a quality control officer and the quality of the widgets that her factory produces. There is a clear sense in which she has no control over the production of the widgets. They are made entirely by other people and machines. Nevertheless, if she is a good quality control officer, we would not say that the absence of faulty widgets coming out of the factory is *lucky* for her. After all, she has the power to remove the faulty widgets, produced by others, from the assembly line before they leave the factory floor. Similarly, we have the ability to withhold commitment to appearances, when we have some reason for doing so. We are not automata who are forced to believe whatever our senses report. When there is reason to be skeptical, we can stop the assembly line and look more closely at the appearance that has prompted our skepticism. If, upon closer inspection, the appearance is faulty, we do not give it our doxastic assent. Insofar as we are doing this kind of epistemic quality control, we are just as responsible for the quality of our perceptual belief output as is the quality control officer for the quality of the widgets produced by the manufacturing plant. Which on my way of looking at things is just another way of saying that it is just as little a matter of luck for us when we have true perceptual beliefs and it is for the quality control officer when her manufacturing plant puts out a quality widget.

## **Conclusion**

Put simply: the control theory of luck is better than the safety theory offered by Pritchard. I have tried to convince you that the safety account is not simply wrong, but fundamentally misguided. That is, I think that it is simply looking in the wrong places for luck. Ironically, my diagnosis here mirrors Pritchard's own diagnosis of control theories like my own. He says that

control theories of luck seem plausible because they mimic in a wide range of cases the results that one gets from the safety theory. I agree, but conclude that this is why the safety theory *appears* plausible. It is the control theory that both gets the cases right and preserves the intuitive conceptual connections among luck, credit, responsibility, and other notions in that family. Consequently, it is the better account of luck to appeal to in an anti-luck theory of knowledge.