

Why epistemologists are so down on their luck¹

Wayne Riggs

© Springer Science+Business Media B.V. 2006

Abstract It is nearly universally acknowledged among epistemologists that a belief, even if true, cannot count as knowledge if it is somehow largely a matter of luck that the person so arrived at the truth. A striking feature of this literature, however, is that while many epistemologists are busy arguing about which particular technical condition most effectively rules out the offensive presence of luck in true believing, almost no one is asking why it matters so much that knowledge be immune from luck in the first place. I argue that the best explanation for the consensus that luck undermines knowledge is that knowledge is, complications aside, credit-worthy true believing. To make this case, I develop both the notions of luck and credit, and sketch a theory of knowledge in those terms. Furthermore, this account also holds promise for being able to solve the “value problem” for knowledge, and it explains why both internal and external conditions are necessary to turn true belief into knowledge.

Keywords Epistemology · Knowledge · Luck · Credit

1 Introductory remarks

Much of the recent literature on knowledge seems to focus on notions like “safety” and “sensitivity.” These are technical notions designed to capture the right sorts of counterfactual conditions under which a true belief should count as knowledge. The usual explanation for why these notions are appropriate conditions for knowledge is

¹ The arguments and ideas in this paper have been in development for a while, and I would like to thank a number of people for their contribution to that development. For many helpful discussions on these topics, I’d like to thank Steve Ellis, Linda Zagzebski and the students in my graduate epistemology seminar—Mary Gwin, Ben Hagy, Matthew Hodge, Robert Johnson and Shyam Patwardhan. And thanks to Karen Antell for her comments on an earlier draft

that each serves to rule out certain classes of “lucky” true belief. It is nearly universally acknowledged among epistemologists that a belief, even if true, cannot count as knowledge if it is somehow largely a matter of luck that the person so arrived at the truth. Thus, conditions like safety and sensitivity are posited to rule out these sorts of lucky instances of true belief.

A striking feature of this literature, however, is that while many epistemologists are busy arguing about which particular technical condition most effectively rules out the offensive presence of luck in true believing, almost no one is asking why it matters so much that knowledge be immune from luck in the first place. After all, the immunity-from-luck requirement is virtually the only thing in the theory of knowledge about which we can claim consensus. Surely this fact alone renders it worthy of explanation. Why do we care so much about the presence of luck in putative instances of knowledge? I think answering this question can tell us more about what knowledge is and why it matters to us than the usual technical collection of necessary and sufficient conditions. This is not to deride the project of providing such conditions for knowledge, but my project here is different. I would like to explain the place knowledge holds in our value system in a way that will both help us understand the phenomenon we are trying to capture with our technical theories, as well as account for the fact that knowledge is, at least sometimes, something of great value.

So the gauntlet has been thrown down. It is now up to me to give an enlightening account of why we care so much about luck-free true believing. Before settling down to this task, there are several misunderstandings prevalent in discussions about so called “epistemic luck” that I wish to forestall. First of all, luck can enter into a situation in a variety of ways. Luck varies both in degree and in kind. If I narrowly avoid being crushed by a falling safe because of a freakishly strong gust of wind, everything I believe truly afterward is a matter of luck, simply because it is a matter of luck that I am alive to have any beliefs at all!² But it is uncontroversial that luck of this sort does not undermine one’s ability to know things. The tricky part, of course, is figuring out how to distinguish the undermining from the non-undermining sorts.

Philosophers have also pointed out that an anti-luck epistemology is doomed to failure because our ability to accomplish anything, whether it be to cross the street safely or to believe truths rather than falsehoods, is always and inevitably due in part to luck. We frail and causally inept humans are never 100% responsible for anything we “accomplish,” so any theory that requires the total absence of luck for knowledge is a recipe for instant skepticism. Agreed. But rather than obviate the possibility of an anti-luck epistemology, this merely demonstrates the need for a bit of subtlety in one’s view. Knowledge is not incompatible with luck full stop. It is incompatible with luck of certain kinds to a certain degree. Further elucidation of this point will have to wait until more has been said about the nature of luck and the role it plays in undermining knowledge.

Why, then, are we so united in our opinion that knowledge is incompatible with luck of certain kinds to a certain degree? Why care about the presence of luck at all? After all, epistemologists are always saying that what we *really* care about is something to do with having true beliefs and avoiding false ones. If so, what difference does it make whether our proficiency at producing such results is due to luck or not? As I have indicated elsewhere,³ I think the best answer to this question is that luck matters

² Unger (1968), p. 160.

³ See Riggs (1998, 2002a, b).

because knowing is an accomplishment. It is something that one deserves credit for. Analyses of knowledge all attempt to ensure that no putative instance of knowing actually counts as such if the believer was too lucky in coming to believe the truth. This unanimity can be explained by the conceptual connections among a family of concepts that includes credit, responsibility, attribution, and luck. In this paper, I will argue briefly for this claim, though the stronger case for it comes from the explanatory power afforded by the hypothesis that “knowing” is credit-worthy true believing. The latter part of the paper will be devoted to articulating this explanatory power, which includes the ability to explain the role of luck in the theory of knowledge, the value that knowledge has over and beyond the value of mere true belief, and, perhaps surprisingly, the reason for the persistence of the longstanding debate over internalism and externalism about epistemic justification.

So, the next two sections of the paper will articulate what I mean by “credit” and “luck,” and how I take these to be related, both in general and specifically with regard to knowledge. Following that, I will address a pair of objections to my account of luck due to Duncan Pritchard, which threaten to derail the account of knowledge that I base upon that conception of luck. And the final section will pull together many of the points argued for in previous sections to illustrate the explanatory power alluded to above.

Before turning to these tasks, I want to acknowledge that I am not the first person to propose and defend a theory of knowledge in terms of credit-worthy true believing. Greco (2003) has made a persuasive case for his version of such a theory in his aptly titled paper, “Knowledge as Credit for True Belief.” Although our views are quite similar in many respects, I will be developing and defending my own view independently of the arguments he gives in that paper. I believe that we are working toward similar goals, and I take my defence in this paper of a “credit-theory” of knowledge to be complementary to the arguments he gives, rather than either opposed to or in rivalry with them. Ideally, interested readers will find the combined case for a credit-theory of knowledge all the more persuasive.

2 Credit: moral and otherwise

The task of this section is to defend my claim that theories of knowledge are designed to eliminate luck from putative instances of knowledge because of our intuitions about the credit-undermining effect of luck. This, if true, would at least suggest that we are intuitively committed to the idea that knowing is an accomplishment—the kind of thing that one deserves credit for.

To begin, it might help to consider another area of philosophy that takes luck to be importantly detrimental to a certain kind of accomplishment. I am speaking, of course, of ethics and the “problem” of moral luck. The classic literature on moral luck originates with two papers on the subject, one by Nagel (1976) and the other by Williams (1976). Both were concerned with the fact that the correct moral evaluation of a person could vary, even when nothing within that person’s control varies. This has the implication that our moral standing is not entirely up to us, but is sometimes at least a matter of luck. In particular, it seems that judgments of a person’s moral responsibility for certain outcomes, whether good or bad, are undercut by the presence of luck. To take a famous example, suppose two people drive home drunk and veer onto the sidewalk. In one case, there happens to be a pedestrian on the sidewalk, and she is struck and killed. In the other case, there is no pedestrian, though had there

been she would have been struck and killed. Here it seems that there is no difference between the two drivers that should make a moral difference. It is purely a matter of luck that one encountered a pedestrian and the other did not. But one is responsible for killing someone and the other is not, despite their apparent moral equivalence.

Ethical theorists have been worried about these sorts of cases because of the plausibility of the Kantian intuition that “ought implies can.” If an outcome is due to something beyond our control, we cannot be held accountable for it. When we compare the two drivers, it is clear that the only difference in the two situations is one that is beyond either driver’s control—the presence or absence of a pedestrian on the sidewalk at the precise time and place that the driver lost control. It appears that we cannot judge the two drivers differently since the only difference in their situations was beyond their control.

As anyone knows who has even a passing familiarity with this literature, there are many different responses to this problem. I am not suggesting that one is forced to agree that the two drivers must be evaluated identically. My purpose is simply to point out what motivates philosophers to worry about moral luck in the first place: a concern that the presence of luck undercuts responsibility. To the extent that some event is due to luck, it is not due to anything attributable to me. Indeed, that is what “luck” means in these cases. To say that it was a matter of luck that the pedestrian happened to be in that particular place is not to say that the event was metaphysically “random,” or uncaused or any such thing. It is not to say that the pedestrian did not have clear reasons for being there or didn’t intend to be there. All of this is perfectly compatible with the (un)luckiness of the pedestrian’s presence. All we mean by saying that the pedestrian’s being there was a matter of luck is that her being there was in no way due to the *driver*. So, in cases of moral luck, the luck is significant because it appears to defeat what would normally be an attribution of moral responsibility.

How does this help us make sense of the importance of luck in epistemology? One possible lesson to draw from the consideration of luck in ethics is that luck plays the same role in epistemology. In other words, perhaps we care so much about eliminating luck from cases we are willing to call knowledge precisely because we consider knowing to be something attributable to the agent, in much the same way as acting rightly or wrongly is something attributable to the agent. Of course, I have already tipped my hand and made it clear that this is precisely the lesson that I draw from the analogy between moral luck and the kind of luck that undermines knowledge.

Resting one’s case on an analogy is always risky, and I would like to have a more convincing and direct argument for the claim that knowledge is credit-worthy true belief. But in the absence of such an argument, I will offer two compensatory considerations.

First, if the significance of luck in the analysis of knowledge is not due to the conceptual connections among the notions of luck, responsibility, attributability, and the like, then to what is it due? (A rhetorical question rushes to the aid of an analogy. Perhaps my argumentative strategy is going in the wrong direction! Nevertheless ...) I cannot think of any situation in which we actually care about the degree of luck involved in producing some outcome in which the reason for our concern is something other than our desire to determine whether, and if so to whom, to credit that outcome. Given that unanimity about anything in epistemology cries out for explanation, having one for the unanimity among epistemologists for the debilitating effect of luck on knowledge should surely count in favour of my view—at least until an alternative explanation is proposed.

Second, in the remainder of the paper I will explain in more detail what I mean by “luck” and “credit” and show that an account of knowledge in terms of these notions allows us to better understand some important issues in epistemology.

3 The anatomy of luck and credit

Let’s consider some simple cases in which we would unhesitatingly attribute an event or outcome to luck. (1) Sam wins 1 million dollars in a fair lottery. (2) Fran throws the basketball across the entire court to make a last-second shot, and it goes in. (3) Three people who operate a small department in a corporation are given promotions for the terrific work of the department. This includes Jerry, who is a complete slacker and rides completely on the coattails of his two talented workmates. (4) Barbara is walking through her house and stumbles on a piece of carpet that has just come loose. Alas, uncontrollably she falls against a table holding an expensive and treasured antique lamp. It falls and shatters.

I take it to be uncontroversial that luck plays a large role in the outcomes in each of these four scenarios. Sam’s case of winning the lottery is probably the purest case of luck there is. In Fran’s case, it is important to keep in mind that nobody is able consistently, or even more than rarely, to make a basketball shot from across the entire court. It simply is not within normal human abilities to do such a thing. Thus, it was almost entirely due to luck that Fran made the shot. If you need more convincing, imagine that Fran is not even a basketball player, but an enthusiastic (but untalented) fan who rushed onto the court after the game and throws the ball in her exuberance. And Jerry’s case is definitely one of undeserved reward. But more importantly, the promotion was undeserved because it was awarded for a performance that Jerry had virtually nothing to do with. Jerry simply lucked out that he had such talented co-workers. (More on Barbara and the final example in a moment.)

There are several points I wish to draw from these examples. The first is that what we mean by “luck” in all these cases is that the agents in question did not bring about the events in question by their own agency. These events were out of the agents’ control. We don’t mean that the events were uncaused, undetermined or even necessarily random in any important sense. Even the lottery case, which we might suppose to be one that has a random element to it, can be changed to make this point. We might imagine that, unbeknownst to Sam, the lottery was fixed by a crime syndicate boss. The boss made the winning number the birthday of his girlfriend. In this case, it isn’t even random which number is picked, but it is just as lucky for Sam that his number was a winner. So it is not the element of randomness in a lottery situation that makes winning it lucky for Sam—it’s the fact that the event (Sam’s ticket winning) is not Sam’s doing.

The same holds true for the other three cases. Since no one has the ability to consistently make a full-court basketball shot, it is not something Fran has control over. It is not something she can *do*, even if she throws the ball and it goes in. That’s why we call it a “lucky shot.” Barbara, by hypothesis, cannot control her fall and thus the shattering of the lamp. And as I have already explained, Jerry is lucky that his co-workers are so good at their job and that his boss does not realize what a slacker he is. Again, we can modify the case to see that it is the lack of control or agency doing the work. Suppose that Jerry is quite crafty, and always looks busy when the boss is around. He also manages to claim credit for work that he does not do, and convinces

the boss that his co-workers' claims to the contrary are the result of envy. Perhaps he even manipulated himself into this department because he knew that his co-workers would do good work that he could claim. Would we still say that Jerry's promotion is "lucky?" Undeserved, yes, but not lucky. His promotion in this case is due to his own efforts and cunning, though not due to his work ethic or competence at the job itself.

The second point to draw from these examples is that luck comes in degrees—it is not an all-or-nothing concept. This is especially clear in the basketball case. As we change the example, bringing the shot closer and closer to the basket, the degree of luck involved when someone makes the shot diminishes. Although there is always *some* element of luck in any basketball shot, there is no question that the degree of luck it is appropriate to attribute to a successful shot varies with the difficulty of the shot for the individual.

The third point is the obvious one that luck comes in both good and bad varieties. In other words, "luck" is itself value-neutral, and needn't imply that the event in question is either good or bad for anybody. It was a matter of luck for Sam that he won the lottery, but it was also a matter of (bad) luck that Barbara broke her favourite lamp. This goes slightly against common usage, since when we talk about "luck" *simpliciter* we tend to mean good luck, as in "he is a lucky guy." But it is also common to use it value-neutrally, as in "if it weren't for bad luck, I'd have no luck at all." In this paper, I will use "luck" value-neutrally.

The final point I wish to bring out with these examples is that there are two importantly different ways in which an event can be lucky for us. It comes down to two importantly different ways in which an event can fail to be something we "do." The first three cases highlight situations in which the event is not due to the agents' control because their actions simply have no significant effect on the outcome of those actions. Merely buying a lottery ticket is not a causally efficacious way to win the lottery (though it's better than not buying a ticket!). Nor is throwing the basketball the length of the court an efficacious way to make a basket (though, once again, it's better than not throwing it at all). Note that merely being in the causal chain that leads to the event, even being a necessary element of the causal chain leading to the event, is not sufficient to make an *agent* causally efficacious in bringing something about. So one way for an event to be lucky for agent *A* is for the agent to be causally inefficacious in bringing about the event.

But there is another way that an event can be a matter of luck for someone, and it is illustrated by the fourth example. In this example, Barbara is unlucky enough to stumble and break her lamp. Here, the main reason this event is unlucky for Barbara is *not* that stumbling into a lamp is not a causally efficacious way to break it, but rather that her stumbling (which caused the breaking) was inadvertent. She did not *mean* to stumble. She had neither the intention nor the desire either to stumble or to break the lamp. To make it plain that it is the inadvertence that is doing the work here, we can again change the example. Suppose that Barbara finds herself through no fault of her own in a room full of delicately balanced, expensive antiques. There is very little room to move, and any misstep will be very likely to cause something valuable to crash to the concrete floor. In this case, it is fairly clear that stumbling is a highly efficacious way to bring about something's destruction. But if Barbara stumbles in this scenario, it is still bad luck when something breaks (so long as she is being sufficiently careful).

So in the realm of action, to say that an outcome or event is "lucky" for someone is to say that it was, to some important degree, out of his or her control—it is not something that the agent brings about. But it is also important that there are two

significantly different ways for an outcome to be lucky. It can be a matter of luck because you, the agent, were insufficiently causally efficacious in either of two senses: (1) the outcome was not causally due to the agent's *abilities*, or (2) the outcome was not intended by the agent.⁴ The basketball example shows one way that these two kinds of luck can come apart, and the example of the room full of antiques shows another.

I have argued for this distinction elsewhere, so I do not want to belabour the point.⁵ But this distinction is important for two related reasons. First, in order for an agent to be fully responsible for some positive event, and thus deserve full positive credit, her connection to the event must not be undermined by the presence of either kind of luck just distinguished. With respect to her, the event must be neither causally indifferent to her abilities nor inadvertent. Interestingly, there is an asymmetry here between what is necessary for positive credit and what is necessary for negative credit. One can be held fully blameworthy (negative credit) for something one did not intend, if one's inadvertence was due to something like culpable ignorance. But one can never be granted full positive credit for bringing about some good end if one was not aiming to bring about that end. Fortunately, I need only consider the simpler case—that of the conditions for positive credit. Since my claim is that knowing is credit-worthy true believing, the credit involved will always be positive credit.

Second, there is an epistemic analogue to both these kinds of luck, and so any definition of knowing in terms of luck must take both into account. Here is a rough sketch of what I have in mind:

S knows that *p* iff:

- (1) *S* believes *p*,
- (2) *p* is true,
- (3) *S* is sufficiently deserving of credit for the fact that she has come to hold a true belief in this instance.

Clause (3) must be spelled out in terms of avoiding both the kinds of luck mentioned above. So,

S is sufficiently deserving of credit for the fact that she has come to hold a true belief in this instance iff

- (a) *S*'s coming to hold a true belief in this instance is the product of *S*'s actual abilities and
- (b) *S*'s coming to hold a true belief in this instance is not inadvertent.

Putting these two definitions together, we get:

S knows that *p* iff:

- (1) *S* believes *p*,
- (2) *p* is true,
- (3a) *S*'s coming to hold a true belief in this instance is the product of *S*'s actual abilities and
- (3b) *S*'s coming to hold a true belief in this instance is not inadvertent.

⁴ I don't mean to suggest here that the agent must have formed a conscious intention to do *X* in order for her to be creditable with *X*. But if there is nothing even remotely intentional involved in bringing about the event (e.g., an intention, desire, practical reason), then the event is lucky in an important respect with respect to the agent.

⁵ See Riggs (1998).

Again, roughly speaking, in order for an agent's true belief that p to meet condition (3a), it must be the case that forming true beliefs of the sort that p is something that S is ordinarily able to do. One way to fill this out more specifically is to say that condition (3a) is met when and only when S 's belief that p is the product of a reliable belief-forming process. Since I want to keep my analysis of knowing at a general enough level to be compatible with a variety of more specific theories of knowledge, I will refrain from committing myself to any more specific formulation of condition (3a).

Condition (3b) is a little trickier to explain. When is S 's coming to hold a true belief in this instance inadvertent? It is not necessary that S have the conscious intention of forming a true belief when coming to believe p . Ordinary folks rarely, if ever, meet such a strange condition. But something like a "desire for one's beliefs to be true rather than false" must be sufficiently operative in regulating S 's doxastic performance in the production of p for S to be credited with knowing, rather than merely believing p . The classic example of a failure of this sort of regulatory influence is wishful thinking. When someone believes something mainly because he wishes it were true, he does not have the intention of believing something false. Nor does he have the intention of believing something because he wishes it were true. What has happened is that the psychological mechanism of wishful thinking has short-circuited his normal belief-forming mechanisms, and so has screened off the influence of his desire for truth. He still has the desire, but it is not exerting its usual causal regulatory influence in the production of the belief that p . If the belief happens to be true, he is due no credit for it because his believing something true in this case was inadvertent.

4 Whose luck, which credit?

Before going on to argue for the advantages of construing knowing as true belief for which one deserves credit, I want to address an objection that strikes at the very heart of my project. As I have made clear, there are deep conceptual connections between the notion of credit, as I am using it, and the notion of luck. I have argued that our distaste for lucky true belief is what motivates our contemporary analyses of knowledge, and in particular motivates the analysis that I prefer. If my account of luck is wrong in a substantive way, then the analysis of knowledge which I motivate in terms of that notion of luck is similarly cast into doubt. In my discussion and my examples, I have consistently treated "lucky" events as those events that are outside the control of the agent for whom they are lucky. And the intuitive connection between luck and credit depends upon this feature of luck. But Pritchard (2005), in one of the very few philosophical discussions of luck in the literature, considers a conception of luck in terms of this sort of lack of control, and concludes that it will not do. He argues that an adequate account of luck need not impose a "lack of control" condition at all. He prefers a modal conception of luck according to which

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

and

- (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). (Pritchard, 2005, pp. 121–125).

Both (L1) and (L2) express only necessary conditions for luck, but while he admits these conditions are somewhat vague, Pritchard seems to be committed to taking their conjunction to amount to a set of necessary and sufficient conditions for an event's being lucky. I want to defend my "lack of control" thesis against Pritchard's objections, as well as point out a few problems with his own modal analysis of luck.

Pritchard gives two arguments against the "lack of control" account of luck. The first reiterates a point made by Latus (2000, p. 167) that, while the Sun's coming up this morning was an event that was out of my control, we would not ordinarily consider its rising to be *lucky*. Thus, lack of control over an event can at best be a necessary condition that allows in all kinds of events that are not intuitively counted as lucky. But he denies lack of control even this role in an analysis of luck, arguing instead that his modal account captures all the intuitive force of the lack of control provision, without being subject to its flaws.

Pritchard's second objection addresses *epistemic* luck directly, by considering the status of beliefs, and in particular simple perceptual beliefs. We do not have any sort of "direct control" over these beliefs, and so it would seem that our having these beliefs is always out of our control, hence lucky. But if that is the case, on my view one would never have perceptual knowledge, because the having of such beliefs would always be a matter of luck for me, and hence never creditable to me. I will address these objections in order.

The first thing to ask about Pritchard's appeal to Latus's example is whether we do, in fact, ever consider the rising of the Sun to be lucky. I think there are cases where we do, but I must admit that as a general rule, we do not think of the Sun's rising in the morning as a matter of luck. Given that the rising of the Sun is beyond my control, I must also admit that my lack of control over event *E* is not a sufficient condition for the occurrence of *E* being a matter of luck for me. But neither of Pritchard's two conditions (L1) and (L2), is sufficient by itself for luck either. The question is, why does Pritchard abandon lack of control as one of the necessary conditions that are jointly sufficient for luck? The closest he comes to giving an answer to this is in the following passage:

A further motivation for employing (L1) as a condition on luck is that it can explain why ... lack of control [is] closely related to, but not essential to, luck. This is because if one has control over a certain event, such that one is able to (typically) determine that a certain outcome obtains, then that is naturally understood as implying that in a wide class of relevant near-by possible worlds that outcome is realized and therefore not lucky (just as (L1) would predict). (Pritchard, 2005, p. 123)

Here, Pritchard seems to argue that his condition (L1), which requires that the lucky event be absent in a wide class of relevant near-by possible worlds, will do a better job of ruling in and out the right cases than a lack of control condition. Furthermore, Pritchard seems to believe that we think that a lack of control is somehow constitutive of the luck relation between an individual and an event because we are conflating "having control" with "high likelihood of occurrence." It would then be natural to

conflate “lack of control” with “unlikelihood of occurrence.” Thus, we come to think that lack of control is necessary to luck, when really it is unlikelihood that matters.

But a brief consideration of some examples suffices to show that Pritchard’s two conditions on luck are not themselves sufficient, and moreover that it is precisely the lack of control condition that is missing. Indeed, the lesson of these examples is that Pritchard has gotten the significance of unlikelihood and control precisely backward. In an admirably clear paper Coffman (2007, DOI: 10.1007/s11229-006-9046-8) offers an example that purports to show that an account of luck that fails to incorporate a “lack of control” clause is insufficient. In fact, it is an entire class of examples.

Suppose that, at time t , S performs a morally significant action such that there was just before t a large chance that S would not perform that action at t . For concreteness, suppose that S chooses at t to make a large donation to Oxfam, where there was just before t a large chance that S would not so choose at t . [*Pritchard’s account of luck*] entails that S is lucky with respect to her choosing at t to make a donation to Oxfam [...] More generally, [*Pritchard’s view*] entails that *any* morally significant action such that there was a large chance its agent would not perform it is such that its agent is lucky with respect to its occurrence. (Coffman, 2007, DOI: 10.1007/s11229-006-9046-8)

This looks like a class of counterexamples to Pritchard’s view, if we agree that an agent who chooses on the spur of the moment, against her ordinary inclinations, to do a morally significant action is not, in fact, lucky to have done so. While I am inclined to agree with this diagnosis, I think that Pritchard has a response that threatens to pull the teeth of this purported set of counterexamples. Recall that Pritchard’s principle (L1) states that a lucky event is one that occurs in the real world, but “does not occur in wide class of the nearest possible worlds *where the relevant initial conditions for that event are the same as in the actual world*” (*emphasis added*). It seems that Pritchard has only to insist that the circumstances, whatever they were, that led our reluctant philanthropist to donate the money are part of the “relevant initial conditions” that must be the same in the nearest possible worlds in which the lucky event must not occur. But if the prompting events are present, then the action will be present. So, the event will not count as lucky on Pritchard’s account, thus saving his view from the purported counterexamples.

Yet we cannot let Pritchard off so easily. After all, why should we grant that the relevant initial conditions include the prompting events of the donation? It seems that Pritchard owes us some account of how we determine the relevant initial conditions. If it were not for such purported counterexamples, perhaps we could make do with a rough and intuitive account of these conditions, but here some precision and objectivity are called for. Pritchard acknowledges this shortcoming of principle (L1), and is uncomfortable relying on bare intuition to provide us with the specification of which conditions are to be preserved in the cross-world comparisons. But, he concludes that the problem is not fatal, because the “relevant initial conditions” clause can be jettisoned if need be.

There are, of course, problems with this partial specification of luck, one of which is the inherent vagueness involved in the demand that the relevant initial conditions of the event should be the same in all the near-by possible worlds under consideration [...] One way around this problem could be to drop this clause and simply consider the (unrestricted) class of near-by possible worlds, on the grounds that this class of worlds will tend to be dominated by worlds

in which the ‘relevant initial conditions’ as we intuitively understand them are the same [...] Given this possible escape route, the objection is no longer fatal. (Pritchard, 2005, pp. 124–125).

But if this clause is jettisoned, then Pritchard has no reply to Coffman’s example.

But it is instructive to consider how we would intuitively determine which initial conditions are relevant to the determination of the “luckiness” of an event for an agent. To do this, I think we need a slightly different example. Imagine a young basketball player who has tremendous natural and developed skills and displays them proficiently every day in practice. Unfortunately, he also has a terrible fear of failure which causes him to “choke” when he is actually playing a game against an opposing team. A typical performance during a game would be for him to take, say, 20 shots and miss all 20. And this is due to his fear interfering with his ability to deploy his impressive skills. But every once in a while, for no reason the player has ever been able to determine, he finds himself confident and calm for a moment or two during a game. One night he happens to have the ball in his hands when this occurs, and he shoots the ball. Absent his usually crippling fear, he makes a skilful shot, which goes in.

Are we tempted to say that this was a “lucky shot?” I don’t think so. If he had taken the shot under the conditions of his typical fearfulness, and it had gone in, we would be tempted to call it a lucky shot. But in this case, the ball went in by virtue of his skills and abilities. Thus, we give him credit for making the shot, rather than calling it lucky. This looks like another case in which Pritchard’s principle (L1) goes wrong. In a large class of near-by possible worlds (now unrestricted), our player misses the shot because he doesn’t happen to have his moment of calm right then. If there were some principled, non-question-begging way to determine that the relevant initial conditions that had to be the same across the possible worlds we are comparing included whatever produced the calm confidence that the player experienced, then Pritchard might still have a way out. But in the example we are not even given an account of what causes the player’s occasional periods of calm confidence. So it’s very hard to imagine how we could be making any kind of intuitive judgment about whether that cause, whatever it is, is among the relevant initial conditions of the event. Yet I submit that we are in no similar uncertainty about whether the event was lucky. Our judgment about the luckiness of the event is conditioned on our judgment that the player was in control of the event, not on some nuanced understanding of what conditions are or are not included in a cross-world comparison.

Now this does not mean that a principle like (L1) is not still necessary in an adequate analysis of luck. It simply shows that such a principle, even in conjunction with (L2), is not sufficient. What is needed is a third principle:

- (L3) If an event E is lucky for S , then E was either:
- (1) not the product of S ’s actual abilities, or
 - (2) inadvertent with respect to S .

Notice that this principle incorporates the elements of the account of epistemic credit offered previously. Principle (L3) says, in short, that if E is lucky for S , then S does not deserve credit for E . This is precisely the relationship I think holds between luck and credit in general.

If we add this principle to Pritchard’s first two, then we have a more adequate analysis of luck. In short, Pritchard was wrong to have abandoned the lack of control

condition on luck. But if we add (L3), why do we still need (L1)? It seemed as if they were rivals to be decided between, not independently necessary conditions. But the initial objection raised by Pritchard shows why we must keep something like (L1). Even though the Sun rising this morning was not under my control, it was nonetheless not a matter of luck with respect to me. Pritchard is right to point out that lucky events must be modally remarkable, though I disagree with him on the details of what (L1) should look like. I think that the proper condition (L1) should be capturing is actual probabilistic unlikelihood, rather than his more sophisticated condition involving “close” possible worlds. But that is beyond the scope of this project. For now, I will rest content with having shown that Pritchard’s account of luck is unsuccessful without the addition of a principle that incorporates the lack of control on the part of the lucky agent.

Pritchard’s second objection to the “lack of control” thesis about luck is that we seem to lack precisely the kind of control over our beliefs that would be necessary for them to fail to be lucky with respect to us. He takes perceptual beliefs to be particularly problematic, because we seem to get them “willy nilly,” without any sort of control at all. I recognize that this is one of the most difficult objections facing an account of knowledge in terms of credit. Responding adequately to it would require at least a paper of its own. However, I will offer two brief responses to indicate how I think such a response should go.

First, it is important to realize that perceptual believings are not as passive as epistemologists sometimes try to make out. Sosa (1991) has pointed out that we do not automatically believe everything that our senses report. When there is a cue in our environment that something may be abnormal about our current perceptual situation, we do not take our sensations at face value. Rather, we check to see what the abnormality is, and we compensate as best we can for it in what we believe. Thus, even knowledge that flows more or less directly from our faculties of perception, memory or introspection is still what Sosa calls “reflective knowledge,” rather than a more brutish, non-reflective kind he aptly calls “animal knowledge.”

Note that no human blessed with reason has merely animal knowledge of the sort obtainable by beasts. For even when perceptual belief derives as directly as it ever does from sensory stimuli, it is still relevant that one has *not* perceived the signs of contrary testimony. A reason-endowed being automatically monitors his background information and his sensory input for contrary evidence and automatically opts for the most coherent hypothesis even when he responds most directly to sensory stimuli. For even when response to stimuli is most direct, *if* one were also to hear or see the signs of credible contrary testimony that would change one’s response. The beliefs of a *rational* animal hence would seem never to issue from *unaided* introspection, memory or perception. For reason is always at least a silent partner on the watch for other relevant data, a silent partner whose very *silence* is a contributing cause of the belief outcome. (Sosa, 1991, p. 240, *emphasis in original*)

Sosa is arguing that beliefs resulting from perception, memory or introspection are categorically the product of reason, though there are differences of degree between these beliefs and more deliberative beliefs in terms of how large a role reason plays in their production. But even if we grant this, Pritchard may still contend that beliefs in general are not the sorts of things that one can have control over. Thus, the more

general issue of “voluntarism” is still to be addressed, which leads me to my second brief response to this line of objection.

It is important to understand the kind of control that is required for me to be able to make my case. It is not necessary that we have so-called “direct control” over our beliefs in order for my account of knowledge to be successful. It is a commonplace that we cannot choose to believe that p in the same way that we can choose to raise our arm or close our eyes. But such direct control over events is not necessary for one to deserve credit for them. All that is necessary is that one has the kind of control over one’s beliefs that one typically has over other kinds of happenings that we rightly attribute to individuals. For example, we might attribute the success of a debate team to the hard work and brilliance of its coach, and credit her for that success. Yet having a successful debating year is not something the coach has direct control over.

I take the problem of “doxastic voluntarism” to be one of the most important objections to the analysis of knowledge in terms of credit that I offer here, and it deserves a paper unto itself. Although I cannot address this issue as effectively as I would like here, I will at least point out that a number of philosophers writing on this subject have argued that we do have the kind of control over our beliefs necessary to make them something for which we are rightly attributed credit. Interestingly, these philosophers come at this issue from a variety of perspectives, some from moral psychology, some from the free will debate, some from more traditional philosophy of mind, and some from mainstream epistemology.⁶ I will have to settle for a promissory note that sense can be made of our having the requisite control over our beliefs necessary for the analysis of knowledge I provide. The promise, however, is at least backed by the support of the influential philosophers who have made similar cases recently.

I conclude that Pritchard has not succeeded in making the case that we can leave out a “lack of control” condition from our account of luck, at least as a necessary condition. This is crucial for my account of knowledge, because I have argued that knowing that p depends upon one’s having control (of certain specified kinds) over the fact that one came to hold a true belief in the circumstances, and this part of the account of knowing explains the widely held intuition that knowing is incompatible with (certain kinds of) luck. If an event’s being lucky for S did not imply that S lacks control over E , then this explanatory function of my account of knowing would fail.

5 The payoff

Having laid out in somewhat more detail what luck and credit amount to, and how an analysis of knowledge can be given in terms of these notions, I want to point out some of the advantages of such an analysis. To begin with, I have already shown that this analysis can account for the significance of luck in putative instances of knowledge. As far as I know, no other account of knowledge has a good answer to the question, “Why must knowledge be (sufficiently) immune from luck?” unless it concedes, either implicitly or explicitly, that knowing is an achievement for which the knower deserves credit.

There are two other major problems in the theory of knowledge that this analysis sheds light on. First, there is what has been variously called the “value problem,” the

⁶ For a representative sample, see Hieronymi (2006), Raz (1999), Adler (2002), Owens (2000) and Audi (2001).

“Meno problem,” and the “swamping problem.”⁷ All of these amount to the problem of accounting for the value that “knowing that p ” has over and above the value that “believing- p -when- p -is-true” has. Several epistemologists have addressed this problem recently,⁸ mostly with an eye towards showing that certain theories of knowledge (usually reliabilism) cannot solve the problem. But others have argued that a view much like my own *can* solve the problem. Greco (2003) defends an account of knowledge that, like my own, is given in terms of credit-worthy true belief, and he claims that this view can solve the value problem. And I have argued that a version of reliabilism, suitably revised to reflect the notion of credit, can also make some progress toward solving the value problem.⁹

The reason that credit-worthiness views of knowledge can solve the value problem is that they introduce a new vector of value: credit. Traditional reliabilist theories have trouble accounting for the extra value knowledge has over true belief because the only elements of the theory are true belief and reliability. But the value of reliability derives entirely from the product that is reliably produced. Thus, any value that reliability might have is “swamped”¹⁰ by the value inherent in any true belief. A true belief is not made more valuable by being produced reliably. But if knowing that p always entails that one deserves credit for having achieved a true belief, then this introduces something besides true belief that is valuable. Even though deserving such credit requires that a belief be produced by a reliable process, deserving credit for the thing so produced is a further fact that confers additional value on the state of affairs of which it is a component.¹¹

This point raises obvious and important questions about the nature and source of the value of credit, and these are questions that the credit-worthiness account of knowledge should address and, hopefully, satisfactorily answer. But for now, I will have to make do with the intuitive pull of a representative pair of examples.

Scenario 1. Suppose that a small child has fallen into a raging river, and will soon drown if not rescued. I leap into the dangerous waters, swim to the child and bring her safely to the shore.

Scenario 2. Once again, a small child is in danger of drowning in the river. I leap into the dangerous waters, but I’m unable to reach the child, despite exerting just as much effort and displaying just as much courage and skill as in Scenario 1. It simply is not enough in this scenario to allow me to reach the child. However, in leaping into the water, I dislodge a fallen branch which falls in to the water. The swirling current carries it to the child, who is thereby able to stay afloat until she is carried ashore further downstream.

In both scenarios, the good outcome comes to pass—the child survives. But I would prefer Scenarios 1 to 2, because in Scenario 1 the good outcome is attributable to me. Of course, the difference in value between the two scenarios pales beside the value of the child’s survival, but the difference is nonetheless there.

⁷ See Kvanvig (2004) and Zagzebski (2000)

⁸ See Kvanvig (1998, 2004), Jones (1997) and Zagzebski (2000)

⁹ Riggs (2002b)

¹⁰ Kvanvig (2004)

¹¹ For a defence of this claim, see Riggs (2002b). For a more extended defence of a similar claim, see Sosa (2003).

Besides offering hope of solving the value problem, the credit-worthiness account of knowledge also explains both the source of the internalist/externalist controversy and the fact that neither side has ever achieved victory. The controversy arises because, while proponents of each side agree that “knowing” must preclude lucky true belief in some way, they focus on different *kinds* of luck that undermine knowledge.¹² Until it is understood that an adequate account of knowledge must preclude both kinds of knowledge-undermining luck, these approaches will be seen as rivals rather than complementary elements of a single account of knowledge.

But if we see that knowing is an achievement, and that for *S* to achieve *X* is for *S*'s abilities and intentions to be sufficiently credit-worthy for bringing *X* about, then it is clear why both internalist and an externalist conditions on whatever is needed for knowledge in addition to true belief are necessary. An agent cannot be said to be credit-worthy for having a particular true belief, *p*, unless the process, method, habit, trait or what have you that produced it is objectively reliable. That is because objective reliability is a mark of having an ability. This may need to be suitably indexed to an environment, as Sosa (1991) does on his version of reliabilism, but the fact remains that it is an externalist, objective condition that must be met for an agent to be sufficiently credit-worthy for her belief.

Furthermore, the credit-worthiness account of knowing explains why *S*'s desire or intention to have true beliefs must also be operative in the production of the true belief as well. This ensures that *S* is employing her own best epistemic standards in coming to believe *p*—that she is being epistemically responsible. While meeting one's own best epistemic standards (or being epistemically responsible) is not a universally accepted condition of internalist justification, it is at least a common one. So the credit-worthiness account of knowing demands both external and internal conditions to be met for an agent to be sufficiently credit-worthy to be said to know that *p*, and does so in a way that ties those conditions to what makes knowing more valuable to us than merely being right.

So the credit-worthiness account of knowing has the advantages of explaining the relevance of luck to knowledge claims, of showing promise in solving the value problem, and of explaining the need for both internalist and externalist conditions on whatever turns true believing into knowing. I think these are powerful advantages and warrant taking seriously this view of knowing. Admittedly, it is an analysis at a very high level of abstraction. As I have pointed out at several points, this rubric for knowing is compatible with a variety of different specific theories of knowledge. Yet, at the same time, it does impose a substantive constraint on such theories.

References

- Adler, J. (2002). *Belief's own ethics*. Cambridge, MA: MIT Press.
- Audi, R. (2001). Doxastic voluntarism and the ethics of belief. In M. Steup (Ed.), *Knowledge, truth, and duty* (pp. 93–111). Oxford: Oxford University Press.
- Coffman, E. J. (2007). Thinking about luck and making it pay off. *Synthese*. DOI: 10.1007/s11229-006-9046-8 (this issue).
- Greco, J. (2003). Knowledge as credit for true belief. In: M. DePaul, & L. Zagzebski (Eds.), *Intellectual virtue: Perspective from ethics and epistemology* (pp. 111–134). Oxford: Oxford University Press.
- Hieronymi, P. (2006). Controlling attitudes. *Pacific Philosophical Quarterly*, 87(1), 45–74
- Jones, W. (1997). Why do we value knowledge? *American Philosophical Quarterly*, 34, 423–439
- Kvanvig, J. (1998). Why should inquiring minds want to know? *The Monist*, 81, 426–451

¹² See Riggs (1998) for a fuller defence of this claim.

- Kvanvig, J. (2004). *The value of knowledge and the pursuit of understanding*. Cambridge: Cambridge University Press.
- Latus, A. (2000). Moral and epistemic luck. *Journal of Philosophical Research*, 25, 149–172
- Nagel, T. (1976). Moral luck. *Proceedings of the Aristotelian Society*, 76, 136–150
- Owens, D. (2000). *Reason without freedom*. London: Routledge
- Pritchard, D. (2005). *Epistemic luck*. Oxford: Oxford University Press.
- Raz, J. (1999). *Engaging reason*. Oxford: Oxford University Press.
- Riggs, W. (1998). What are the chances of being justified? *The Monist*, 81, 452–472
- Riggs, W. (2002a). Beyond truth and falsehood: the real value of knowing that *P*. *Philosophical Studies*, 107, 87–108
- Riggs, W. (2002b). Reliability and the value of knowledge. *Philosophy and Phenomenological Research*, 64, 79–96
- Sosa, E. (1991). *Knowledge in perspective*. Cambridge: Cambridge University Press.
- Sosa, E. (2003). The place of truth in epistemology. In M. DePaul, & L. Zagzebski (Eds.), *Intellectual Virtue: Perspective from ethics and epistemology* (pp. 155–179). Oxford: Oxford University Press.
- Unger, P. (1968). An analysis of factual knowledge. *The Journal of Philosophy*, 65, 157–170
- Williams, B. (1976). Moral luck. *Proceedings of the Aristotelian Society*, 76, 115–135
- Zagzebski, L. (2000). From reliabilism to virtue epistemology. In G. Axtell (Ed.), *Knowledge, Belief and Character* (pp. 113–122). Lanham, MD: Rowman & Littlefield.