When Is a Belief True Because of Luck?

Greene, Preston

2013

Greene, P. (2013). When Is A Belief True Because Of Luck? The Philosophical Quarterly, 63(252), 465-475.

<https://hdl.handle.net/10356/79720>

<https://doi.org/10.1111/1467-9213.12037>

© 2013 The Author The Philosophical Quarterly. This is the author created version of a work that has been peer reviewed and accepted for publication by The Philosophical Quarterly, The Author The Philosophical Quarterly. It incorporates referee's comments but changes resulting from the publishing process, such as copyediting, structural formatting, may not be reflected in this document. The published version is available at: [<http://dx.doi.org/10.1111/1467-9213.12037>].

Downloaded on 26 Apr 2023 20:04:44 SGT

When Is a Belief True Because of Luck?

Preston Greene, Rutgers University

Forthcoming in *The Philosophical Quarterly*

PENULTIMATE DRAFT: Please cite from the published draft.

Abstract

Many epistemologists are attracted to the claim that knowledge possession excludes luck. Virtue epistemologists attempt to clarify this idea by holding that knowledge requires *apt belief*: belief that is true because of an agent's epistemic virtues, and not because of luck. Thinking about aptness may have the potential to make progress on important questions in epistemology, but first we must possess an adequate account of when a belief is true because of luck. Existing treatments of aptness assume a simple and natural view of luck attribution, according to which the success of a performance is attributable to luck if one of the principal causes of the success is a lucky event. I show that this view is false, and should be replaced. This has major implications for virtue-theoretic accounts of knowledge, as well as the role of luck in epistemology more generally.

Keywords: knowledge; virtue epistemology; luck; Gettier; Sosa; Greco

I. APTNESS AND THE SIMPLE VIEW OF LUCK

Knowledge is a type of cognitive achievement, and the influence of luck is incompatible with genuine achievements. Thus, if an agent knows something, it is not the case that she was merely lucky to believe the truth. We might call this thought the 'epistemic luck platitude'.¹ Many believe the epistemic luck platitude provides a crucial insight into the nature of justification and knowledge. Among other things, it is claimed that it can help us better understand the main lesson of the Gettier problem. Jonathan Dancy, for example, tells us: 'Justification and knowledge must somehow not depend on coincidence or luck. This was just the point of the Gettier counter-examples; nothing in the tripartite definition excluded knowledge by luck'.²

¹ D. Pritchard, *Epistemic Luck* (Oxford UP, 2005), p. 1.

² J. Dancy, *Introduction to Contemporary Epistemology* (Oxford: Basil Blackwell, 1985), p. 134.

It is no surprise then that much contemporary work has been devoted to specifying the conditions under which an agent gets lucky, epistemically speaking. Thanks in large part to the work of Ernest Sosa, *virtue epistemology* offers one of the most promising attempts to create an account of knowledge and justification that is adequately responsive to the idea that knowledge possession excludes luck. A common element of virtue-theoretic accounts is the idea that if an agent knows that p , then the truth of her belief is attributable to the exercise of her epistemic virtues, and not to luck. As early as 1988 Sosa wrote, ‘Knowledge is true belief out of intellectual virtue, belief that turns out right by reason of the virtue and not just by coincidence’,³ and that knowledge requires that ‘one’s belief non-accidentally reflects the truth of p through the exercise of...virtue’ (p. 184). This idea has subsequently been developed by Sosa and several other proponents of virtue epistemology. John Greco writes that a subject S knows that p if and only if S believes the truth because S ’s belief is produced by intellectual ability.⁴ And this idea is then used to explain why Gettier subjects fail to know. Linda Zagzebski, for example, holds that a Gettier agent may form a belief using her epistemic virtues, and the belief may be true, but ‘she does not have the truth *because of* her virtues’.⁵ The importance of specifying that a belief be true ‘because of’ one’s epistemic virtues---or, in other words, that the truth of the belief be ‘attributable to’ one’s epistemic virtues---has become an important element for virtue theories in epistemology.⁶

Sosa takes his view to be located within a more general theory of performance normativity. He illustrates this with an example of an archer shooting at a target.⁷ When the archer aims and shoots, Sosa explains, there are at least three important ways to evaluate the shot. An *accurate* shot is successful in hitting the target. An *adroit* shot is skillful: it manifests the archer’s competence relative to archery. Finally, an *apt* shot is accurate because it is adroit. The accuracy of an apt shot is creditable to the skill of the

3 E. Sosa, ‘Beyond Scepticism, to the Best of our Knowledge’, *Mind*, 97 (1988), pp. 153–88, at p. 175.

4 J. Greco, *Achieving Knowledge* (Cambridge UP, 2010), p. 12.

5 L. Zagzebski, *Virtues of the Mind: An Inquiry into the Nature of Virtue and the Ethical Foundations of Knowledge* (Cambridge UP, 1996), p. 297 (emphasis added).

6 See also J. Greco ‘Knowledge as Credit for True Belief’, in M. DePaul and L. Zagzebski (eds.), *Intellectual Virtue: Perspectives from Ethics and Epistemology* (Oxford UP, 2002), pp. 111–34; J. Greco, ‘Virtue and Luck, Epistemic and Otherwise’, *Metaphilosophy*, 34 (2003), pp. 353–66; W. Riggs, ‘Reliability and the Value of Knowledge’, *Philosophy and Phenomenological Research*, 64 (2002), pp. 79–96; D. Pritchard et al., *The Nature and Value of Knowledge* (Oxford UP, 2010); E. Sosa, *A Virtue Epistemology: Apt Belief and Reflective Knowledge* (Oxford UP, 2007); and E. Sosa, *Knowing Full Well* (Princeton UP, 2011).

7 Sosa, *A Virtue Epistemology*, p. 22 and *Knowing Full Well*, p. 4.

archer. Beliefs, Sosa thinks, can be evaluated in the same way. We can ask whether a belief is accurate (i.e., true); whether a belief is adroit (i.e., epistemically competent); and finally, whether a belief is apt (i.e., true because competent). Apt beliefs are creditable to the skill of the believer, and such credit is required for knowledge.

Sosa provides an influential example of an accurate and adroit shot that fails to be apt, which I'll refer to as the 'double-gust case'. In the double-gust case, an archer's competently shot arrow is first blown off course by a gust of wind. A second gust of wind then blows the arrow back onto its original course. In this case, Sosa argues, due to the intervening gusts the accuracy of the shot is not attributable to the archer's skill, but rather to luck (*A Virtue Epistemology*, p. 23).. As John Greco writes, 'We will not cite [the archer's] skill as the cause of the bull's eye, although clearly a manifestation of skill was involved' (*Achieving Knowledge*, p. 75). The double-gust case is important because it appears to explain what is happening in Gettier cases. A Gettier case involves a subject who forms a true and competent belief that is nevertheless not true because competent, due to a lucky event similar to that of the countervailing gust. Both the shot and the belief, it is claimed, are accurate because of luck. Hence, the use of a notion of aptness, coupled with the double-gust case as an elucidation of Gettier cases, has the potential to offer progress on some of the most important questions in epistemology.

This line of thought assumes a simple and quite natural view about the nature of aptness. The view is that the success of a performance is attributable to luck if one of the principal causes of the success is a lucky event. This supports the following conclusions: the archer may have shot competently, but one of the main causes of the accuracy of the shot is a lucky event (the appearance of the second gust). Therefore, the success of the shot is attributable to luck, rather than to the archer's competence. Similarly, a Gettier agent may have believed competently, but the truth of her belief is ultimately caused by a lucky event, and so the success of the belief is attributable to luck, rather than to the believer's competence.

In section II, I argue that this simple view, as well as John Greco's sophisticated extension of it invoking causal explanatory salience, are false. In Section III, I argue that whether a performance is apt primarily depends on the *basis* for the performance. According to this *basis-relative* view, in order to determine whether a performance is successful because of luck, we need to know on what possibilities the agent bases his performance. Realizing this has some large implications for virtue-theoretic accounts of knowledge, as well as the role of luck in epistemology generally, which I draw out in Section IV.

II. THE POKER CASE

Compare the following case to the double-gust case. Annie Duke is an expert poker player playing against an opponent who has just bet all of his chips. While Annie's hand is weak, by picking up on cues from her opponent's betting pattern and subtle hints suggested by his body language she has concluded that her opponent's hand is even weaker. She chooses not to fold, instead calling his bet. After flipping over the cards, it is revealed that she was right: her opponent's hand is in fact much weaker than her hand. The situation is this: if neither of the next two cards to be revealed is the king of spades, or if either is an ace, Annie wins. This gives her more than a 95% chance of winning the hand. Let us now consider two ways in which the story could develop. In case one, neither of the next two cards is the king of spades. The next card is a two and the last is a three. In this case, we are quick to attribute Annie's success in winning the hand to her expert skill. Whereas a novice would not have picked up on the subtle cues in her opponent's betting pattern and body language and would have folded, Annie called, and this decision manifested an expert competence and led to the winning of the hand. In case two, however, the first card to appear after Annie's decision is the king of spades. 'Oh no', she might cry, 'how unlucky'. Given that the first card is the king of spades, she now can only win the hand if the last card is an ace, which will happen only 7% of the time. But the last card actually does turn out to be an ace, and she jumps for joy. Now it is time for her opponent to cry, 'How unlucky!' In this case, as in the double-gust case, some may be quick to attribute Annie's success to luck. While her decision certainly manifests her competence, it may be claimed that her success in winning the hand does not: the success is not sufficiently attributable to her skill since luck plays an intervening role. But this is a mistake.

There are many ways the hand could play out, and this can be represented by a table of possible card combinations. In our example, after Annie makes her decision, there are only two values left to be decided---the two cards to be revealed---and so for our purposes the table would consist of all the different two-card combinations. Most of the combinations on this list won't include the king of spades {e.g., $(2_C, 3_D)$, $(10_S, Q_D)$ },⁸ but a few will. Of the combinations that do include the king of spades, an even smaller number will include an ace {e.g., (K_S, A_D) , (A_H, K_S) }, and these combinations would be highlighted were we to mark out the subset that are winners for Annie. If we were to

8 The letter or number in normal script is the card's rank, while the letter in subscript is the card's suit.

actually create such a table, we would have thousands of boxes containing the possible combinations, and 95% of which would be highlighted as winning combinations. Visualizing the possibilities in this manner puts us in a better position to understand what it means, exactly, to say that she had a 95% chance of winning after the crucial decision. It means that after deciding to call the bet the situation was such that 95% of the possible two-card combinations would result in her winning. Any combination is as good as the other—Annie wins all the same whether the actual combination turns out to be $(2_C, 3_D)$ or, as in case two, (K_S, A_D) . In light of this, it appears incorrect to say that Annie won because of luck in case two. After all, the number of winning combinations represents a high percentage of the total possibilities partly because combinations like (K_S, A_D) are highlighted. These combinations are in part responsible for the fact that she had a 95% chance of winning, which is a fact that is partly responsible for explaining her manifestation of skill when she decides to call. Given this, it makes no sense to claim that Annie wins due to skill in cases like $(2_C, 3_D)$ but due to luck in cases like (K_S, A_D) .

We can see this even more clearly by imagining a variant of case two in which the ace appears first, thus sealing the victory, before it is revealed that the last card is the king of spades. It is certainly true that the *experience* of watching this hand develop would be very different. But given that there are no more moves either player can make, it is hard to see how the order of the cards could affect whether Annie wins because of luck. A similar point can also be made by imagining that the cards are flipped over simultaneously.

It is therefore a mistake to think that Annie won the hand because of luck. This is important because, if true, it shows that the simple view is false. One of the principal causes of Annie winning the hand is a lucky event: the appearance of the ace.⁹ And yet the winning of the hand is not attributable to luck.

Perhaps at this point one might have the feeling that the poker case is importantly different from the archery case because the random selection of an ace is a far more

9 A popular account takes lucky events to be significant chancy events that lie beyond the agent's control. See E. Coffman, 'Thinking About Luck', *Synthese* (2007), pp. 385–98 and Pritchard, *Epistemic Luck*. But however one understands what a lucky event is, it seems that the ace appearing as the last card is going to qualify as one. Skepticism regarding the plausibility of creating a satisfactory account of lucky events only serves to strengthen the argument of this paper, since I offer an account that dispenses with the notion of a lucky event in thinking about apt and inapt performance (also see footnote 13 below). However, I will continue to assume that there exists an adequate account.

likely event than a gust of wind knocking the arrow back on course. However, this is not a significant difference, because we can imagine *million-card poker* and get the same result. In million-card poker Annie is playing with a million-card deck and again has a 95% chance of winning when she decides to call with two cards to come. The state of play is such that, if a certain type of card appears first, then only a single card, say, the '1,000,000 of spades', can save her from losing. All the arguments that applied in the original poker case also apply here. In million-card poker, the shock of seeing the 1,000,000 of spades appear as the final card will be several times more intense than the shock of seeing the ace in the original poker example. Indeed, this event, given its decreased chance, will be significantly more lucky than in the original example. But, again, given that Annie has the same chance of winning after she calls and no subsequent control over the hand, and given that the 1,000,000 of spades possibility is partly responsible for her manifestation of skill, we should not claim that the winning of the hand is due to luck.

We can perhaps better explain our reaction to the poker case by using one of John Greco's ideas: what is important in determining whether we attribute the success of a performance to luck is the *explanatory salience* of the lucky event versus that of the competent performance.¹⁰ Greco agrees with the simple view that the 'because of' or 'attributable to' relation should be thought of as a causal explanation, and he adds that our causal explanations pick out salient parts of the causal story. He takes competences to have 'default salience' in explaining successes, but holds that such default salience can be easily overridden by abnormal events. In the archery case, Greco takes the presence of the countervailing gusts to be 'deviant', and explains that 'salience goes to what is deviant, and away from what is normal or usual' (p. 75). Indeed, a normal way of hitting the bull's eye involves the arrow flying relatively straight, or perhaps slightly curved if the skilled archer accounts for a steady breeze. The countervailing gusts are abnormal, and, therefore, they have more explanatory salience than the archer's competence. But the poker case follows the same pattern. The usual way for Annie to win the hand would be for the next two cards to be inconsequential, as in case one. But when the unusual sequence presents itself, the influence of the cards gains salience and we feel the urge to attribute the winning of the hand to the appearance of the ace, rather than to Annie's competence.

All that said, the winning of the hand *is* attributable to Annie's competence, even though Greco's theory predicts (perhaps correctly) that we will feel the urge to attribute the winning of the hand to luck due to the salience of the unusual sequence of cards.

¹⁰ Greco, *Achieving Knowledge*, pp. 73--5.

Accordingly, Greco's proposal should be taken to be a descriptive account of what is controlling our sometimes mistaken judgments about luck. It should not be taken to be a normative account, and it certainly should not be taken to show that Annie's performance is attributable to luck. Instead of embracing the workings of causal explanatory salience, we should work to overcome it.

What are we left with if we reject both the simple view of aptness and Greco's more sophisticated view invoking causal explanatory salience? The poker case shows that if certain other conditions are in place, then the fact that a success is caused by a lucky event does not preclude the success from being attributable to a competence. But what are these other conditions? Answering this question will provide us with a new account of performance-aptness to replace the simple view.

III. NON-EPISTEMIC APTNESS

What does seem to be of legitimate importance in thinking about aptness is the relation between the possibilities that obtain and the basis for the agent's performance. A natural way to interpret the archery case is the following: the archer aims his shot directly at the bull's eye assuming that there will be no gusts. There then is a gust, which is corrected by a second gust. If this is how one interprets the archery case, then it is indeed perfectly appropriate to attribute the accuracy of the shot to luck. But the poker case is different. Poker is explicitly a game of probabilities, and being skilled at poker requires that one understand these probabilities. This in turn requires that one anticipate even remote possibilities in deciding how to act. Therefore, in deciding to call, it is reasonable to assume that Annie is perfectly aware that the possibility which obtains is a way that she could win the hand. It is further reasonable to assume that this possibility plays at least an indirect role in her reasoning about whether to call. After all, as noted above, this possibility is partly responsible for explaining why Annie has a 95% chance of winning.

Removing these elements shows their importance. Imagine that Annie decides to call on the basis that, if her opponent's hand is weak, she will, without question, win. The king then appears, and Annie is shocked to see that it is now very unlikely that she will win. She is then saved by the ace, and wins after all. In this variant, it seems perfectly appropriate to attribute the winning of the hand to luck, and this provides nearly decisive evidence against the simple view and in favor of a basis-relative account of aptness.

We can further test for the need for a basis-relative account by considering a variant of Sosa's case in which an archer does in fact base his shot partly on the possibility of countervailing gusts. Imagine that the archer is shooting at a target through a long tunnel with two holes, one on each side, placed at different points along the way. There are fans positioned to blow through the holes, and the archer is informed that there is a 10% chance, for each fan, that the fan will be activated. The archer reasons that shooting straightly (shooting so that the arrow would hit the target with no gusts) is the best approach to take. That way, as long as his shot is well-placed, he has an 82% chance of hitting the target, since 81% of the time there will be no gusts and 1% of the time there will be countervailing gusts. Should we attribute the accuracy of the shot to luck when there are countervailing gusts? It seems not, and, indeed, this situation is almost exactly like the one that Annie faces.

If the above is true, then whether the archer's shot is attributable to luck, rather than his archery competence, does not depend on whether one of the main causes of the shot's accuracy is a lucky event, but rather what role the obtaining possibility---dueling gusts---played in his decision about how to shoot. Here we can see a legitimate role for salience to play in a theory of aptness. Rather than influencing causal explanations, as in Greco's account, recognizing what is salient when an agent is deciding how to perform is important for identifying the possibilities on which we are likely to think that a performance is based. In tunnel archery, the placement of the tunnels and fans makes the double-gust possibility very salient, and so it is very likely that the archer would consider that possibility in aiming his shot. But in normal archery, the double-gust possibility is not salient, and so we are less likely to think that the shot is partly based on that possibility.

Does this mean that the accuracy of the shot could never be attributable to luck, no matter how bizarre the influence on the arrow, so long as it is shot competently and the archer partly bases his shot on the possibility of that influence? No, for there are clearly limits to what being skilled in archery entails. If a dog intercepts the arrow and runs around with it for a while before depositing it on the target, it would be incorrect to attribute the accuracy of the shot to competence in archery.¹¹ But this is not because the event is incredibly lucky; rather, it is because dealing with the interference of dogs is not relevant to archery competence. While archers can be expected to understand and anticipate the influence of wind, they are not expected to anticipate the influence of dogs.¹² But we can indeed imagine a broader competence that mixes both archery and

11 The example is due to Pritchard et al., *The Nature and Value of Knowledge*, p. 28.

12 Another way to get a grasp on this is to think about which sorts of influence require shots to be retaken in a competition, and which do not.

dog handling abilities, which may be manifested in this instance. If the competition requires the archer to both shoot accurately and train a dog to intercept the arrow and place it on the bull's eye, then the success of such a performance may indeed be attributable to this sort of broader competence.

Accepting this still leaves open several formulations of a basis-relative theory of aptness. One question concerns the attitude one must take toward the obtaining possibility. Must an agent merely consider that possibility in deciding how to act, or must it play a larger role in their reasoning? This will determine how permissive we are in allowing an agent's anticipation of lucky events to confer aptness. One possibility is to take the intuition generated by million-card poker seriously, and hold that an agent's performance is apt so long as she bases her performance on the obtaining possibility, no matter how unlikely that possibility is (assuming, of course, that the possibility is relevant to the competence). According to such a maximally permissive view, the archer's performance is apt if he considers the double-gust possibility in deciding how to shoot, even if he decides that such a possibility should not alter how he shoots. A maximally permissive view could be motivated by the thought that whatever initially leads us to incorrectly attribute Annie's win to luck also leads us astray in cases in which the obtaining possibility is more unlikely, as in the archery case. If this view is correct, then the simple view gets both cases wrong. However, less permissive views that take seriously the intuition that the archer's shot is inapt, even when he considers the double-gust possibility, are also possible. A less permissive basis-relative view might hold that the double-gust possibility is too improbable to play an appropriate role in the archer's reasoning. Therefore, it is possible for a less permissive view to hold that the accuracy of the archer's shot is attributable to luck, while Annie's win is attributable to her competence. This result is unavailable to the simple view.¹³

IV. EPISTEMIC APTNESS

We have been lead to the following: in order to understand when the success of a performance is attributable to luck, we should determine the relation between the obtaining possibility and an agent's basis for performing, rather than determining whether a “lucky” event caused the performance to succeed. When applied to epistemic performances, this has the potential to emend our understanding of the workings of luck

¹³ Basis-relative views are united by the idea that aptness is primarily determined by the relation between the obtaining possibility and an agent's basis for performing, and not the status of events as “lucky” or “unlucky.” They therefore offer a further benefit in that they do not give a foundational role to concepts like luck and chance in thinking about apt performance.

in Gettier cases. Consider Zagzebski's Mary case, in which Mary sees someone who looks just like her husband sitting in his favorite chair in the living room and comes to believe that her husband is in the living room. However, the man is not her husband, but her husband is indeed in the living room after all, in a corner unseen by Mary.¹⁴ Here, similar to Sosa's archery case, it is not unreasonable to assume that the obtaining possibility—that Mary is looking at someone who only has a similar appearance to her husband, but that her husband is somewhere else in the room nonetheless—is not one that Mary partly bases her belief on. Rather, she bases her belief just on her visual data. It is therefore appropriate to attribute the truth of her belief to luck.

Compare this now with existing treatments of Zagzebski's case. Kelly Becker provides a clear example of a common way that the case is interpreted. He contends that the influence of 'world luck' is responsible for Mary's lack of knowledge.¹⁵ Mary's belief, he holds, is formed via a reliable general process, but 'the world "conspires" in a somewhat unusual way to make her belief only luckily true' (p. 355). But this analysis is too simple because it ignores the significance of Mary's basis for believing. According to a maximally permissive basis-relative view, if an agent affected by world luck anticipates, in forming her belief, the influence of such luck, and this becomes part of the basis for her so believing, then world luck does not preclude the agent from believing aptly. According to less permissive views, even if an agent does anticipate this possibility, it is not suitable to play the appropriate role in her basis for believing. In neither case should we think that whether an agent's belief is attributable to luck solely depends on the way the world 'conspires'.

This criticism also applies to Zagzebski's proposed general procedure for constructing Gettier cases (*Virtues of the Mind*, pp. 288--9):

Start with a case of justified (or warranted) false belief. Make the element of justification (warrant) strong enough for knowledge, but make the belief false...due to some element of luck. Now amend the case by adding another element of luck, only this time an element that makes the belief true after all...The situation might be described as one element of luck counteracting another. We now have a case in which the belief is justified (warranted) in a sense strong enough for knowledge, the belief is true, but it is not knowledge.

14 Adapted from Zagzebski, *Virtues of the Mind*, pp. 285--7.

15 K. Becker, 'Epistemic Luck and the Generality Problem', *Philosophical Studies*, 139 (2008), pp. 353--66, at p. 365.

Zagzebski's procedure ignores an agent's basis for believing, and instead focuses exclusively on elements associated with 'world luck'. But, as is revealed by taking the Mary case and juxtaposing it with that of Annie, it is an agent's basis for believing, rather than world luck, that Gettier cases exploit. Therefore, Zagzebski's procedure may work, but it is not for the reasons she thinks. The procedure is a way to get us to think that an agent did not base her belief on the obtaining possibility, or to make the obtaining possibility so unlikely that it cannot play the appropriate role in an agent's basis for belief. This is ultimately responsible for the intuition that Gettier subjects do not know.

V. CONCLUSION

An apt performance is both accurate and adroit, and at least partly based on the possibility that obtains. A maximally permissive basis-relative view allows for the basing to be very weak, involving just the consideration of the obtaining possibility, while less permissive views would demand that the basing be stronger, and thus have the potential to differentiate between the archery and poker cases. Either way, the simple view of aptness is insufficient, and should be replaced by a basis-relative view.¹⁶

¹⁶ With grateful thanks to Andy Egan and Ernest Sosa for very helpful comments on earlier drafts of this paper. For additional helpful comments and discussions, I thank Robert Audi, Earl Conee, Trent Dougherty, Sanford Goldberg, Michael Huemer, Ben Levinstein, Errol Lord, Lisa Miracchi, Ram Neta, and Mary Salvaggio. Earlier versions of this paper were presented at the 7th Biennial Rochester Graduate Epistemology Conference and the 3rd Annual Northwestern/Notre Dame Graduate Epistemology Conference. I thank both audiences for helpful discussion.