



# Is the folk concept of luck normative?

Mario Attie-Picker<sup>1</sup>

Received: 18 February 2018 / Accepted: 11 February 2019  
© Springer Nature B.V. 2019

## Abstract

Contemporary accounts of luck, though differing in pretty much everything, all agree that the concept of luck is descriptive as opposed to normative. This widespread agreement forms part of the framework in which debates in ethics and epistemology, where the concept of luck plays a central role, are carried out. The hypothesis put forward in the present paper is that luck attributions are sensitive to normative considerations. I report five experiments suggesting that luck attributions are influenced by the normative features of the case, and not merely by descriptive notions like agential control or modal fragility. I discuss the implications of this effect for theories of luck and debates in ethics and epistemology.

**Keywords** Luck · Epistemic luck · Moral luck · Normative mismatch · Lack of control

A person gets really drunk at a friend's party and decides, against everyone's advice, to drive home. She arrives safely and has a good night sleep. There is a strong intuition that she was lucky. It was due to luck, someone might say, that she didn't hurt somebody (or herself).

Now suppose that another friend at the party, as drunk as the first one, decides to drive home as well. He hits a passerby and spends the night in jail. By parity of reasoning, one might claim that he suffered from bad luck, i.e., that he was unlucky. I will be arguing, however, that this is not the conclusion that follows from the ordinary concept of luck. On the ordinary concept, it seems wrong to claim that the accident was due to (bad) luck. Indeed, a case in which an accident was caused by drunk driving appears to be the very *opposite* of a case in which a bad outcome is simply due to luck.

The hypothesis put forward in the present paper is that luck attributions are influenced by a normative evaluation of an agent's action and its relation to the ensuing outcome. Luck attributions are sensitive to whether the valence of the

---

✉ Mario Attie-Picker  
mario.attie@yale.edu

<sup>1</sup> Philosophy Department, Yale University, New Haven, CT 06520-8306, USA

action matches the valence of the outcome. It is the badness of driving drunk in connection with the goodness of avoiding an accident that partly explains the intuition that the first driver was lucky. And similarly, the match of valences between driving drunk and causing harm contributes to our reluctance to ascribe luck to the second driver.

The hypothesis that ordinary luck attributions are sensitive to normative considerations implies that the folk concept of luck is different from what philosophers have traditionally thought about luck. As we shall see, existing philosophical accounts conceptualize luck as an objective, non-normative phenomenon. The present findings therefore have important consequences for debates in ethics and epistemology where the concept of luck plays a major role.

Here is how the paper is structured. In Sect. 1, I introduce the main problem that luck brings out in ethics and epistemology. I then present, in Sect. 2, the main theories of luck that have been proposed in the literature. In Sect. 3, I develop the hypothesis that luck attributions are sensitive to normative considerations and, in Sect. 4, I present empirical evidence in support of it. In Sect. 5, I discuss the relevance of the results for theories of luck. I lay out some of the issues which need to be worked out to give a fuller account of the folk concept of luck. Finally, I consider some of the implications of the results for debates in ethics and epistemology.

## 1 The incompatibility theses

The concept of luck plays a similar role in ethics and epistemology. In ethics, the intuition that luck is incompatible with blameworthiness leads to the question of when, if at all, we are justified in making moral judgments. In epistemology, the analogous intuition, that luck is incompatible with knowledge, brings with it the threat of skepticism. The challenge is to accommodate luck in a framework that leaves space for knowledge and genuine blameworthiness without ignoring the possibility that luck can, at times, undermine them.

The central problem that luck generates in ethics arises from the claim that any pair of agents differing only in their degree of luck should be judged as moral equals. A different way to put it is that morality is immune to luck: responsibility is undermined when luck is the chief influence on morally relevant outcomes. This thought expresses what I shall call the incompatibility thesis.

Moral Incompatibility Thesis: Blameworthiness<sup>1</sup> is incompatible with luck.

The alleged puzzle is that our moral practices are themselves incompatible with the incompatibility thesis. We don't tend to judge like actions in like manner. Rather, we assign more blame to agents whose actions led to harm, even if other agents performed the very same actions (Cushman 2008). Consider the following example: a surgeon, about to perform a relatively mild procedure, negligently forgets to wash

<sup>1</sup> For simplicity, I restrict the incompatibility thesis to blameworthiness, though similar arguments might apply for praiseworthiness.

his hands. The risk of an injury increases, but the nature of the surgery makes its occurrence very unlikely. As it happens, the patient does suffer an injury. There is a strong intuition here that the surgeon is blameworthy and hence deserves some kind of punishment. Now imagine another surgeon, about to perform the exact same surgery as the former, who also forgets to wash her hands. This time, however, the patient leaves the operation unharmed. Should we judge the second surgeon as we judged the first? Supporters of the incompatibility thesis answer positively. However, most people have the intuition that the former surgeon deserves more blame and punishment (Cushman 2008; Young et al. 2010; but see Kneer and Machery 2019 for evidence that people have the intuition only in the case of punishment). It is this intuition, in combination with the incompatibility thesis, that creates the problem of moral luck.

The incompatibility thesis has its epistemic analogue.

Epistemic Incompatibility Thesis: Knowledge is incompatible with luck.

There is widespread agreement among epistemologists that something in the vicinity of this thesis is true. Pritchard (2005) notes the “seemingly universal intuition that *knowledge excludes luck*” (Pritchard 2005, p. 1). The force of the intuition comes from everyday examples of lucky guesses. Suppose that while playing poker, I suddenly get a feeling that I’m about to draw a pair of aces. I warn my friends that a pair of aces is coming, and as a matter of fact, I draw the pair of aces. Did I know I was going to draw a pair of aces? I certainly did not. I just happened to guess correctly. The plausibility of the incompatibility thesis, moreover, is amplified by the usual explanation given to standard Gettier cases. When an agent is Gettierized, her true belief is due to a coincidental, and thus lucky, causal chain. Commentators are mostly in agreement that it is the lucky element in the connection between belief and truth that prevents a Gettierized agent from knowing what she truly, and justifiably, believes.

As they stand, the incompatibility theses pose an existential threat to the possibility of blameworthiness and knowledge. Once luck *simpliciter* has been deemed incompatible with blameworthiness, “the area of genuine agency, and therefore of legitimate moral judgment, seems to shrink under this scrutiny to an extensionless point” (Nagel 1979, p. 35). The predicament is not different in epistemology. Engel puts it boldly: “the incompatibility thesis entails skepticism” (Engel 1992, p. 59).<sup>2</sup>

The reaction has therefore been to amend the theses. The strategy is to say, plausibly, that luck *simpliciter* is *not* incompatible with knowledge [blameworthiness], and that only *some* instances of luck undermine knowledge [blameworthiness]. Consider what Nagel calls ‘constitutive luck’, i.e., luck “in the kind of person you are” (Nagel 1979, p. 28). Suppose the genetic lottery blessed you with a great memory. Does such luck prevent you from knowing that David Bowie was born on January 8? Quite the opposite, it seems. Constitutive luck is therefore an example of knowledge-compatible luck—since not only it doesn’t undermine it, but it can be conducive to knowledge.

<sup>2</sup> The thought seems to rely on the assumption that luck is a ubiquitous and irremovable element of human life.

What kind of luck, then, is thought to be problematic? There is consensus among philosophers that one kind of luck threatens knowledge and blameworthiness, namely, luck in the consequences of one's actions and epistemic practices (usually called 'resultant luck' in ethics Nagel 1979 and 'veritic luck' in epistemology Engel 1992).<sup>3</sup> The idea is that if two agents behave in the same way, morally or epistemically, any difference in how things turned out is due to luck, and in virtue of this fact, they should be assessed equally. Restricting the incompatibility theses to resultant/veritic luck implies that an agent should not be considered blameworthy for any harm resulting from bad luck in the consequences of her action. Similarly, knowledge is undermined when the connection between an agent's belief and its being true is due to luck. This is the case in lucky guesses and, as we have seen, in standard Gettier cases.

In focusing on luck in one's consequences as potentially incompatible with blameworthiness and knowledge I don't mean to suggest that it is the only kind of luck that poses a threat.<sup>4</sup> But it certainly has been the main focus of philosophers in both ethics and epistemology. Following that tradition, I will come back to resultant and veritic luck in testing my hypothesis about ordinary luck attributions.

Notice, however, that before we can evaluate the thesis that luck is incompatible with knowledge [blameworthiness], we need to say what luck is. The incompatibility theses thus show the importance of exploring the nature of luck. In the next section, I review the accounts of luck in the existing literature; this will not only help us see what the incompatibility theses amount to, but more importantly, it will provide a contrast between philosophical accounts of luck and my hypothesis about ordinary luck attributions.

## 2 Current accounts of luck

Three accounts of luck have been defended in the existing literature. My purpose on this section is to present the main idea behind each of them. I don't claim to be exhaustive here: the following should be read as a basic introduction to each account.

- (1) *The lack of control account* In its simplest formulation, the lack of control account states that if an agent had no control over an event, and the event is of significance to the agent, then such event is (un)lucky for that agent (for a characterization along these lines see Statman 1991).<sup>5</sup> The main idea is that a lucky

<sup>3</sup> Most discussions of moral luck trade in examples of resultant luck, and I think it is fair to say that the jury is still out regarding its compatibility with blameworthiness (Nagel 1979; Statman 1993). In epistemology, veritic luck is widely held to be incompatible with knowledge (Axtell 2001; Engel 1992; Gettier 1963; Pritchard 2005, 2008; though see Baumann 2014; Greco 2003 for a different position).

<sup>4</sup> Another type of luck that has been claimed to undermine blameworthiness and knowledge is luck in the circumstances (or evidential situation, in epistemology) one encounters. Again, there is debate on how real the threat is. (For discussion in ethics, see, e.g., Doris 2002; Harman 2003; Kamtekar 2004, for discussion in epistemology, see, e.g., Axtell 2001; Engel 1992; Harman 1973; Vahid 2001).

<sup>5</sup> A flurry of counterexamples (e.g., Lackey 2008; Latus 2000) has led to ever more complex formulations of the account. For example, Latus notes that the rising of the sun is a daily event that is beyond my control and is highly significant for me, but it seems wrong to say that I am lucky every morning. Riggs

event or outcome is one that was beyond the agent's powers to bring about or affect in any significant way.

With this core intuition in mind, consider the surgeon of Sect. 1: was the patient's injury, once the surgeon has forgotten to wash his hands, an unlucky event *for the surgeon*?<sup>6</sup> The lack of control account answers positively. If we assume that the occurrence of the injury depends on how the patient's body reacts to the surgeon's dirty hands, then it seems clear that the injury lies too far out the surgeon's control, and is thus an example of bad resultant luck.

The lack of control account, as the surgeon's example illustrates, is especially suitable in generating problems of moral luck. Therefore, it is hardly surprising that it has been the theory of luck implicitly assumed by most moral luck theorists. Williams (1981), in his seminal treatment of the subject, writes, "what is not in the domain of the self is not in its control, and so is subject to luck" (Williams 1981, p. 20). Nagel followed him in locating luck in the absence of control, "Where a significant aspect of what someone does depends on factors beyond his control, yet we continue to treat him in that respect as an object of moral judgment, it can be called moral luck" (Nagel 1979, p. 26). Zimmerman (1993) takes the principle that "Something which occurs as a matter of luck is something which occurs beyond anyone's control" as the starting point for his putative solution to the problem of moral luck (Zimmerman 1993, p. 219). Similar identifications of luck with lack of control in treatments of moral luck can be found in Richards (1993, pp. 167–168), Srinivasan (2015, p. 284) and Statman (1991, p. 146). The prevalence of the lack of control account in discussions of moral luck makes its truth a matter of great importance. Most of the arguments in Nagel (1979), for example, would lose much of their bite were we to adopt a different account of luck.

- (2) *The modal account* The modal account states that an event is lucky iff it occurs in the actual world but—keeping the initial conditions fixed—there are close possible worlds where it fails to obtain (Pritchard 2005, 2008, 2015).<sup>7</sup> The modal account sees luck in terms of how different the world would have to be for the target event to have failed to obtain. Lucky events are those that would not have happened had the world been slightly different.

---

Footnote 5 (continued)

(2009) amended the account to deal with these sorts of counterexamples. According to his formulation, "E is lucky for S iff (a) 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 S" (Riggs 2009, p. 220), where exploitation is something like using the belief (or knowledge) that E will obtain as part of one's future plans (see Riggs 2009, p. 216, for the thought experiment that purports to establish condition (b), see also Broncano-Berrocal 2015; Coffman 2007; Levy 2009 for so-called 'mixed-accounts,' that is, accounts with both modal and control conditions).

<sup>6</sup> It seems obvious that it is for the patient.

<sup>7</sup> Note that Pritchard's formulation does not make reference to a subject. He argues that a philosophical theory of luck should deal with "luck as an objective phenomenon" (Pritchard 2015, p. 154).

Pritchard's modal account was devised as part of his project of developing an anti-luck epistemology. Pritchard's analysis of knowledge aims at giving a precise meaning to the incompatibility thesis. The theory roughly states that knowledge is *safe* true belief, where safety is formulated so as to eliminate veritic luck. A belief is safe iff it is true in the actual world, and in most nearby possible worlds, where the believer forms the belief in the same way; or in other words, *if it is not vertically lucky*. Pritchard's modal account, then, is purposely tied with his epistemology (Pritchard 2005).

The modal account has been recently used to give a novel account of moral luck. Focusing on resultant luck, Whittington (2015) offers the following formulation: "S is morally lucky that E iff (1) S's action had been performed in the same way as in the actual world but the results (E) would have been different in a wide set of relevant nearby worlds. (2) The results (E) of S's action are of positive or negative moral value (where the moral value is defined by the moral theory that we are using)" (Whittington 2015, p. 216). It is worth noting that lack of agential control and modal fragility usually go together. An outcome tends to be more stable when an agent is in control of it: the less control the agent has, the more variable the outcome becomes. But this relationship doesn't generalize. There are cases where an event is safe but no agential control is present, and vice versa.

- (3) *The probability account* According to the probability account, an event is lucky for an agent iff the event's occurrence is improbable relative to some condition. In epistemology, for example, a true belief would be lucky iff relative to a given circumstance and method of belief-formation, the probability of its being truth is below some threshold (Baumann 2014). Baumann gives the following example: "Julie's true belief that there is an Air France airplane flying by is lucky insofar as the probability of getting this right, given that she just looks at the sky, is very low" (Baumann 2014, p. 543). What goes into determining the relevant threshold and delimiting the appropriate context is up for debate. For our purposes, the key claim is that whether an event is lucky is entirely dependent on how antecedently probable it was given a set of conditions.

### 3 The normative mismatch effect

An important feature of the existing accounts is their descriptive character. Their central concepts—lack of control, modal fragility and probability—attempt to describe an objective condition in the world. I'm going to argue that luck attributions are sensitive to normative considerations in a way that (at face value) is not predicted by any of these accounts. Now, it is important to note that value does come into the picture in some of the existing theories. As we have seen, some theorists posit a 'significance' condition, to the effect that for an event to count as (un)lucky it has to *matter* for someone. Nevertheless, the significance condition treats good and bad outcomes symmetrically. According to the lack of control account, for example,

if something bad happens to me beyond my control, then it is a case of bad luck. And if something good happens to me beyond my control, it is a case of good luck. My claim is that value comes into the picture in a different way: luck attributions depend not only on the valence of the outcome, but also on the valence of the action that brought it about. Specifically, they are influenced by the *normative relation* between action and outcome.

The core idea is that luck attributions are sensitive to whether the valence of the action matches the valence of the outcome: people are more willing to attribute luck to an event when the valence of the outcome *does not match* the valence of the action.<sup>8</sup>

Consider the surgeon case as an example.

**Negligent Surgeon** About to perform a relatively mild procedure, a surgeon forgets to wash his hands. As a result, the chances of a serious injury rise to about 30%.

As a matter of fact, the patient does not suffer a serious injury.

It is the normative mismatch, or so I say, between the badness of the surgeon's action and the goodness of the outcome that (at least in part) explains our intuition that he is lucky (and that the success of the procedure was due to luck). The explanation is fundamentally normative: it is the evaluation of the surgeon's action as *bad* in conjunction with the evaluation of the result of the procedure as *good* that influence our intuition that luck was involved.

Now if I'm right, and people in fact attribute luck to the success of the procedure, this would be a case in which the effect of normative mismatch trumps the influence of probabilities, for the event had a 70% chance of occurring. And even if one takes a probability of 70% to be low enough for the outcome to count as lucky in this context, the probability account would predict that the case in which the patient does suffer an injury would be seen as *more* lucky, given its lower antecedent probability (30%). But this seems false. It seems wrong to say that the surgeon is unlucky if he is negligent and something bad happens as a result of his negligence.

The putative effect of normative considerations on luck attributions is not restricted to ethics. The mismatch relation can potentially function in any domain that lends itself to normative evaluation. In epistemology, for example, the prediction is that luck attributions are influenced by a mismatch between the normative evaluation of an agent's method of belief formation (e.g., virtuous, flawed) and the valence of its epistemic result (e.g., true or false belief). Consider the following case.

**Virtuous Mathematician** Martha, a mathematician, is trying to prove a complicated theorem. After weeks of unsuccessful attempts, where she kept pushing the same line of reasoning, she starts looking at the problem from a different perspective. The change of focus, added to her proven mathematical abilities, enables her to discover new ways to approach the

<sup>8</sup> The relevant notions of "match" and "mismatch" turn out to be subtler to spell out than one might at first expect; I discuss some of the issues that arise in Sect. 5.1.

problem. As a result, the chances of Martha proving the theorem rise to about 30%.

As a matter of fact, Martha proves the theorem.

Again, the prediction is that the normative match between her epistemic practices and the solving of the theorem prevails over the improbability of the event in people's ascriptions of luck. If this is right, then people should be unwilling to describe Martha (and the solving of the theorem) as lucky.

The last two cases show that the normative mismatch effect makes interesting novel predictions regarding luck attributions. An intriguing possibility, however, is that the effect also explains the cases in which existing accounts seem to give the right answer. Consider lucky guesses. Imagine a fortuneteller who claims to have the power to correctly predict the winner of the next five Super Bowls by looking at the position of the stars. She takes a good look at the night sky, makes her predictions (and comes to believe in them), and 5 years after, they all come out true. Was the fortuneteller lucky? The uncontroversial answer is *yes*, and existing theories have a good explanation of why that is. The modal account points to the fragility of the result: in most nearby worlds where the fortuneteller forms her prediction by looking at the stars, it comes out false. The probability account, in turn, says that relative to her method of belief-formation and the nature of football, her getting the right answer was objectively improbable. Finally, the lack of control account notes that the Super Bowl is an event outside her control. But the result is also predicted by the normative mismatch effect. For the valence of her method of belief formation is epistemically *flawed*, and hence doesn't match the valence of the outcome (the goodness of a true belief). Whether this result generalizes is an open question. For now, it is worth keeping in mind that the normative mismatch effect could account for the cases that made existing theories appealing in the first place.

Of course, the opposite could also be true. That is, it may be that the normative mismatch effect is explained by the other theories. It is possible, in other words, that once we understand how people think about agential control or probabilities, for example, we would see that the normative mismatch effect is in fact predicted by the existing theories. I come back to this view in Study 4.

Finally, a different possibility is that there are numerous factors that have an impact on luck attributions. It could be, for example, that both modal fragility and normative mismatch play a role in luck judgments. On this picture, there are multiple *independent* considerations that go into the folk concept of luck.

To sum up, there appears to be an effect of normative considerations on luck attributions that is not (at least not at first) predicted by existing accounts of luck. If this is true, further questions arise: Can the effect explain the cases that make existing accounts appealing? Or can existing theories, contrary to first appearances, explain the effect? Is the effect simply one of multiple factors that impact luck attributions? Or, finally, is there another concept that explains both the existing accounts and the normative mismatch effect? The next section is a first attempt at answering these questions.



## 4 Testing the effect of normative mismatch on luck attributions

The key hypothesis is that normative mismatch—the mismatch between the valence of an agent’s action and the valence of the resulting outcome—has an impact on luck attributions. The goal of the experiments is to see whether this is true, and if so, to explore the nature of this effect and its relation to existing accounts of luck.

### 4.1 Study 1: normative mismatch and resultant luck

Study 1 tested whether luck attributions in cases of resultant moral luck are sensitive to a relation of normative mismatch using the Negligent Surgeon case introduced in the last section.

#### 4.1.1 Method

*Participants* This and all subsequent samples were recruited from Mechanical Turk. In Study 1, 327 participants (mean age = 34, 42% female) were recruited. Three participants were excluded because they failed to complete the study.

*Procedure* The study used a 2 (Action-Valence: negligent vs. virtuous) × 2 (Outcome-Valence: good vs. bad) × 3 (Vignette: ‘Surgeon’ vs. ‘Commander’ vs. ‘Lawyer’) between-subjects design.

Here is an example of one vignette (brackets indicate the different conditions, see “Appendix” for all vignettes).

**Negligent action** About to perform a relatively mild procedure, a surgeon forgets to wash his hands. As a result, the chances of a serious injury rise to about 30%.

As a matter of fact, the patient suffers [does not suffer] a serious injury.


**Virtuous action** About to perform a complicated procedure, a surgeon takes special precautions, reviewing each part of the procedure carefully. As a result, the chances of a serious injury drop to about 30%.

As a matter of fact, the patient suffers [does not suffer] a serious injury.

*Dependent variables* Two different items were used to measure luck attributions.

*Luckiness of the event* Participants indicated their agreement with the following statement, “It was due to luck that the patient got injured [didn’t get injured],” using a 7-point Likert scale ranging from “disagree” to “agree.”

*Luckiness of the agent* Participants were asked the question, “To what extent would you describe the surgeon as [un]lucky”? They answered using a 7-point Likert scale ranging from “not at all” to “great extent.”



**Fig. 1** Luckiness of the event attributions: mean responses by condition in Study 1 collapsing across vignettes. Error bars indicate standard error of the mean

#### 4.1.2 Results

Mean responses for each condition are displayed in Fig. 1.

The results were analyzed using a 2 (Action-Valence: negligent vs. virtuous)  $\times$  2 (Outcome-Valence: good vs. bad)  $\times$  3 (Vignette: 'Surgeon' vs. 'Commander' vs. 'Lawyer') ANOVA on luckiness of the event attributions. There was a main effect of Outcome-Valence,  $F(1, 312) = 23.42$ ,  $p < .001$ ,  $\eta^2 = 0.06$ , such that participants in the bad conditions were less inclined to attribute luck to the event (e.g., the injury) than those in the good conditions (e.g., the patient not getting injured). Critically, there was a significant Action-Valence  $\times$  Outcome-Valence interaction,  $F(1, 312) = 57.53$ ,  $p < .001$ ,  $\eta^2 = 0.15$ . As predicted, participants were more inclined to attribute luck when the valence of the action did not match the valence of the outcome.

The interaction was decomposed to further understand the effect. Participants in the bad outcome conditions were less inclined to attribute luck to the event when the action that brought it about was negligent ( $M = 1.79$ ,  $SD = 1.17$ ) than when it was virtuous ( $M = 2.90$ ,  $SD = 1.51$ ),  $t(312) = 4.84$ ,  $p < .001$ ,  $d = 0.67$ ,

whereas participants in the good outcome conditions showed the opposite pattern. They attributed less luck when the action leading to the good outcome was virtuous ( $M = 2.45$   $SD = 1.32$ ), than when it was negligent ( $M = 3.77$ ,  $SD = 1.74$ ),  $t(312) = -5.87$ ,  $p < .001$ ,  $d = 0.81$ .

Further analyses were carried out to see if there was a difference between vignettes. There was a significant two-way interaction between Outcome-Valence and Vignette,  $F(2, 312) = 4.64$ ,  $p = .01$ ,  $\eta^2 = 0.02$ . The interaction was due to the fact that the ‘Commander’ vignettes, unlike ‘Surgeon’ and ‘Lawyer’, did not show a main effect of Outcome-Valence,  $F(1, 312) = 0.18$ ,  $p = .66$ . Importantly, there was no three-way interaction.<sup>9</sup>

### 4.1.3 Discussion

This first study aimed to test the normative mismatch effect by giving participants scenarios where it makes different predictions from the probability (and arguably also the modal) account. The vignettes thus kept the probability of bad outcomes constant and manipulated the valence of the action (negligent vs. virtuous) and the valence of the outcome (good vs. bad). Note that according to the probability account, bad outcomes should be judged equally lucky (given their 30% chance of occurring) whereas good outcomes should not be judged as involving luck (given their 70% chance of occurring).<sup>10</sup> And crucially, whether the outcome was brought about by a negligent or a virtuous action should not matter at all.

The results suggest that normative mismatch has an effect on luck attributions. Participants were more willing to attribute luck when the valence of the action didn’t match the valence of the outcome. For example, participants attributed more luck (both to the agent and to the event) when the surgeon acted negligently, but performed the operation successfully (an event with 70% chance of occurring) than when the patient suffered an injury (an event with 30% chance of occurring). The opposite pattern was true for the virtuous surgeon: luck attributions decreased when the patient didn’t suffer an injury.

Surprisingly, participants were in general more willing to attribute luck to good outcomes than to bad outcomes, despite the former’s higher antecedent probability. I discuss a possible explanation for this effect in Sect. 4.4.

<sup>9</sup> The results were virtually identical when the dependent variable was the luckiness of the agent. For the sake of space, I will not discuss these results any further. Suffice it to say that the pattern of results observed for luckiness of the event was observed for luckiness of the agent. This was true as well for Study 3 (luckiness of the agent was not measured in Studies 2 and 4).

<sup>10</sup> In other words, the probability account predicts a main effect of outcome (where bad outcomes should be judged luckier than good outcomes) and no interaction. In contrast, the normative match effect predicts an interaction reflecting the impact of normative mismatch.

## 4.2 Study 2: understanding the effect of probability

The results of the previous study support the hypothesis that luck attributions are sensitive to normative mismatch. However, a possible confound threatens the validity of the results: in Study 1, virtuous actions always lowered, and negligent actions always raised, the antecedent probability of a bad outcome. It is possible that it is the raising of the probability, and not the valence of the action, that is driving luck attributions.

To disassociate these variables, a further condition was added to the design used in Study 1. Study 2 included a condition where a virtuous action *increases* the probabilities of a bad outcome obtaining (and vice versa). One of the vignettes, for example, describes a Mayor who, by acting rightly (he decides to expand the construction of a public hospital), makes a bad outcome (a family losing their house) more likely. The goal was to test whether the results of the previous study can be accounted by the normative mismatch effect, independently from the effects of probability.

### 4.2.1 Method

*Participants* 410 participants (mean age = 32, 41% female) were recruited. Four participants were excluded because they failed to complete the study.

*Procedure* The study used a 2 (Action-Valence: corrupt vs. virtuous) × 2 (Outcome-Valence: good vs. bad) × 2 (Probability: increasing chances of bad outcome vs. decreasing chances of bad outcome) between-subjects design.

The following vignette structure was used (brackets indicate the different conditions).

**Virtuous action, probability of bad outcome decreases** A mayor is considering suspending an ongoing construction of a private golf course. The construction was forcing some families to be relocated. The mayor, known for his responsible behaviour, and deep commitment to the wellbeing of his people, orders an independent agency to consider whether the benefits outweigh the costs. The agency's detailed investigation reveals that the suspension is the best option for the overall good of the town. With the interest of each citizen in mind, and determined to keep serving the people to the best of his abilities, the mayor orders suspending the construction.

Thus, as a result of the order, the chances of the Smiths losing [keeping] their house drop [rise] to about 30%.

As a matter of fact, the Smiths lose [keep] their house.

**Corrupt action, probability of bad outcome decreases** A mayor is considering expanding an ongoing construction of a public hospital. The construction was forcing some families to be relocated. The mayor, known for his lack of responsibility, and utter detachment from the people he represents, orders an independent agency to consider whether the benefits outweigh the costs. The agency's detailed investigation reveals that the expansion is the best option for the overall good of the town.

But thinking about the power of the businessmen involved, and recognizing the benefits of doing them a favor, the mayor decides to forego the expansion. Thus, as a result of the decision, the chances of the Smiths losing [keeping] their house drop [rise] to about 30%.

As a matter of fact, the Smiths lose [keep] their house.

**Virtuous action, probability of bad outcome increases** A mayor is considering expanding an ongoing construction of a public hospital. The construction was forcing some families to be relocated. The mayor, known for his responsible behavior, and deep commitment to the wellbeing of his people, orders an independent agency to consider whether the benefits outweigh the costs. The agency's detailed investigation reveals that the expansion is the best option for the overall good of the town. With the interest of each citizen in mind, and determined to keep serving the people to the best of his abilities, the mayor orders expanding the construction.

Thus, as a result of the order, the chances of the Smiths losing [keeping] their house rise [drop] to about 30%.

As a matter of fact, the Smiths lose [keep] their house.

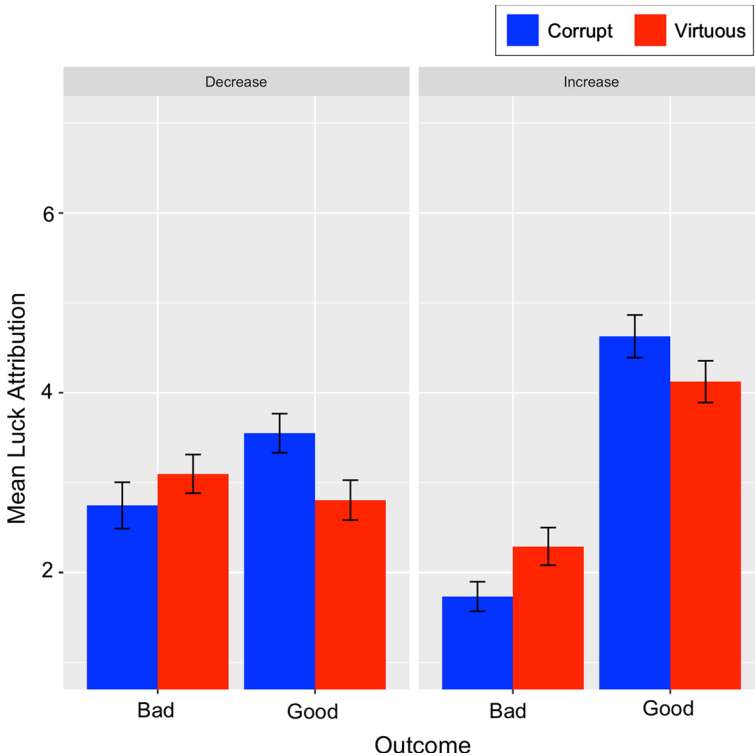
**Corrupt action, probability of bad outcome increases** A mayor is considering suspending an ongoing construction of a private golf course. The construction was forcing some families to be relocated. The mayor, known for his lack of responsibility, and utter detachment from the people he represents, orders an independent agency to consider whether the benefits outweigh the costs. The agency's detailed investigation reveals that the suspension is the best option for the overall good of the town. But thinking about the power of the businessmen involved, and recognizing the benefits of doing them a favor, the mayor decides to forego the suspension. Thus, as a result of the decision, the chances of the Smiths losing [keeping] their house rise [drop] to about 30%.

As a matter of fact, the Smiths lose [keep] their house.

Note a methodological modification with respect to Study 1. In Study 1 the probability of the bad outcomes obtaining was fixed to 30% (and hence the probability of the good outcomes was always 70%). In Study 2, the probability of the outcome that obtained (whether good or bad) was fixed to 30%. That is, each participant, regardless of condition, read about a scenario where the event that obtained had an antecedent probability of 30%. The modification was made to further control the effects of probability.

*Dependent variables* Since the mayor in the corrupt conditions was described as lacking any interest in the well-being of the citizens, it didn't make sense to ask the luckiness of the agent question. Thus, only luckiness of the event attributions were measured.

*Luckiness of the event* Participants indicated their agreement with the following statement, "It was due to luck that the Smiths lost [kept] their house," using a 7-point Likert scale ranging from "disagree" to "agree."



**Fig. 2** Mean responses by condition in Study 2. Error bars indicate standard error of the mean

#### 4.2.2 Results

Mean responses for each condition are displayed in Fig. 2.

The results were analyzed using a 2 (Action-Valence: corrupt vs. virtuous)  $\times$  2 (Outcome-Valence: good vs. bad)  $\times$  2 (Probability: increasing chances of bad outcome vs. decreasing chances of bad outcome) ANOVA on luckiness of the event attributions. There was a main effect of Outcome-Valence,  $F(1, 398) = 69.93$ ,  $p < .001$ ,  $\eta^2 = 0.14$ , such that participants in the bad conditions were less inclined to attribute luck to the event (e.g., the Smiths losing their house) than those in the good outcome conditions (e.g., the Smiths keeping their house). Crucially, there was no main effect of Probability,  $F(1, 398) = 0.85$ ,  $p = .35$ . Importantly, the effect of the previous studies was observed here. There was a significant Action-Valence  $\times$  Outcome-Valence interaction,  $F(1, 398) = 11.83$ ,  $p < .001$ ,  $\eta^2 = 0.02$ .

Follow-up tests showed that the interaction followed the relation of normative mismatch. Participants in the bad outcome conditions were less inclined to attribute luck to the event when the action that brought about the bad outcome was corrupt ( $M = 2.25$ ,  $SD = 1.61$ ) than when it was virtuous ( $M = 2.69$ ,  $SD = 1.57$ ),  $t(398) = 2.04$ ,  $p = .04$ ,  $d = 0.25$ , whereas participants in the good outcome

conditions showed the opposite pattern. They attributed less luck when the action leading to the good outcome was virtuous ( $M=3.45$ ,  $SD=1.73$ ), than when it was corrupt ( $M=4.08$ ,  $SD=1.70$ ),  $t(398)=-2.81$ ,  $p<.01$ ,  $d=0.34$ .

There was also a significant interaction between Outcome-Valence and Probability,  $F(1, 398)=45.27$ ,  $p<.001$ ,  $\eta^2=0.10$ . Participants in the bad outcome conditions were less willing to ascribe luck when the agent's action increased the probability of its occurrence ( $M=2.01$ ,  $SD=1.37$ ) than when it decreased it ( $M=2.92$ ,  $SD=1.70$ ),  $t(398)=-4.11$ ,  $p<.001$ ,  $d=0.50$ . Similarly, participants in the good outcome conditions were less willing to ascribe luck when the action increased the probability of its occurrence ( $M=3.17$ ,  $SD=1.60$ ) than when it decreased it ( $M=4.38$ ,  $SD=1.67$ ),  $t(1, 398)=5.39$ ,  $p<.001$ ,  $d=0.66$ . There was no three-way interaction.

### 4.2.3 Discussion

The results of the current study buttress the results of Studies 1: luck attributions were again sensitive to the relation of normative mismatch. The design of the previous study was such that virtuous actions always decreased, and negligent actions always increased, the probability of bad outcomes obtaining. These variables were disassociated in Study 2.

The results provided further evidence for the existence of the normative mismatch effect. Even when controlling for whether the action raised or lowered the antecedent probability of bad outcomes, luck attributions again tracked the relation of normative mismatch. That is, participants were more inclined to ascribe luck when the valence of the action didn't match the valence of the outcome.

Note, however, that regardless of the valence of the action, increasing the probability of the bad outcome (and vice versa) did impact luck attributions: participants were less willing to attribute luck when the action increased the probability of the outcome that obtained. Nevertheless, normative mismatch had a significant impact above and beyond the effects of probability.

### 4.3 Study 3: confounding variables?

The results of studies 1–2 are consistent with the hypothesis that luck attributions are sensitive to normative considerations. However, they are not sufficient to rule out an alternative explanation, namely, that the manipulation of normative mismatch affects other kinds of non-normative judgments that in turn affect luck attributions. This sort of explanation has been proposed in other domains (e.g., intentional action) in which normative considerations have been shown to influence concepts previously thought to be descriptive (see, e.g., Uttich and Lombrozo 2010). Consider the manipulation used in Study 1. With the goal of manipulating action-valence, a doctor either forgets to wash his hands or takes special precautions with the medical procedures. However, all sorts of non-normative factors might be confounded with the doctor's actions (and the same is true for the manipulations used in the other studies). Perhaps, then, it is these other factors that account for the effect of normative mismatch on luck attributions.

To address this concern, Study 3 employed a different design. The perceived valence of the action (and hence the perceived normative mismatch relation) was not manipulated across conditions (as in previous studies), but rather depended on the moral views of the participants themselves. Participants read a story about a university president faced with the task of deciding whether or not to cancel an upcoming talk by a controversial speaker. The perceived valence of the president's action thus varied with individual differences in judgments about what the president should do. For example, if certain participants think that she did the wrong thing, the normative mismatch effect predicts that those participants would see good outcomes as lucky, for the valence of the outcome would not match the valence of the action. By contrast, those who think that she did the right thing should not see the good outcome as due to luck, for the valence of the outcome would match the valence of the action.

This design thus enabled a stronger test of the normative mismatch effect without manipulating the normative mismatch relation. It thus avoided the confounds that might potentially befall the previous studies.

#### 4.3.1 Method

*Participants* The study was pre-registered (<https://aspredicted.org/65j7a.pdf>). As indicated in the pre-registration, the goal was to recruit at least 240 participants answering that the president should cancel the event, and 240 with the view that the president should let the speaker give the talk. To that effect, an initial sample of 550 participants was recruited. Subsequently, additional batches of 50 participants were recruited until the target of 240 per moral view was reached.

762 participants (mean age = 36, 42% female) were recruited. One hundred and thirteen participants were excluded because they failed to pass an attention check.

*Procedure* The study used a 2 (Action: president allows vs. president cancels)  $\times$  2 (Outcome-Valence: success vs. failure) between-subjects design.

The following vignette was used (brackets indicate the different conditions).

A student organization at a university has invited a controversial speaker to give a talk. Many students, however, regard the speaker as a bigot. They claim that the he has a well-documented history of racist and misogynistic views, and argue that the university should have no business giving him a platform to express such views. A large portion of the student body signs an open letter addressed to the university president asking her to cancel the event. An excerpt of the letter reads, “the speaker’s deplorable views pass from protected free speech to incitement, harassment and defamation once they publicly target individuals in his audience or on campus. Such actions are protected neither by free speech nor by academic freedom. For this reason, the university should not provide a platform for such harassment.”<sup>11</sup>

On the other hand many students want the president to let him speak. They argue that a university should foster dialogue and encourage students to confront and evaluate their beliefs in light of new ideas. If universities suppress minority views,

<sup>11</sup> From, members of the UC Berkeley faculty, “Open letters calling for cancellation of Milo Yiannopoulos event,” *The Daily Californian*, January 10, 2017.



they think, how are they supposed to train their students to think for themselves? Moreover, they contend that the first amendment was designed to protect unpopular speech. They, too, pen an open letter calling the president to let him speak. Here's an excerpt from the letter, "There's a distinction between deciding whether to invite a controversial speaker, or silencing him once he has already been invited. Certainly the second type of speech is protected. But the present speaker had already been invited by a student group and funded by the university. While peacefully protesting an event is certainly justified, attempting to prevent it from even being held is deeply troubling."<sup>12</sup>

A great deal of pressure surrounds the university president, for she has the final saying on whether the speaker is allowed to give the talk. She is well aware of the importance of this decision and the serious consequences it might bring about. The president aims, above all, to create at the university an environment of communal tolerance and intellectual freedom. She carefully reflects on the arguments from both sides and, with the goal of promoting stimulating dialogue across and mutual respect between students with differing viewpoints, she decides to let the speaker give his talk [to cancel the event].

There was a lot of disagreement about what the consequences of the president's decision would be, with some people predicting that things would work out quite well and others predicting that things would work out very badly. A period of great uncertainty ensued in the days surrounding the decision. As it happened, her decision was a total success [failure]. In the months after the event, the university became a paradigm of tolerance and freedom, recognized by all for the exemplary relations among students and its invigorating intellectual climate [the university became a paradigm of intolerance and fear, recognized by all for the deplorable relations among students and its deteriorated intellectual climate].

*Dependent variables* Two different items were used to measure luck attributions.

*Luckiness of the event* Participants indicated their agreement with the following statement, "The success [failure] of the president's decision was due to *luck*," using a 7-point Likert scale ranging from "disagree" to "agree."

*Luckiness of the agent* Participants were asked the question, "To what extent would you describe the president as [un]lucky" for her success [failure]? They answered using a 7-point Likert scale ranging from "not at all" to "great extent."


*Moral view* Participants indicated their moral view by answering the following question: "What do *you* think was the morally right thing to do?" They were asked to select between "cancel the event" and "let the speaker give the talk." All participants answered this question after providing answers on the two dependent measures.

### 4.3.2 Results

Mean responses for each condition are displayed in Fig. 3.

The results were analyzed using a 2 (Action: president allows vs. president cancels)  $\times$  2 (Outcome-Valence: success vs. failure)  $\times$  2 (Participant's moral view:

<sup>12</sup> From Anna Mitchell, "Let Milo Speak," *The Stanford Review*, February 12, 2017.



**Fig. 3** Mean responses by condition in Study 3. Left: participants with the view that the president should have let the speaker talk. Right: participants with the view that the president should have canceled the event. Error bars indicate standard error of the mean

**Table 1** In every condition, participants with differing moral views significantly differed in luck attributions in the way predicted by the normative match account

	Should cancel	Should allow	<i>t</i>	<i>p</i>	<i>d</i>
President allows					
Success	4.44 (1.79)	3.24 (1.77)	- 3.79	.0002	0.67
Failure	2.18 (1.30)	2.83 (1.87)	2.16	.03	0.40
President cancels					
Success	3.15 (1.54)	3.87 (2.01)	2.66	.007	0.40
Failure	3.03 (1.76)	2.35 (1.48)	- 2.38	.01	0.41

should allow vs. should cancel) ANOVA on luckiness of the event attributions. As recorded in the pre-registration, a three-way interaction was predicted: perceived normative mismatch was hypothesized to vary according to what participants regarded as the right action for the president to take.

There was a main effect of Outcome-Valence,  $F(1, 641)=53.87$ ,  $p < .001$ ,  $\eta^2=0.07$ , such that participants in the failure conditions were less inclined to attribute luck to the event than those in the success conditions. No other main effect or two-way interaction reached significance, all  $p$ 's  $> .07$ . Crucially, the predicted three-way interaction between Action x Outcome-Valence x Participant's moral view was significant,  $F(1, 641)=30.74$ ,  $p < .001$ ,  $\eta^2=0.04$ .

Planned contrasts, displayed in Table 1, showed that the interaction followed the relation of normative mismatch as realized by the participants' own moral views.

### 4.3.3 Discussion

The results of this study once again provided evidence for the normative mismatch effect. Participants were more willing to attribute luck when the valence of the outcome did not match the valence of the action. The design of the present study, however, was more complex. Instead of manipulating action-valence (e.g., acting negligently vs. acting virtuously), the study relied on individual differences across participants in the evaluation of the action. Therefore, perceived normative mismatch was not experimentally manipulated, but rather depended on participants' own moral views.

The results are therefore further evidence against the hypothesis that non-normative confounding variables, and not the normative mismatch relation, accounted for the results from the previous studies. For luck attributions differed significantly among participants with different moral views responding to the same scenario. For example, when the president decided to let the speaker give the talk and the decision was successful in creating a positive environment at the university, participants who disagreed with the decision judged the outcome as lucky. Those who judged the president's action as morally right, however, did not attribute the success to luck.

Of course, it might be that other factors are confounded with normative mismatch. If that were the case, those factors would have to vary with individual differences on moral views. That is, one would have to hold that the confounding variable, whatever that may be, is itself correlated with what the participant judges as morally right. To be sure, the present study does not rule out *that* further hypothesis. It does, however, make the case that luck attributions are directly sensitive to normative considerations more compelling.

## 4.4 Normative mismatch and theories of luck

The results of Studies 1–3 support the claim that normative considerations play an important role in ordinary luck attributions: across a variety of different scenarios, luck attributions were highly sensitive to the normative mismatch relation. Does that show that existing philosophical theories are at odds with the folk notion of luck? Not necessarily. It might be that the pattern observed in the experiments falls out of one of the theories discussed in Sect. 2. This is because the three concepts that are at the heart of such accounts—agentic control, modality, and probability—have been shown to be influenced by normative considerations in people's thinking (Alicke

2000; Phillips and Cushman 2017; Dalbauer and Hergovich 2013). For example, Dalbauer and Hergovich (2013) show that subjective probabilities are influenced by outcome valence: people take bad outcomes to be more likely. If, in general, bad outcomes subjectively feel as more likely to occur, and the probability account of luck is right, then people should be less willing to judge bad outcomes as due to luck. The effect might remain even when people have information about the objective probabilities of such outcomes. And in fact, a main effect of outcome-valence was observed in all the experiments reported above—people indeed rated bad outcomes as less due to luck. Contrary to first appearances, then, the results may be in line with the probability account of luck (note, however, that the account only predicts the main effect of outcome, and thus leaves unexplained the interaction effect, i.e., the normative mismatch relation). A similar story could perhaps be told by proponents of both the lack of control and modal accounts of luck.<sup>13</sup>

The goal of the present experiment was to directly test whether the effect of normative mismatch is explained by its impact on notions of control, modal fragility or probability. The experiment, in other words, enables us to determine whether the results of the previous studies are consistent with existing accounts of luck.

#### 4.4.1 Method

*Participants* 408 participants (mean age = 35, 47% female) were recruited. *Procedure.* The study used a 2 (Action-Valence: negligent vs. virtuous) × 2 (Outcome-Valence: good vs. bad) × 3 (Vignette: ‘Surgeon’ vs. ‘Commander’ vs. ‘Lawyer’) between-subjects design.

Here is an example of one vignette (brackets indicate the different conditions. The vignettes were virtually identical to those of Study 1<sup>14</sup>):

**Negligent action** About to perform a complicated procedure, a surgeon forgets to wash his hands. As a result, the chances of a failed [successful] procedure rise [drop] to about 30%.

As a matter of fact, the procedure fails [succeeds].

**Virtuous action** About to perform a complicated procedure, a surgeon takes special precautions, reviewing each part of the procedure carefully. As a result, the chances of a successful [failed] procedure rise [drop] to about 30%.

<sup>13</sup> Proponents of the lack of control account might suggest that bad outcomes lead us to ascribe more control to the agent, and consequently less luck to the event, explaining the main effect observed in the results (cf. Alicke 2000). Fans of the modal account, in turn, could propose that given a bad action, the possibility of a bad outcome seems modally closer. And thus, if a good outcome obtains instead, people would be led to posit luck. This explanation would be able to account for the effect of normative mismatch on luck attributions.

<sup>14</sup> There were two changes with respect to the vignettes of Study 1. First, the probability of the outcome that obtained was fixed to 30%. Second, the description of the outcome was modified to avoid negative constructions. Instead of saying, ‘the patient does not suffer an injury’ (as in Study 1) the vignette read, ‘the procedure succeeds.’ To make it parallel, the description of the bad outcome was modified from ‘the patient suffers an injury’ to ‘the procedure fails.’ Analogous changes were made to the other vignettes (see ‘Appendix”).

As a matter of fact, the procedure succeeds [fails].

*Dependent variables* Participants indicated their agreement with the following statements using a 7-point Likert scale ranging from “disagree” to “agree” (participant always answered the luck question first. The order of the other questions was randomized):

*Luck* “It was due to luck that the procedure failed [succeeded].”

*Control* “The procedure’s failure [success] was outside the surgeon’s control.”

*Modal fragility* “[Even] given the surgeon’s action, the procedure could easily have succeeded [failed].”

*Probability* Participants were asked the question, “Given the surgeon’s action, how likely do you think it was that the procedure would fail [succeed]?” using a 7-point Likert scale ranging from “very unlikely” to “very likely.”

#### 4.4.2 Results

Mean responses for each condition are displayed in Fig. 4.

The goal of the present study was to test if existing theories of luck can account for the effect of normative mismatch on luck attributions. The question, in other words, was whether normative mismatch has an impact on luck attributions independently from its influence on lack of control, modal fragility and probability. To test this, I ran an ANCOVA on luck attributions with Action-Valence (negligent vs. virtuous), Outcome-Valence (good vs. bad) and Vignette (‘Surgeon’ vs. ‘Commander’ vs. ‘Lawyer’) as fixed factors and Control, Modal Fragility, and Probability as covariates. There was indeed a significant Action-Valence  $\times$  Outcome-Valence interaction after controlling for the rest of the variables,  $F(1, 393) = 25.42, p < .001, \eta^2 = 0.06$ . The interaction followed the relation of normative mismatch (see note 15 for details). However, both Control ( $\beta = 0.43, p < .001$ ) and Modal Fragility ( $\beta = 0.14, p = .002$ ) significantly predicted luck attributions after controlling for the normative mismatch effect (i.e., after controlling for the Action-Valence  $\times$  Outcome-Valence interaction).<sup>15</sup>

<sup>15</sup> There was a main effect of Outcome-Valence,  $F(1, 393) = 28.34, p < .001, \eta^2 = 0.02$ , such that participants in the bad conditions were less inclined to attribute luck than those in the good conditions. There was also a main effect of Vignette,  $F(2, 393) = 4.21, p = .01, \eta^2 = 0.02$ . Participants attributed more luck in ‘Lawyer’ than in both ‘Surgeon’ and ‘Commander.’

A significant Outcome-Valence  $\times$  Vignette interaction obtained,  $F(2, 393) = 3.94, p = .02, \eta^2 = 0.01$ . The interaction was due to the lack of an Outcome-Valence effect in the ‘Lawyer’ vignette. As mentioned in the text, there was a significant Action-Valence  $\times$  Outcome-Valence interaction,  $F(1, 393) = 25.42, p < .001, \eta^2 = 0.06$ . The interaction was decomposed to further understand the effect. Participants in the bad outcome conditions were less inclined to attribute luck when the action was negligent ( $M = 2.16, SD = 1.59$ ) than when it was virtuous ( $M = 3.84, SD = 1.92$ ),  $t(393) = 6.01, p < .001, d = 0.95$ , whereas participants in the good outcome conditions showed the opposite pattern. They attributed less luck when the action leading to the good outcome was virtuous ( $M = 3.04, SD = 1.90$ ), than when it was negligent ( $M = 4.44, SD = 1.76$ ),  $t(393) = -5.30, p < .001, d = 0.76$ .

There was a significant Action-Valence  $\times$  Outcome-Valence  $\times$  Vignette interaction,  $F(2, 393) = 3.55, p = .02, \eta^2 = 0.01$ . The Action-Valence  $\times$  Outcome-Valence interaction was significant in both the ‘Surgeon’ and ‘Lawyer’ vignettes (both  $p$ ’s  $< .01$ ), but failed to reach significance in the ‘Commander’ vignette.




Fig. 4 Mean responses by condition for each measure in Study 4 collapsing across vignettes. Error bars indicate standard error of the mean

### 4.4.3 Discussion

Past research has shown that normative considerations impact attributions of control, modal reasoning and probability judgments. This gives rise to the possibility that normative mismatch affects luck attributions only *through* its effects on one of these concepts. For example, it may be that negligent actions lead people to take bad outcomes as more probable, and subsequently to attribute less luck if they obtain. The goal of Study 4 was to test this hypothesis.

The results suggest that normative mismatch has an effect on luck attributions *independently from* its effect on judgments on control, modality and probability. That is, the effect remained significant even after controlling for the other (descriptive) concepts. Of course, the fact that normative mismatch remained significant is not conclusive evidence that it has an independent effect on luck attributions, for different factors (e.g., unreliable measurements) might inflate the odds of a Type I error (Westfall and Yarkoni 2016). Nevertheless, the findings yield support to the idea that luck attributions are directly sensitive to normative considerations. On the other hand, the results also provided evidence that they are sensitive to non-normative factors as well. For judgments of control and modality also had a significant effect after controlling for normative mismatch. The results thus point to a more complex picture of ordinary luck attributions. I explore some of the complexities involved in Sect. 5.

## 4.5 Study 5: luck attributions and causal judgments

The results of Study 4 are evidence that the impact of normative considerations on luck attributions is not explained by any of the central concepts posited by existing theories of luck. But there is another candidate to consider, namely, judgments of causation.

In recent years, there has been a great deal of work on the influence of normative judgments on causal cognition. It has been shown, for example, that attributions of causality are sensitive to whether the agent acted in violation of a norm (Alicke 1992). Indeed, there is evidence that valence matching is an important aspect of causal judgments: people tend to infer good causes from good outcomes and bad causes from bad outcomes (LeBoeuf and Norton 2011).

Ever since the first results came to light, there has been debate concerning the explanation of these findings (Halpern and Hitchcock 2014; Hitchcock and Knobe 2009; Knobe and Fraser 2008; Samland and Waldmann 2016; Shepard and Wolff 2013; Sytsma et al. 2012). Specifically, a consensus has yet to be reached about whether (and to what extent) the influence of normative judgments on attributions of causation shows the latter to be themselves normative.

Regardless, the impact of normative considerations on causal judgments leads to an alternative explanation for the results of studies 1–4. The idea is that luck attributions are just a different way to express judgments of causal efficacy. To judge an agent lucky, that is, is simply to deny that he or she caused the event in question.

This hypothesis has its root in attribution theory, which analyzes luck attributions as attributions of external (i.e., environmental) and unstable causes (see Pritchard and Smith 2004, pp. 6–7 for a discussion of attribution theory and luck). The present study was designed to test this alternative hypothesis.

Crucially, there is an important difference between past research on the relation between normative judgments and causal cognition and the research presented here. As Studies 1–4 show, action-valence, in itself, does not have a consistent effect on luck attributions. Instead, it is the matching relation between the valence of the action and the valence of the outcome that I have argued explains the results. In causal cognition, however, there is ample evidence that norm-violating actions are seen as highly causal even when they lead to good outcomes (see, e.g., the ‘drug vignette’ in Hitchcock and Knobe 2009, pp. 603–604).<sup>16</sup> There is thus reason to think that luck attributions are sensitive to different considerations from those influencing judgments of causation. If this is right, it should be possible to attribute *both* luck and causal efficacy to a single agent. In particular, past research on causal cognition together with the effect of normative mismatch observed in the previous studies, predicts that agents who act in norm-violating ways that result in good outcomes would be seen as both causal and lucky. This would tell against the hypothesis from attribution theory discussed above. The goal of this study was to decide between the two hypotheses.

#### 4.5.1 Method

*Participants* 202 participants (mean age = 36, 44% female) were recruited.

*Procedure* The study used a 2 (Outcome-Valence: good vs. bad) × 2 (Vignette: ‘chemist’ vs. ‘detective’) between-subjects design.

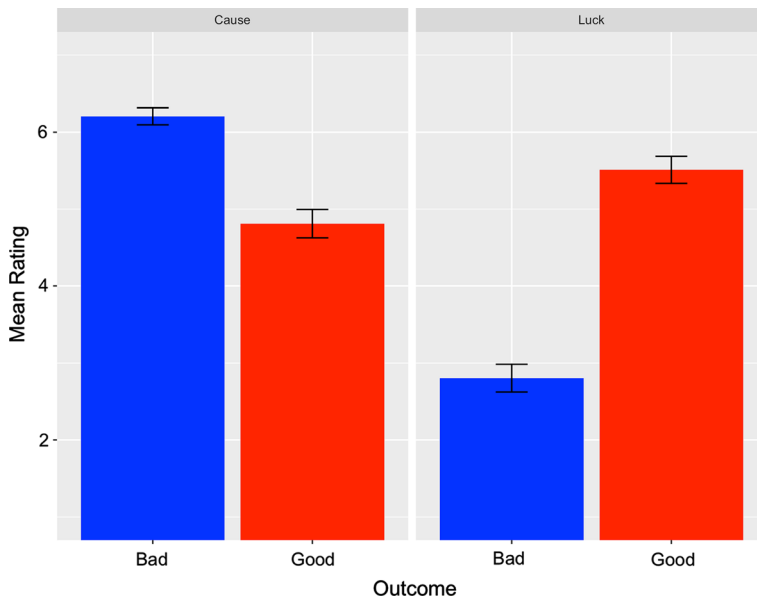
Here is an example of one vignette (brackets indicate the different conditions, see “Appendix” for all vignettes).

A chemist has spent the last months developing a theory—theory X—which, if proven, would be a nice contribution to the field of theoretical chemistry. He believes that he is at the brink of an important discovery. He assembles a team of graduate students to work together in designing an experiment to corroborate theory X. The team works for a couple of weeks and finally comes up with an experiment. The scientist gives the lab’s technician a list with detailed and systematic instructions to carry out the experiment.

The technician has no background in theoretical chemistry and therefore does not even comprehend what the experiment is about. He decides, however, to ignore the chemist’s instructions, for he believes there is an easier way to get the chemical reaction he thinks the experiment is designed to produce. He adds an extra chemical to

<sup>16</sup> There is debate concerning whether, and to what extent, outcome-valence has an impact on causal judgments. Alicke et al. (2011) present evidence suggesting that it does. Their theory of culpable causation posits that actions resulting in bad outcomes are seen as more causal. Recently, however, Cova et al. (2018) failed to replicate their results (the replication was conducted by Ivar Hannikainen). The replication’s results showed an effect of outcome-valence on judgments of blame, but not on causal judgments. For the details of the replication, see <https://osf.io/4yuyml/>.





**Fig. 5** Left: causal attributions. Mean responses by condition collapsing across vignettes. Right: luck attributions. Mean responses by condition collapsing across vignettes. Error bars indicate standard error of the mean

the compound. The experiment does not work and, in fact, there is a huge explosion produced by the chemical the technician added. As it happens, the explosion reveals something about the original compound the chemist had never noticed before. The explosion leads to the discovery the chemist was after. Theory X is proven. [The experiment does not work and, in fact, there is a huge explosion produced by the chemical the technician added. As it happens, the explosion prevents the chemist from making the discovery he was after. Theory X remains unproven].

*Dependent variables* Each participant answered the following questions (the order of the questions was randomized).

*Luck* Participants were asked the question, “To what extent would you describe the technician as [un]lucky”? They answered using a 7-point Likert scale ranging from “not at all” to “great extent.”

*Causation* Participants indicated their agreement with the following statement, “The technician’s action caused [the failure of] the discovery,” using a 7-point Likert scale ranging from “disagree” to “agree.”

#### 4.5.2 Results

Mean responses for each condition are displayed in Fig. 5.

To test whether normative mismatch has an impact on luck attributions independently from its influence on causal judgments, I ran an ANCOVA on luck attributions with Outcome-Valence (good vs. bad) and Vignette (‘Chemist’ vs. ‘Detective’) as fixed factors and Causal Judgments as a covariate. There was a significant main

effect of Outcome-Valence after controlling for Causal Judgments,  $F(1, 198)=98.21$ ,  $p < .001$ ,  $\eta^2=0.33$ . There was also a significant Outcome-Valence  $\times$  Vignette interaction,  $F(1, 198)=5.31$ ,  $p=.02$ ,  $\eta^2=0.02$ . Importantly, the main effect of Outcome-valence was significant in both vignettes (both  $p$ 's  $< .001$ ). The interaction was due to the larger effect size in the 'chemist' vignette ( $d=1.49$ ) compared to the 'detective' vignette ( $d=.96$ ).

### 4.5.3 Discussion

Past research on causal cognition has shown that judgments about whether an action caused an event are influenced by the normative status of the action. This finding suggests that manipulating action-valence in the manner of the present studies affect people's intuitions about the causal efficacy of the agents' actions. An alternative explanation of the results is thus that causal judgments explain luck attributions. Drawing from attribution theory, the idea is that judgments of luck are the converse of causal judgments: attributing luck to an agent is a denial of her causal efficacy, and vice versa.

The results from Study 5, however, are evidence that this is not the case. Even when controlling for causal judgments, luck attributions were sensitive to normative mismatch.<sup>17</sup> Indeed, when an agent acted irresponsibly but the resulting outcome was good, participants attributed luck (as in the previous studies) but also causal efficacy to the agent (as past research has shown). Both judgments were above the midpoint of the scale. Moreover, the design was such that each participant judged both the luckiness of the agent and its causal contribution to the outcome, meaning that participants were willing to say, at the same time, that the agent caused the outcome *and* that he was lucky. It thus seems that luck attributions are not simply denials of (agent) causation.<sup>18</sup>

## 5 The folk concept of luck

The results of the experiments show that normative considerations influence luck attributions. In particular, luck attributions seem to be responsive to whether the valence of the agent's action matches the valence of the resulting outcome. Furthermore, existing accounts of luck seem unable to explain this result. The effect of normative mismatch on luck attributions remains significant even after controlling for lack of control, modal fragility, probability, and causal judgments.

The question arises: what explains this effect? A striking possibility is that nothing else explains it because normative mismatch itself is all there is. In other

<sup>17</sup> Again, it is worth keeping in mind that this result does not *guarantee* that normative mismatch has an independent effect on luck attributions, other factors might be at play (see p. 36).

<sup>18</sup> It is worth noting that the results supported the claim by Alicke et al. (2011) that causal judgments are sensitive to outcome-valence. Contrary to Cova et al.'s (2018) failed replication, irresponsible agents were rated as more causal when the outcome was bad than when it was good. More research is needed to elucidate the conditions in which the valence of the outcome has an effect on causal judgments.

words, it may turn out that the correct account of luck is one in which normative mismatch plays the central role. On such an account, luck attributions are not merely sensitive to normative considerations, but the latter are what the former are all about. On this view, luck attributions are not descriptions of, say, an objective condition of chance in the world. Rather, they are evaluations tracking a normative mismatch between the valence of an action and that of its result.

Is this a plausible alternative to the existing accounts? Is the mismatch relation all there is to our folk concept of luck? The experiments suggest a negative answer to both questions. In Study 2, luck attributions were sensitive to probabilities independently of the normative mismatch effect. And in Study 4, both lack of control and modal fragility had an impact after controlling for normative mismatch. It seems, therefore, that all these concepts are part of our folk concept of luck. Moreover, luck attributions across all five studies were quite low, falling below the midpoint of the scale in the majority of conditions. This suggests that other factors beyond normative mismatch are involved when we attribute luck to an event.

So we can rule out two of the possibilities raised in Sect. 3. Normative mismatch can neither explain the relevance of control, modality and probability nor can the latter fully explain the former. We are in turn left with two different pictures of what the folk concept of luck might look like.

On the first picture, there is a correct theory that, once understood, would explain the relevance of the concepts discussed so far (i.e., control, modality, probability and normative mismatch). On this view, one concept—call it Concept X—lies at the heart of luck attributions and the effect of other factors is explained in terms of it.

A different approach rejects the existence of such a theory. According to the second picture, there is no need to posit a Concept X—no need to explain everything in terms of one thing. Rather, there are different factors that are all independently relevant for ordinary luck attributions. This approach views the normative mismatch effect as one among various factors that go into determining whether an event was due to luck, and thus takes existing accounts of luck as *incomplete* (rather than unequivocally wrong). Further work is needed to elucidate the factors that are in fact relevant and how they interact among them to form our folk concept of luck.

## 5.1 Normative mismatch: unresolved questions

Suppose something like the second picture is right and normative mismatch is one of the relevant factors in ordinary luck attributions. Still, even if this much is granted, more work is needed to give a more precise account of what it is. In the rest of this section, I present two kinds of cases that shed light on the issues that need to be resolved in order to give a complete description of the normative mismatch effect.

### 5.1.1 Match and luck

One obvious problem is that there seem to be cases in which there is normative *match* and yet we are inclined to attribute luck. This is best illustrated by standard Gettier cases, where an agent seems to get lucky in *justifiably* forming a *true* belief.

What this shows, I think, is not that the mismatch relation is irrelevant for luck attributions, but that we need a better understanding of such a relation. Gettier cases illustrate that a mere (mis)match of valences is insufficient. *How* the valences between action and outcome are connected seems to be essential for the relation to hold. A possible approach would thus be to deny that Gettier cases are instances of the relevant normative relation. Perhaps the normative match relation, properly understood, requires the outcome to occur *as a result of* the agent's epistemic competence, rather than in spite of, or coincidental with it.<sup>19</sup>

The lesson from Gettier cases is that a fuller characterization of normative mismatch is needed. In particular, they reveal the need to go beyond the mismatch of valences and pay attention to whether (and if so in what ways) the action explains (or brings about) the outcome.

### 5.1.2 Determining the action's valence

Suppose you contrive a very elaborate lie in a job interview and get the job at least in part because of it. Was your success due to luck? The answer appears to be in the negative, despite the fact that the valence of the action (lying) does not match the valence of the outcome (succeeding at a job application). This is only true, however, from the point of view of morality. For it seems clear that, from that of personal interest, constructing a carefully designed lie and getting the job both have positive valences. Indeed, it is a familiar fact that an action considered bad from the standpoint of morality may be good from the standpoint of personal interest (the same, of course, is true of outcomes). More generally, it is often the case that an agent's action has a positive valence when evaluated from one standpoint, but negative under another. And thus, what may seem as a case of normative mismatch from one perspective may be a case of match under another. This creates a problem, for it illustrates that positing normative mismatch is often not enough to derive a concrete prediction regarding which events would be seen as involving luck. We still need a way to specify which valence is the relevant one for luck attributions.

## 6 Implications for ethics and epistemology

I want to close by discussing some of the implications that follow if we recognize normative mismatch as an important part of our folk concept of luck. Recall the discussion towards the end in Sect. 1. The puzzle of moral luck is best understood as a

<sup>19</sup> As Sosa (2007) notes, in Gettier cases, an agent's epistemic virtues explain the *existence* of a given belief, but they don't explain the *truth* of the belief.

puzzle about consequences. On the one hand, when two agents act in the same way, we think they should be assessed equally—regardless of the consequences of their actions. On the other hand, we think that agents who actually caused harm deserve more blame and punishment. As Cushman and Young (2009) put it, “it seems perverse to let the *unlucky* homicidal driver off with nothing more than a ‘driving under the influence’ charge, or to punish the *lucky* driver as if he had killed” (Cushman and Young 2009, p. 12, emphasis mine). The incompatibility thesis supports the first claim: it argues that (a) resultant luck is incompatible with blameworthiness and thus that (b) the drunk driver who kills a passerby deserves to be treated just as any other drunk driver. Ordinary intuition, however, supports the second claim: people do in fact assign more blame to the homicidal driver (Cushman 2008).

Is there a way to reconcile the two claims? Notice what our results imply for the puzzle: they suggest that when negligent actions result in bad outcomes, the impact of the matching relation trumps that of probabilities. The upshot is that people do not think that the drunk driver who hits a passerby is unlucky. That is, they do not treat the accident as a case of resultant luck. If we no longer describe the homicidal driver as *unlucky*—if we take luck out of the picture—then there is no luck to accommodate within our moral assessment, and our judgments of blame and punishment can proceed without regard to the incompatibility thesis. In other words, the claims that (1) resultant luck is incompatible with blameworthiness and (2) the homicidal driver deserves blame are no longer in tension.<sup>20</sup>

A similar dialectic ensues in epistemology. Greco (2003) has argued that since ordinary inductive knowledge involves veritic luck, any epistemology supporting the incompatibility thesis has to explain how inductive knowledge is possible (or else to get rid of inductive knowledge altogether, a high cost to pay). Greco concedes that veritic luck is sometimes incompatible with knowledge, as in Gettier cases, but argues that sometimes it clearly isn’t, as the case of inductive knowledge shows. This creates a problem, since “any plausible epistemology must somehow explain two powerful yet incompatible intuitions: (a) that veritic luck is incompatible with knowledge, and (b) that there can be inductive knowledge, even with veritic luck” (Greco 2003, pp. 361–362). The problem seems nicely resolved by the normative mismatch effect. In cases where knowledge results from an inductive inference reflecting the agent’s epistemic virtues, the resulting belief is not lucky. These are paradigmatic cases of normative match. If this is right, it follows that people would not consider inductive knowledge as an instance of knowledge *and* luck, but rather an instance where luck is out of the equation. Here again, a puzzle arises by positing

<sup>20</sup> What this shows is that blaming the homicidal driver is not in conflict with the incompatibility thesis because the outcome involves no luck. This of course does not, in any way, solve the problem of moral luck. For the problem is whether the two drivers deserve the *same* blame. As a matter of fact, I think the problem can be formulated without mentioning the term ‘luck’ at all. For example, if two persons shoot somebody in the same way and for the same reasons, but one of them survives because the ambulance arrives faster, don’t they deserve to be blamed in the same manner, since the time the ambulance arrives has nothing to do with them? The point is that these sorts of questions remain *prima facie* compelling, even without talking about luck. Thanks to Dan Greco for discussion on this point. The example is Greco’s.

luck where there is none. It is a project for future research to further elucidate how the folk conception of luck interacts with ordinary knowledge attributions.

The results also help us dissolve a common assumption made in the literature, namely, that whenever an agent lacks control, luck judgments are symmetric across outcomes. Call this the ‘symmetry assumption.’ As the above quote from Cushman and Young exemplify, there is a temptation to move from the intuition that the first driver is lucky, to the claim that the second one is unlucky, i.e., to state that luck stays constant across cases and the valence is all that changes (good vs. bad luck). This symmetry in luck attributions gives rise to the question of why judgments of blame shouldn’t be similarly symmetrical. This, of course, is the puzzle of moral luck. But if people think about luck as the results suggest they do, no such symmetry in luck attributions exists. That is, people only attribute luck to the first driver. The results show that ordinary luck attributions are *not* symmetric in cases of negligence, and this may perhaps explain why attributions of blame are not symmetric as well. *Of course, the pressing problem is not whether the homicidal driver deserves blame, but whether she deserves more blame than the driver who harmed no one.* The suggestion is that the asymmetry in blame judgments may be partly explained by the asymmetry in luck judgments that we see in the experiments. It is for further research to test this idea.

Interestingly, the symmetry assumption is sometimes made in epistemology. Statman (1991) relies on it in making his case that most of our beliefs are lucky. Suppose Bob believes truly that his neighbor is schizophrenic and ought to be treated by a psychiatrist. Had he lived a few centuries earlier, Bob would have believed that his neighbor is the victim of the gods’ wrath. Statman concludes, “Hence, the fact that one holds certain beliefs about reality and not others, is a matter of luck. If they happen to be true, one has good luck; if they happen to be false, one has bad luck” (Statman 1991, p. 148). Here again the normative mismatch effect denies the existence of this symmetry, and by extension, Statman’s awkward conclusion. If Bob’s belief is the result of the proper exercise of his epistemic virtues, then there is no relevant sense in which he is lucky to believe truly.

The symmetry assumption is also employed in discussions of justification. Engel (1992) makes great use of it in his exposition of the new evil demon problem. In a demon world, my epistemic twin has the exact same mental states that I have, but he is systematically deceived about the external world. On an internalist framework, if I am justified in believing  $p$ , then so is my twin. Poor twin Mario is unlucky to be in a demon world, and hence, by the symmetry assumption, I must be lucky to be in *this* world. Engel writes, “Obviously, it is simply a matter of luck that my justified beliefs are true, for it is simply a matter of luck that I am in this world” (Engel 1992, p. 61). It is only obvious, however, if one has a theory of luck that supports the symmetry assumption. But if the results are to be trusted, we don’t think of luck in this way, and so we have reason to reject the symmetry and thereby to reject Engel’s conclusion that luck is present when our justified beliefs happen to be true. The upshot is that luck is not as ubiquitous and problematic as Statman and Engel would have us believe.<sup>21</sup>

<sup>21</sup> Engel doesn’t think that the type of luck involved in the new evil demon is problematic. He merely thinks that it is the type of luck that we must recognize as compatible with knowledge. My point is that we don’t even have to recognize it as luck in the first place.

The normative mismatch effect, then, is well suited to explain away some of the counterintuitive conclusions that philosophers have defended using a notion of luck that differs from the ordinary one.

**Acknowledgements** I want to thank Stephen Darwall, Keith DeRose, Dan Greco, Shaun Nichols, Duncan Pritchard, Anat Shor, and three anonymous referees for their insightful feedback on earlier drafts. I am also grateful to members of Josh Knobe's Experimental Philosophy Lab at Yale for stimulating discussion and advice. Above all, I thank Josh Knobe for his tremendous help and constant encouragement.

## **Compliance with ethical standards**

**Conflict of interest** The author declares that there are no conflicts of interest.

## **Appendix**

### **Vignettes used in Study 1**

#### **Commander**

##### **Negligent action**

About to perform a military mission to rescue a group of hostages, the commander of the team decides to ignore crucial new information that would have altered the original plan. As a result, the chances of a failed mission rise to about 30%.

As a matter of fact, the mission fails [does not fail].

##### **Virtuous action**

About to perform a military mission to rescue a group of hostages, the commander of the team decides to alter the original plan in light of crucial new information. As a result, the chances of a failed mission drop to about 30%.

As a matter of fact, the mission fails [does not fail].

#### **Lawyer**

##### **Negligent action**

The night before defending an innocent person in court, a lawyer decides to watch TV instead of preparing the case. As a result, the chances of the lawyer losing the case rise to about 30%.

As a matter of fact, the lawyer loses [does not lose] the case.

**Virtuous action**

The night before defending an innocent person in court, a lawyer decides to go over each and every piece of evidence in order to prepare the best possible case. As a result, the chances of the lawyer losing the case drop to about 30%.

As a matter of fact, the lawyer loses [does not lose] the case.

**Vignettes used in Study 4****Commander****Negligent action**

About to perform a military mission to rescue a group of hostages, the commander of the team decides to ignore crucial new information that would have altered the original plan. As a result, the chances of a failed [successful] mission rise [drop] to about 30%.

As a matter of fact, the mission fails [succeeds].

**Virtuous action**

About to perform a military mission to rescue a group of hostages, the commander of the team decides to alter the original plan in light of crucial new information. As a result, the chances of a failed [successful] mission drop [rise] to about 30%.

As a matter of fact, the mission fails [succeeds].

**Lawyer****Negligent action**

The night before defending an innocent person in court, a lawyer decides to watch TV instead of preparing the case. As a result, the chances of the lawyer losing [winning] the case rise [drop] to about 30%.

As a matter of fact, the lawyer loses [wins] the case.

**Virtuous action**

The night before defending an innocent person in court, a lawyer decides to go over each and every piece of evidence in order to prepare the best possible case. As a result, the chances of the lawyer losing [winning] the case drop [rise] to about 30%.

As a matter of fact, the lawyer loses [wins] the case.



## Vignettes used in Study 5

### Detective

A detective has spent the last couple of years working on a case involving an unsolved murder. He believes that he is at the brink of an important discovery. After many hours of interviews with people who knew the victims, and testimonies from those who were at the scene of the crime, the detective narrows down the list of suspects to ten people. He asks to run DNA tests on all ten suspects to finally determine who committed the crime. The district attorney approves the requests and DNA tests are ordered.

The scientist in charge of the DNA tests, who does not know anything about the crime, is tired and sleepy at the moment of analyzing the results. He does not double-check the analyses or makes sure the samples are not contaminated, as is common procedure. He messes up the results and mistakenly informs the detective that sample number 2 matches the target DNA. The detective thus arrests suspect number two. As it happens, suspect number two is the only person who knows who committed the crime and, at the prospect of being charged with the murder himself, tells the detective who the murderer is. The detective thus solves the crime. [He messes up the results and mistakenly informs the detective that sample number 2 matches the target DNA. The detective thus arrests suspect number two. After extensive questioning, it becomes clear that he is innocent. The detective thus fails to solve the crime].

### References

- Alicke, M. D. (1992). Culpable causation. *Journal of Personality and Social Psychology*, 63(3), 368–378.
- Alicke, M. D. (2000). Culpable control and the psychology of blame. *Psychological Bulletin*, 126(4), 556–574.
- Alicke, M. D., Rose, D., & Bloom, D. (2011). Causation, norm violation, and culpable control. *The Journal of Philosophy*, 108(12), 670–696.
- Axtell, G. (2001). Epistemic luck in light of the virtues. In A. Fairweather & L. Zagzebski (Eds.), *Virtue epistemology: Essays in epistemic virtue and responsibility*. Oxford: Oxford University Press.
- Baumann, P. (2014). No luck with knowledge? On a dogma of epistemology. *Philosophy and Phenomenological Research*, 89(3), 523–551.
- Broncano-Berrocal, F. (2015). Luck as risk and the lack of control account of luck. *Metaphilosophy*, 46(1), 1–25.
- Coffman, E. J. (2007). Thinking about luck. *Synthese*, 158(3), 385–398.
- Cova, F., Strickland, B., Abatista, A., Allard, A., Andow, J., Attie, M., et al. (2018). Estimating the reproducibility of experimental philosophy. *Review of Philosophy and Psychology*. <https://doi.org/10.1007/s13164-018-0400-9>.
- Cushman, F. (2008). Crime and punishment: distinguishing the roles of causal and intentional analyses in moral judgment. *Cognition*, 108(2), 353–380.
- Cushman, F., & Young, L. (2009). The psychology of dilemmas and the philosophy of morality. *Ethical Theory and Moral Practice*, 12(1), 9–24.
- Dalbauer, N., & Hergovich, A. (2013). Is what is worse more likely?—The probabilistic explanation of the epistemic side-effect effect. *Review of Philosophy and Psychology*, 4(4), 639–657.
- Doris, J. M. (2002). *Lack of character: Personality and moral behavior*. Cambridge: Cambridge University Press.

- Engel, M. (1992). Is epistemic luck compatible with knowledge? *The Southern Journal of Philosophy*, 30(2), 59–75.
- Gettier, E. L. (1963). Is justified true belief knowledge? *Analysis*, 23(6), 121–123.
- Greco, J. (2003). Virtue and luck, epistemic and otherwise. *Metaphilosophy*, 34(3), 353–366.
- Halpern, J. Y., & Hitchcock, C. (2014). Graded causation and defaults. *The British Journal for the Philosophy of Science*, 66(2), 413–457.
- Harman, G. (1973). *Thought*. Princeton: Princeton University Press.
- Harman, G. (2003). No character or personality. *Business Ethics Quarterly*, 13(1), 87–94.
- Hitchcock, C., & Knobe, J. (2009). Cause and norm. *The Journal of Philosophy*, 106(11), 587–612.
- Kamtekar, R. (2004). Situationism and virtue ethics on the content of our character. *Ethics*, 114(3), 458–491.
- Kneer, M., & Machery, E. (2019). No luck for moral luck. *Cognition*, 182, 331–348.
- Knobe, J., & Fraser, B. (2008). Causal judgment and moral judgment: Two experiments. *Moral Psychology*, 2, 441–448.
- Lackey, J. (2008). What luck is not. *Australasian Journal of Philosophy*, 86(2), 255–267.
- Latus, A. (2000). Moral and epistemic luck. *Journal of Philosophical Research*, 25, 149–172.
- LeBoeuf, R. A., & Norton, M. I. (2011). Consequence-cause matching: Looking to the consequences of events to infer their causes. *Journal of Consumer Research*, 39(1), 128–141.
- Levy, N. (2009). What, and where, luck is: A response to Jennifer Lackey. *Australasian Journal of Philosophy*, 87(3), 489–497.
- Mitchell, A. (2017, February 12). Let Milo Speak. Retrieved from <https://stanfordreview.org/let-milo-speak-a7276e75fa4/>. Accessed 10 Dec 2017.
- Nagel, T. (1979). *Mortal questions*. Cambridge: Cambridge University Press.
- Phillips, J., & Cushman, F. (2017). Morality constrains the default representation of what is possible. *Proceedings of the National Academy of Sciences*, 114(18), 4649–4654.
- Pritchard, D. (2005). *Epistemic luck*. Oxford: Clarendon Press.
- Pritchard, D. (2008). Virtue epistemology and epistemic luck, revisited. *Metaphilosophy*, 39(1), 66–88.
- Pritchard, D. (2015). The modal account of luck. In D. Pritchard & J. L. Whittington (Eds.), *The philosophy of luck* (pp. 143–167). Hoboken: Wiley.
- Pritchard, D., & Smith, M. (2004). The psychology and philosophy of luck. *New Ideas in Psychology*, 22(1), 1–28.
- Richards, N. (1993). Luck and desert. In D. Statman (Ed.), *Moral luck*. Albany: State University of New York Press.
- Riggs, W. (2009). Luck, knowledge, and control. In: A. Haddock, A. Millar, D. Pritchard (Eds.), *Epistemic value* (pp. 204–221). Oxford: Oxford University Press.
- Samland, J., & Waldmann, M. R. (2016). How prescriptive norms influence causal inferences. *Cognition*, 156, 164–176.
- Shepard, J., & Wolff, P. (2013). Intentionality, evaluative judgments, and causal structure. In: M. Knauff, M. Pauen, N. Sebanz, I. Wachsmuth (Eds.), *Proceedings of the 35th annual conference of the cognitive science society* (pp. 3390–3395).
- Sosa, E. (2007). *A virtue epistemology: Volume I: Apt belief and reflective knowledge*. Oxford: Oxford University Press.
- Srinivasan, A. (2015). Normativity without Cartesian privilege. *Philosophical Issues*, 25(1), 273–299.
- Statman, D. (1991). Moral and epistemic luck. *Ratio*, 4(2), 146–156.
- Statman, D. (Ed.). (1993). *Moral luck*. Oxford: Oxford University Press.
- Sytsma, J., Livengood, J., & Rose, D. (2012). Two types of typicality: Rethinking the role of statistical typicality in ordinary causal attributions. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 43(4), 814–820.
- UC Berkeley Faculty. (2017, January 10). Open letters calling for cancellation of Milo Yiannopoulos event. Retrieved from <http://www.dailycal.org/2017/01/10/open-letter-calling-cancellation-milo-yiannopolous-event/>. Accessed 10 Dec 2017.
- Uttich, K., & Lombrozo, T. (2010). Norms inform mental state ascriptions: A rational explanation for the side-effect effect. *Cognition*, 116(1), 87–100.
- Vahid, H. (2001). Knowledge and varieties of epistemic luck. *Dialectica*, 55(4), 351–362.
- Westfall, J., & Yarkoni, T. (2016). Statistically controlling for confounding constructs is harder than you think. *PLoS ONE*, 11(3), e0152719.

- Whittington, L. J. (2015). Getting moral luck right. In D. Pritchard & J. L. Whittington (Eds.), *The philosophy of luck* (pp. 205–218). Hoboken: Wiley.
- Williams, B. (1981). *Moral luck: philosophical papers 1973–1980*. Cambridge: Cambridge University Press.
- Young, L., Nichols, S., & Saxe, R. (2010). Investigating the neural and cognitive basis of moral luck: It's not what you do but what you know. *Review of Philosophy and Psychology*, 1(3), 333–349.
- Zimmerman, M. (1993). Luck and moral responsibility. In D. Statman (Ed.), *Moral luck* (pp. 217–234). Oxford: Oxford University Press.

**Publisher's Note** Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.