Collective Epistemic Luck

Moisés Barba¹ · Fernando Broncano-Berrocal²

Received: 17 October 2020 / Accepted: 5 August 2021 / Published online: 28 October 2021

© The Author(s) 2021, corrected publication 2022

Abstract

A platitude in epistemology is that an individual's belief does not qualify as knowledge if it is true by luck. Individuals, however, are not the only bearers of knowledge. Many epistemologists agree that groups can also possess knowledge in a way that is genuinely collective. If groups can know, it is natural to think that, just as true individual beliefs fall short of knowledge due to individual epistemic luck, true collective beliefs may fall short of knowledge because of *collective epistemic luck*. This paper argues, first, that the dominant view of epistemic luck in the literature, the modal view, does not yield a satisfactory account of lucky collective beliefs. Second, it argues that collective epistemic luck is better explained in terms of groups lacking (suitably defined) forms of control over collective belief formation that are specific to the different procedures for forming collective beliefs. One of the main implications of this, we will argue, is that groups whose beliefs are formed via internal deliberation are more vulnerable to knowledge-undermining collective luck than groups that form their beliefs via non-deliberative methods, such as non-deliberative anonymous voting. The bottom line is that the greater exposure to knowledge-undermining luck that deliberation gives rise to provides *a* reason (not a conclusive one) for thinking that non-deliberative methods of group belief formation have greater epistemic value.

✉ Fernando Broncano-Berrocal
f.broncano-berrocal@ub.edu

Moisés Barba
moises.barmag@gmail.com

¹ Autonomous University of Madrid, Madrid, Spain

² University of Barcelona – LOGOS, Barcelona, Spain

1 Collective Epistemic Luck and Collective Gettier Cases

Attributing knowledge to groups is a common phenomenon in ordinary parlance. For instance, we say that a political party knows that their voters are left-wing voters, that a firm knows that its gross margin is 3 million, or that a hiring committee knows that a particular applicant is qualified.¹ Many epistemologists consider such knowledge attributions meaningful, to the point that they believe groups can have knowledge in a way that is genuinely collective.² Here are two examples of group belief formation, the output of which (a collective belief) can be properly considered group knowledge:

Hikers

A group of four hikers, wanting to reach the top of the mountain in the shortest possible time, try to settle whether path A, B, C, or D is the quickest route to the top. Each has, respectively, walked one (and only one) of those four paths in the past. Therefore, for each path, there is only one hiker that *knows* how long it takes to walk it. To decide what path is the fastest, the hikers engage in group discussion revealing their private evidence, and they collectively conclude, correctly, that A is the fastest path.

Hiring committee

A hiring committee needs to decide who is the most qualified candidate for a vacant teaching position. Being very competent, all committee members meticulously revise the candidates' CVs, and all interview each candidate personally, one-to-one. After that, each makes a knowledgeable assessment of each candidate and a corresponding competent choice. The committee's rules stipulate that committee members can neither deliberate nor reveal their individual assessments or choices to other members. Instead, they must take an anonymous vote. The outcome of the voting is unanimous: candidate 6 is the most qualified for the job. This is the case, indeed. The committee's spokesperson notifies the Dean that the hiring committee believes that candidate 6 is the most qualified, and the Dean appoints that person.

Intuitively, the group of hikers *knows* that A is the fastest path. Similarly, the hiring committee *knows* that candidate 6 is the most qualified for the job. Now, if groups can have knowledge, it is natural to think that, just as true individual beliefs fall short of knowledge because of *individual epistemic luck* (e.g., in Gettier cases),³ by parity of reasoning, true collective beliefs may fail to qualify as knowledge

¹ The forthcoming discussion is restricted to collective *propositional* knowledge.

² For discussion of collective knowledge, see Bird 2010; Gilbert 1987; 2002; 2013; Mathiesen 2006; Tollefsen 2006; Tuomela 1995; and Wray 2007.

³ Here is one such case. At 10 pm, Bertrand looks at a generally reliable digital clock that says "22:00" and forms the true belief that it is 10 pm on that basis. Unbeknown to Bertrand, however, the clock's batteries ran out yesterday exactly at 10 pm. Thus, Bertrand's belief is true by luck, but does not amount to knowledge. See Russell (1948/2009).

because of *collective epistemic luck*.⁴ Here are some examples of knowledge-undermining collective luck:

Lucky Hikers

A group of four hikers, wanting to reach the top of the mountain in the shortest possible time, try to settle whether path A, B, C, or D is the quickest route to the top. Each *believes* to have walked one (and only one) of those four paths in the past. Therefore, for each path, there is only one hiker that *believes* how long it takes to walk it. To decide what path is the fastest, the hikers engage in group discussion revealing their private evidence, and they collectively conclude, correctly, that A is the fastest path.

This time, however, the hikers' epistemic situation is more complicated. Hikers 2, 3, and 4 did, respectively, walk paths B, C, and D, and each not only believes but also *knows* the time it took them: respectively, 2, 3, and 4 hours. By contrast, hiker 1 makes a *double mistake*. First, she claims to have walked path A, but she's misremembering and in reality took path B (the one hiker 2 walked). Second, she incorrectly measured how long it took her to walk path B. It normally takes 2 hours, but she measured 1h 30 min, which is, coincidentally, the time it takes to walk A. Now, here is the *lucky twist*: from her false belief that she walked path A (hiker 1 walked B), and from her false belief that it took her 1h 30 min (it took her two hours), she comes to believe that it takes 1h 30 min to walk path A. As things turn out, this belief is true, but it's just a coincidence that it is. During group deliberation, hiker 1 communicates her belief to the others, and they collectively agree, correctly, that path A is the fastest.

Lucky Hiring Committee

A hiring committee needs to decide who is the most qualified candidate for a vacant teaching position. Each of their gender-biased members revises the candidates' CVs in a sloppy, prejudiced way and interviews each candidate personally, one-to-one, in an unfair and often dismissive manner. After that, each makes an individual biased assessment of the candidates and a corresponding arbitrary choice. The committee's rules stipulate that committee members can neither deliberate nor reveal their individual assessments or choices to other members. Instead, they must take an anonymous vote.

As things turn out, all committee members vote for candidate 9, the only male candidate who is also utterly unqualified for the job. To vote, they just need to press "9" on a numerical keyboard. Unbeknown to them, however, the janitor

⁴ See Carter (2015) and Lackey (2014) for a parallel argument in the case of epistemic defeat: just as individual knowledge can be defeated, they argue, so it should be in the case of collective knowledge. In general, it is to be expected that collective knowledge is subject to the same kind of epistemic phenomena (e.g., epistemic luck, defeat, etc.) as individual knowledge. This doesn't mean that these collective phenomena need to work and hence be spelled out in the same way as their corresponding individual phenomena. In fact, as we will see, collective epistemic luck, although the same kind of phenomenon as individual epistemic luck (they're both instances of luck), is different in important respects that need to be spelled out with different theoretical resources. This is, at any rate, one of the key points of the paper.

had accidentally turned keys “6” and “9” upside down the day before when cleaning the keyboard. Consequently, when committee members anonymously and unanimously press key “9” to vote for candidate “9,” the computerized voting system registers “6” and thereby casts a vote for candidate 6, a female candidate who also happens to be the most qualified for the job. In this way, the voting results in the committee’s unanimous true opinion that candidate 6 is the most qualified. The committee’s spokesperson notifies this verdict to the Dean, and the Dean appoints that person.

The hikers’ collective belief that path A is the fastest is true by luck. After all, the individual belief (hiker 1’s) upon which the hikers reach consensus is also true by luck. Collective epistemic luck is thus the reason the group ignores that proposition. It should also be clear that the hiring committee, despite holding a true collective opinion about who’s the most qualified candidate, ignores this. The culprit, again, is knowledge-undermining collective luck. While in the former case, the luck in play is due to an individual group member’s lucky true belief being the basis for the group’s view, in the latter, luck intervenes in the procedure for aggregating individual beliefs (the voting system). In both cases, the upshot is the same: group knowledge is undermined by the epistemic luck in play. To better grasp the structure of the cases, let’s compare them in Table 1 with the previous examples of group knowledge:

Granted, as *Lucky Hikers* and *Lucky Hiring Committee* exemplify, groups may fail to possess collective knowledge because of the phenomenon we’ve called collective epistemic luck. But there’s a potential objection lurking around to calling such a phenomenon collective “epistemic luck.” As the objection goes, Gettier cases are paradigmatic cases of individual epistemic luck; if collective epistemic luck is the same kind of phenomenon as individual epistemic luck, there should be Gettier cases at the collective level too, but the proposed cases do not look like Gettier cases. After all, in *Lucky Hikers* hiker 1’s true belief that it takes 1 h 30 min to walk path A is *not* epistemically justified—recall: this is the belief upon which the group’s true belief is based. If so, one might argue that the lack of justification of hiker 1’s belief renders such a collective belief unjustified. But *collective Gettier cases* must surely be cases of *justified* collective belief. Analogously, in *Lucky Hiring Committee*, the voting system is unreliable. If individual beliefs are aggregated unreliably, the resulting collective beliefs are unjustified (on a reliabilist conception of justified group belief).⁵ The same reasoning applies: it’s not a Gettier case either.

As a rejoinder, these factors—the lack of justification of hiker 1’s individual belief and the voting system’s unreliability—are contingent features of the cases. The preceding examples can be easily turned into collective Gettier cases by making such factors epistemically appropriate while retaining the intuition that the relevant collective beliefs are lucky. Consider the following versions of *Lucky Hikers* and *Lucky Hiring Committee*:

Gettierized Hikers

⁵ See Goldman (2014) and Lackey (2016) for recent work on justified group belief.

Everything is like in *Lucky Hikers* except for the fact that when hiker 1 walked path B, she had very good justification for thinking that it was path A: everyone, including several rangers, Google Maps, and the National Park Guide, referred to that path as “A.” Besides, the hiker’s very reliable stopwatch was, unbeknownst to her, suddenly restarted when exactly 1h 30min of walking was left—e.g., unbeknownst to the hiker, a squirrel accidentally dropped a nut that bounced on her stopwatch’s on/off button two times (one to switch it off, one to switch it on) exactly at that time.

Gettierized Hiring Committee

Everything is like in *Lucky Hiring Committee*, except for the fact that the voting system is reliable: the keys of the keyboard are fine and nothing amiss happens with the way votes are aggregated. However, the *output* of the reliable voting procedure is rigged by a feminist activist to ensure that the committee ends up making a fair decision by reaching the true collective judgment that candidate 6 (a woman) is the most qualified person for the job.

Therefore, not only collective epistemic luck is a genuine phenomenon: there are also collective Gettier cases. Now, from the fact that true justified collective beliefs can be Gettierized and, more generally, from the fact that luck can prevent a group’s true belief from qualifying as knowledge, we should not so readily assume that the conditions that give rise to epistemic luck at the collective level are the same as those at the individual level. At best, the previous cases give *prima facie* reason to think that individual and collective epistemic luck are the same *kind* of epistemic phenomenon, in that they both involve luck (rather than, say, defeat, or lack of luminosity for that matter). However, albeit being the same kind of phenomenon, we’ll argue that they behave differently, epistemically speaking. Not only that, one of the main points in our argumentation will be that salient differences between individual and collective belief formation call for a radically different way of theorizing epistemic luck at the collective level: instead of modeling it in modal terms, as per usual in the individual case, collective epistemic luck is best understood in terms of groups lacking (suitably defined) forms of control over collective belief formation that are specific to the different procedures for forming collective beliefs.

Interestingly, one of the main implications of the proposed account will be that groups whose beliefs are formed via internal deliberation (i.e., the process by which group members share their views and evidence with other members and weigh them up with the views and evidence shared by others) are more liable to knowledge-undermining collective luck than groups that form their beliefs via non-deliberative methods (such as non-deliberative majority voting). In fact, as we will argue, this greater exposure to knowledge-undermining luck that deliberation gives rise to can be seen as providing *a* reason (not a conclusive one, though) for thinking that non-deliberative methods of group belief formation have greater epistemic value.

Table 1 Examples

	<i>Hikers</i>	<i>Hiring committee</i>	<i>Lucky hikers</i>	<i>Lucky hiring committee</i>
Does the group have knowledge?	Yes	Yes	No	No
Do group members have knowledge?	All	All	All but one	None
Are the group members' beliefs true?	Yes	Yes	Yes	No
Method of collective belief formation	Deliberation	Voting	Deliberation	Voting

2 Two Competing Accounts of Luck

Luck is present in our everyday lives. It is also a notion that plays a significant role in several areas of philosophy, including philosophy of action, moral philosophy, and epistemology. However, the nature of luck itself has not been a traditional focus of philosophical theorizing until recently. Fortunately, the situation has been amended, and we now count with several views about the nature of luck. Two popular ones are these:

The modal account of luck: An event E is lucky for an agent S if and only if E occurs in the actual world but fails to occur in most nearby possible worlds in which the relevant initial conditions for E are the same as in the actual world.⁶

The lack of control account of luck: An event E is lucky for an agent S if and only if E is beyond S's control.⁷

Both accounts get paradigmatic instances of luck right. For example, winning a fair lottery is typically a matter of luck because (as per the modal account) in most nearby possible worlds one would lose or because (as per the lack of control account), since the lottery is fair, one lacks control over its outputs. Yet, counterexamples have been proposed to both accounts, replies have been given and refinements made, and further versions have been proposed.⁸ Here is not the place to adjudicate between these two competing views. For the purposes here, we'll only mention a critical problem with the modal view. The reason is that this problem is pervasive, in that it spreads to the modal account both of individual epistemic luck (§3) and of collective epistemic luck (§4)—throughout different contexts. This problem thus paves the way for a more reasonable lack of control approach to the phenomenon.

⁶ See Pritchard (2005) for seminal work on the notions of luck and epistemic luck and Pritchard (2014) for a more recent defense of the modal account.

⁷ Most views also include a condition to the effect that the event must be significant to the agent for it to be lucky. See Ballantyne (2012) for several ways of understanding the significance condition of luck. See Broncano-Berrocal (2015) for a defense of the lack of control account of luck.

⁸ See Broncano-Berrocal (2016) for an overview of the different theories of luck as well as the main objections that have been proposed to each of them.

The problem starts with the simple idea that not all lucky events arise out of *coincidence* (we can hence call the problem *the coincidence problem*). If someone tosses a coin and lands tails, it is not a coincidence that it lands tails: after all, someone has tossed the coin intentionally. Yet, that outcome can still count as lucky for someone because, e.g., that person will win a big prize if the coin lands tails. Coincidences, by contrast, are necessarily lucky (if significant).⁹ So coincidences are a subset, perhaps a significant one, of lucky events. As Owens (1992) characterizes the notion, a coincidence is an event that arises out of two or more constituent events which are themselves produced by independent causal factors in such a way that we cannot explain why they come together—e.g., there is no nomological connection between the components of the coincidence or no common nomological antecedent. For example, you desperately need water, you pray that it rains and it rains, but there is no common factor connecting your prayers with the rain. That the two events come together is simply inexplicable, hence a coincidence.

Let's add more details to the case. You are not religious, so you normally don't pray. This means that, although you have prayed (out of desperation), you could *easily* have done otherwise. In addition, it wasn't very likely that it would rain that day, so although it did rain in the end, the sky could *easily* have stayed clear. This extra layer of complexity implies that the relevant components of the coincidence in question (i.e., the prayers and the rain) are *modally fragile* and so is the fact that they come together (i.e., the relevant coincidence). The interesting question is this: is the fact that the components of a coincidence are modally fragile constitutive of or the determining factor that gives rise to the coincidence? Unfortunately for the modal account of luck, the answer is that it is entirely orthogonal to an event being a coincidence that its components are modally fragile or robust.

To bolster the point, consider a version of Jennifer Lackey's treasure case (Lackey, 2008). Hanna buries a treasure at location L. John independently places a rosebush in the ground of L. They neither know each other, nor they have any information about each other's actions or intentions. As John digs, he discovers Hanna's treasure. The discovery is lucky, of course, but also a *coincidence*. Now suppose that L is the only place where Hanna would have buried her treasure. Here is the explanation: she buried her childhood dog there, and she would have *never* chosen any other location to bury the treasure because, in this way, she can be completely sure that she'll never forget where the treasure is. As for John, he has a fixation: he never plants a rosebush in soil whose color is different from the roses. L is the *only* place with red soil in the area and John's roses are red. Given that, we can reasonably stipulate the following: Hanna and John independently had probability 1 (or close to 1) of, respectively, burying her treasure at L (Hanna) and of planting his rosebush at L (John). Accordingly, Hanna's and John's independent actions turn out to be modally robust: in most (if not all) nearby possible worlds, Hanna would still bury her treasure, and John would still plant his rosebush at L. Consequently, John

⁹ Why "if significant"? Because it might be a coincidence that two atoms collide at the other end of the universe, but that's not lucky for anyone. The fact that such a collision fails to be lucky is not because it is a coincidence, but because it is not significant to anyone.

would discover the treasure in most (if not all) nearby possible worlds. Yet, John's discovery is by luck, because it is a coincidence (a modally robust one), and so the modal view of luck is false.¹⁰¹¹

3 Two Competing Accounts of Individual Epistemic Luck

The modal account of individual epistemic luck is a straightforward application of the modal account of luck. Let "S" be any individual agent:

The modal account of individual epistemic luck: S's true belief that p is lucky if and only if S's belief that p is true in the actual world but in most nearby possible worlds in which S forms the belief that p in the same way as in the actual world, p is false.¹²

The modal account gets many clear cases of lucky true beliefs right. For example, you have not prepared for the multiple-choice exam at all, yet by wishful thinking, you come to believe that you have passed it. As luck would have it, you have gotten more than one-half of the answers right despite you answered the questions randomly—this is an unlikely event, of course. Thus, your belief turns out to be true but falls short of knowledge because of the luck in play. The modal view predicts this: you lack knowledge because you could easily have formed a false belief on the same basis (viz., by combining wishful thinking with answering the questions randomly). The point generalizes: epistemically flawed methods of belief formation (like wishful thinking, wild guessing or bad memory) tend to produce lucky beliefs if true. The reason, as per the modal view, is that they *all* could *easily* lead us to error.

However, although the modal view gets things right in many instances of epistemic luck, it also has a significant downside: it inherits the *coincidence problem* from its more general counterpart, the modal account of luck. This is especially problematic in the epistemic case because, as it turns out, central cases of epistemic luck (viz., most Gettier-style cases) do involve coincidences and, as we have seen, it is just a contingent matter of how one fills in the details that the relevant coincidences of the cases become modally robust or fragile. In other words, the problem is that coincidences always involve luck no matter whether they're modally robust or fragile, and the modal view can only account for coincidences whose components are modally fragile. And the epistemic case is no exception.

¹⁰ Notice how naturally the lack of control account would explain John's discovery: the fact that John discovers Hanna's treasure is beyond his control, period. This doesn't mean, of course, that the lack of control view is to be preferred, as it might still fall prey to counterexamples specifically tailored against it (see, e.g., Lackey 2008). However, as mentioned before, here is not the place to adjudicate between the different views about luck.

¹¹ It could be objected that modally fragile coincidences are *luckier* than modally robust coincidences. However, even if that's true, i.e., even if modal fragility might have a bearing on how lucky coincidences are, the modal fragility of the components of a coincidence is not the constitutive of or the determining factor that makes the coincidence lucky, so the modal account is still wrong.

¹² See Pritchard (2005) for this view.

By way of illustration, Smith has strong justification for believing two propositions: (1) that Jones is the man who will get the job and (2) that Jones has ten coins in his pocket. From (1) and (2), he infers the proposition that the man who will get the job has ten coins in his pocket. Orthodoxy says that Smith's belief is justified (Gettier, 1963). The belief is also true, but in an unusual way, because (a) it's Smith who finally gets the job and (b) he also happens to have ten coins in his pocket. (a) and (b) constitute a coincidence, as there is no connection (explanatory or nomological) between them. For this reason, Smith's true belief is lucky, which explains why it doesn't amount to knowledge despite being justified.

Question: would the relevant coincidence cease to be knowledge-undermining if the job was *never* meant to be for Jones, but for Smith only? Further question: what if Smith would *never* leave his home without ten coins in his pocket? If we turn the two components of the coincidence, (a) and (b), into modally robust events, the upshot is that in most (if not all) nearby possible worlds, Smith's belief that the man who will get the job has ten coins in his pocket continues to be true. Consequently, Smith's belief no longer counts as lucky according to the modal account of epistemic luck. But Smith's belief *is* lucky, precisely because it turns out to be true out of sheer coincidence, no matter how modally robust the components of this coincidence are.

Here is another much-discussed case of individual epistemic luck that the modal view gets wrong. Suppose you believe that there is a sheep in the field because you have looked at something that looks like a sheep. In reality, it's a sheep-looking dog. As luck would have it—or more precisely: as coincidence would have it—there is a sheep hidden behind a rock nearby, which makes your belief true (Chisholm, 1966). What if the sheep is chained to the rock or, more bizarrely, inside the rock? Here is the answer: your belief that there is a sheep in the field would continue to be true in most (if not all) nearby possible worlds. However, this doesn't make your belief non-lucky, because the relevant modally robust coincidence—namely, the fact that you've formed the belief that there is a sheep in the field on the basis of misleading evidence while there is a hidden sheep in the field—is lucky. So take your pick and modally robustify your preferred Gettier-style case: the modal view will never be able to explain why it is a case of knowledge-undermining luck.¹³

An alternative is to put forward a lack of control account of individual epistemic luck, such as this:

The lack of control account of individual epistemic luck: S's true belief that p is lucky if and only if S's coming to have the true belief that p is beyond S's control.

¹³ Except for the fake-barn case (Goldman 1976): cases with the structure of this much-discussed case *do not* involve coincidences. In fact, it suffices to make the fact that there is a barn in front of one modally robust for one to end up having knowledge (of the proposition that there is in fact a barn in front of one). This doesn't happen in standard Gettier-style cases. Unsurprisingly, epistemologists treat fake-barn-style cases differently. See Pritchard (Pritchard et al., 2010) for an explanation of the difference in terms of the distinction between intervening and environmental epistemic luck.

Many would deem this view a nonstarter, however. Doesn't a lack of control view of *epistemic* luck entail that most of our true beliefs are lucky? After all, if by "control" we mean *direct voluntary control*, it is fairly obvious that we cannot directly manipulate our doxastic states at will. The proponent of a lack of control view would thus need to offer a suitable account of what it takes to control belief formation. In particular, such a view would need to detach the notion of control from that of voluntariness and apply the resulting notion to individual belief formation. In principle, it is possible to give such a view, but here is not the place to do it, as it requires careful argumentation—recall: we are in the business of showing that knowledge-undermining collective luck exists.

In the absence of a plausible lack of control account of *individual* epistemic luck, perhaps it is better, all things considered, to opt for the modal view, even if—and this is a very big "if"—this view gets *core cases* of individual epistemic luck wrong: namely, any Gettier-style case involving a modally robust coincidence. The alternative, however, is to fall into the clutches of direct doxastic voluntarism. Of course, this means that we don't have a satisfactory account of individual epistemic luck yet. Be that as it may, the interesting question for our purposes is rather this: shall we encounter the same kind of impoverished theoretical situation when it comes to *collective epistemic luck*?

4 Two Competing Accounts of Collective Epistemic Luck

To see whether the dialectical situation replicates at the collective level, let's paraphrase the previous accounts of individual epistemic luck in collective terms and see what follows from the resulting views. Let "G" be any group.

The modal account of collective epistemic luck: G's true belief that p is lucky if and only if G's belief that p is true in the actual world but in most nearby possible worlds in which G forms the belief that p in the same way as in the actual world, p is false.

The lack of control account of collective epistemic luck: G's true belief that p is lucky if and only if G's coming to have the true belief that p is beyond G's control.

We shall focus on the modal account of collective epistemic luck first. Will this modal view inherit the coincidence problem, i.e., the problem that the modal account of individual epistemic luck inherited from the modal account of luck in the first place? Unfortunately, yes. Consider *Gettierized Hikers*. Poor hiker 1 makes two mistakes: mistaking path B for A and incorrectly measuring the time it takes to complete path B (which happens to coincide with the time it takes to walk path A). The problem for the modal view is that we can stipulate that both mistakes are modally robust, so that hiker 1's belief that A is the fastest path would continue to be true (out of modally robust coincidence) in all nearby possible worlds.

By way of illustration, let's make the first mistake modally robust in the following way: suppose there is *no* chance that hiker 1 realizes that the path she has walked is B instead of A. In *Gettierized Hikers*, all her evidence (and we might add: all

the evidence she could possibly acquire given her current epistemic position) points to the fact that she has indeed walked path A—this is false, of course, but reasonable enough for hiker 1 to keep believing it in all nearby possible worlds. Now, let's make hiker 1's second mistake modally robust too: instead of a squirrel accidentally dropping a nut on her stopwatch switching it off and on exactly when 1 h 30 m of walking is left, let's suppose that hiker 1's friend wants to make her believe that she's faster than she is, in such a way that she restarts her stopwatch exactly when 1 h 30 min of walking is left. If we further stipulate that there is no chance that her friend fails to do this action at *exactly* that time while also erroneously thinking that they're walking on path A, then the upshot is a *modally robust coincidence*: in most (if not all) nearby possible worlds, hiker 1 would still (falsely) believe that the path she's walking is A while also (falsely) believing that it has taken her 1 h 30 min to reach the top. On that faulty yet modally robust doxastic basis, hiker 1 infers the *true* belief that it takes 1 h 30 min to complete path A. Such a belief is not only true but also robustly so: in most (if not all) nearby possible worlds where hiker 1 believes that it takes 1 h 30 min to walk path A, that's how long it takes. Yet, it is clearly a matter of luck (actually, of coincidence) that such a belief is true.

Hiker 1's individual belief is also the belief upon which the *group* of hikers comes to form the collective belief that A is the fastest path (after trusting hiker 1's testimony). The group's true belief is, in this way, also lucky. Yet, contrary to what the modal account of collective epistemic luck would predict, in nearly all (if not all) nearby possible worlds in which the group forms the belief that A is the fastest path on the same basis, that proposition continues to be true. So, once again, the coincidence problem reappears in the context of collective epistemic luck for the modal view.

At this point, one might wonder why the modal view encounters the problem once and again in three different contexts (luck, individual epistemic luck, and collective epistemic luck). The reason has to do with the fact that modal views, in general, are solely concerned with the *modal profile* of the bearers of luck thus obviating their specific nature, i.e., it instinctively applies to events, individual opinions, and collective beliefs. Regardless of the nature of such bearers of luck, all that matters to a modal view is whether they could easily have failed to meet success conditions (e.g., of occurrence in the case of events and of correctness in the case of beliefs). When that's the case, the actual satisfaction of such success conditions (the actual occurrence of an event, the fact that a belief is correct) is by luck and *necessarily* so. However, as we have seen, coincidental events and beliefs can have *any* kind of modal profile and still be lucky, for a simple reason: they're coincidences. This means that counterexamples involving *modally robust coincidences* are to be expected in any context where the modal account finds application.

The lack of control account, by contrast, goes in the opposite direction: it focuses on the *nature* of the bearers of luck (events, individual beliefs, collective beliefs) regardless of its modal profile. For a lack of control view, it's not essential whether an event could easily have failed to occur or whether a true belief could easily have been false: if having that true belief or the event's occurrence is beyond one's control, then they are by luck. This also means that each lack of control view needs to understand the relevant notion of control *differently*, in a way that is tailored to the

specific nature of the bearer of luck such a view is about. For it is not the same to control, e.g., an event's occurrence than individual belief formation. For this reason, counterexamples or objections to one lack of control account (e.g., of the nature of luck) need not replicate as counterexamples to other lack of control views (e.g., of the nature of individual epistemic luck). This diversification constitutes a theoretical advantage over the modal account, precisely because it blocks the possibility of pervasive problems such as the coincidence problem.

Still, one might think that, in the epistemic case, the prospects for the lack of control account of *collective* epistemic luck are as dim as they are for the lack of control account of *individual* epistemic luck: just as the lack of control view of individual epistemic luck would strike many as a nonstarter because it leads to unacceptable direct doxastic voluntarism, the collective counterpart cannot get off the ground for exactly the same reason. It is this objection, however, that doesn't get off the ground. An important difference between individual and collective belief formation is that the latter does involve *voluntariness*. Let's elaborate on this.

In the collective epistemology literature, it's an open question whether collective beliefs are best understood as joint acceptances or else as beliefs proper.¹⁴ This, however, is irrelevant for the purposes of distinguishing a collective counterpart of individual epistemic luck. In very idealized terms, the formation of a group's view can be understood as a *three-stage process*: first, individual members bring along their beliefs and acceptances to the group as well as their private evidence; second, the group uses some kind of collective belief-forming procedure (e.g., deliberation, voting) that takes such beliefs or acceptances and evidence as inputs, so that, third, the procedure delivers a collective view as output, which can be conceived either as a collective belief or as a collective acceptance—for simplicity, we'll talk of "collective belief" to refer either to collective acceptances or collective beliefs proper.

Why is collective belief formation *voluntary*? Because paradigmatic methods of collective belief formation involve voluntariness on behalf of the individuals that participate in them.¹⁵ Take group deliberation: group members who engage in group discussion do it voluntarily; not only that, the very notion of reaching a collective agreement (e.g., consensus) following deliberation implies voluntariness. Even in the case that individual members omit to share their beliefs and evidence with their fellow members for, say, strategic reasons, that's a voluntary decision on their behalf, which might of course translate into a suboptimal group decision or belief, but one that involves voluntariness nonetheless. Now consider voting. Groups often vote to settle on certain views, and the typical way to do this is that group members cast a vote (or abstain) voluntarily (see below for a limit case of voting in which voluntariness on behalf of group members is absent).

¹⁴ For a defense of the former view, see Wray 2001 and Hakli 2006; for a defense of the latter, see Gilbert 1987, 2002; Tuomela 1992, 2000; Tollefsen 2002, 2003; and Bird 2019.

¹⁵ This is of course compatible with the fact that there might be a mismatch between what a group believes and the beliefs individually held by its members. See Gilbert (1987) for some examples and relevant discussion.

That both deliberation and voting involve voluntary actions by group members doesn't of course mean that the individual members' voluntary contributions as well as the norms governing them need be the same. Instead, they depend on the specifics of the relevant method for collective belief formation. For example, during deliberation, group members typically share their beliefs and evidence voluntarily. By contrast, during non-deliberative anonymous voting, they don't; instead, group members cast anonymous votes, voluntarily, absent communication among them. There is voluntariness involved in both cases, but the relevant voluntary actions differ.

To put another example, members of a group might realize they're incompetent on an issue they need to settle on urgently, so that they voluntarily decide to defer to an expert for forming a group view. More specifically, in this kind of case, group members voluntarily let an expert's view stand as the group's view. Forming a collective view in this way still involves voluntariness. Even in the limit case in which an external agent coerces group members to put certain inputs in the relevant method of collective belief formation (e.g., share certain evidence with the group during discussion; cast their votes in a particular direction), so that the group ends upholding the belief of this agent, the group members' wills are *replaced with* that agent's will, which means that group belief formation is still voluntary, albeit in a devious way. The bottom line of all this is that the main obstacle to theorizing individual epistemic luck in terms of lack of control—namely, direct doxastic voluntarism—does not apply to the collective case.

This is an important finding because it opens the door to giving a lack of control view of collective epistemic luck, thus allowing us to drop the incorrect modal view (incorrect both in the case of individual and collective epistemic luck). However, note that a lack of control account of collective epistemic luck cannot be spelled out in general terms. The different kinds of procedures groups use to form their collective beliefs voluntarily and thus achieve group knowledge—procedures such as deliberation, voting, deference to an expert, or, even, external coercion—involve very different kinds of actions that are relevant, precisely, to such collective views qualifying as group knowledge (e.g., when competently done). For this reason, in the remainder of the paper, we won't be concerned with distinguishing different forms of voluntary control specific to the different kinds of methods of collective belief formation. Each method requires careful examination.¹⁶ Instead, we'll offer a general argument for the following claim: groups that form their beliefs via internal deliberation (i.e., by exchanging views and evidence among group members) are more vulnerable to knowledge-undermining collective luck than groups that form their beliefs via non-deliberative methods, such as non-deliberative voting. If that's correct, the greater exposure to knowledge-undermining collective luck that deliberation gives rise to provides *a* reason (not a conclusive one, though) for thinking that non-deliberative methods of group belief formation—in particular, voting—have greater epistemic value.

¹⁶ For a detailed account of group knowledge in deliberative settings and how luck may undermine it, see Barba and Broncano-Berrocá (manuscript).

5 Why Deliberation Is More Vulnerable to Knowledge-Undermining Collective Luck than Voting

To put it in general terms, when a group is in the business of settling whether p , groups can deliberate (deliberation cases) or take a vote absent communication among group members (voting cases).¹⁷ More specifically, deliberation cases of group belief formation are cases in which the operative members¹⁸ of a group G hold individual attitudes towards p at t_1 ; at t_2 , G 's operative members deliberate among themselves (i.e., they exchange reasons, evidence, arguments, and so on) with an eye towards settling on whether p , i.e., on whether one answer or another to such a question should stand as G 's view; at t_3 , as a result of this process, they collectively settle on a particular answer which becomes G 's view on whether p . *Hikers* is an example of a deliberation case: to settle on what path is the quickest route to the top; the hikers share their views and evidence and engage in group discussion to collectively conclude, correctly, that A is the fastest path. Such a method, as the case shows, can be conducive to group knowledge.

By contrast, the kinds of voting cases of group belief formation we have in mind are cases in which the operative members of a group G hold individual attitudes towards p at t_1 ; at t_2 , G 's operative members aggregate their views by taking a vote given some voting rule (e.g., majority rule), absent any communication among each other, with an eye towards settling on whether p , i.e., on whether one answer or another to such a question should stand as G 's view; at t_3 , as a result of this process, they collectively settle on a particular answer, which becomes G 's view on whether p . *Hiring Committee* is an example of a voting case: to settle on who is the most qualified candidate for the job, committee members cast individual votes drawing on their private evidence and individual competences, and, as the case shows, voting can be conducive to group knowledge as well.

Therefore, both deliberation and voting (absent communication) are ways for groups to form collective beliefs; not only that, they're both ways of knowing collectively.¹⁹ Which method is more vulnerable to knowledge-undermining luck? To answer this question, it's useful to lay down typical actions that lead to the formation of collective beliefs, both in the case of deliberation and voting.

When a group deliberates or takes a vote to settle a question and thus to let some view stand as the group's view, group members typically start by *collecting evidence* bearing on the different answers to that question. Such evidence becomes the group members' *private evidence*. For example, in *Hikers*, the question is what path is the quickest route to the top, and each hiker has measured how long it takes to walk one (and only one) of the four paths under consideration. That's the individual

¹⁷ Groups can of course deliberate *and* then take a vote, but to evaluate which method of collective belief formation is more vulnerable to knowledge-undermining luck, we will only consider cases of groups that either vote or deliberate but not both.

¹⁸ According to Lackey (2016: 350), operative members are those who "have authority or power to determine certain outcomes for the group as a whole."

¹⁹ For a general comparison of the reliability of group deliberation with that of voting, see Broncano-Berrocal & Carter (2020).

private evidence they bring to group discussion. In *Hiring Committee*, the question is what candidate is the most qualified for the job, and, as the case goes, each committee member reads the CVs of the different candidates and interviews them one-to-one. That's the private evidence they draw on to cast their anonymous votes.

The next action group members typically do both in deliberation and voting cases is to assess the quality and confirmational import of their private evidence. In particular, they first need to assess whether the evidence is good or bad, i.e., whether it comes from trustworthy sources (e.g., trustworthy informants) or reliable methods (e.g., reliable cognitive faculties) or else from untrustworthy sources (e.g., hearsay) or unreliable methods (e.g., wishful thinking, motivated reasoning); second, group members must assess to what extent such evidence confirms or disconfirms the different answers to the question of whether *p*. In *Hikers*, for example, the hikers use their reliable memory to recall the time (*t*) it took them to walk each path. On that basis, they also competently reason whether such evidence supports conditionals of the form "If it takes more than *t* to walk the other paths, then the path I've walked is the fastest route to the top." In *Hiring Committee*, committee members proceed likewise: once they've collected relevant evidence on the candidates by competently and impartially reading their CVs and by having one-to-one interviews with them, they assess to what extent such evidence supports claims of the form "Candidate *n* is the most qualified for the job."

Now, the *key difference* between deliberation and voting cases is that, in voting cases (at least the kinds of cases we're considering), group members do not perform any further actions except for casting a vote. In this way, the epistemic appropriateness of their votes comes down to how competently they have collected the evidence the votes are based on, as well as to how competently they have assessed the quality and confirmational import of such evidence. In deliberation cases, by contrast, group members take further actions. Our contention is that it is mainly for this reason that deliberation is more vulnerable than voting to knowledge-undermining collective luck. But before reaching this conclusion, let's ask: what further actions do members of deliberative groups perform?

In deliberation cases, once group members have individually collected evidence on whether *p* and assessed its quality and confirmational import, they need to decide whether to *share their private evidence* with the rest (along with their views on the question under discussion). This third action is typically followed by another one: once the private evidence is put on the table, group members need to *judge the shared evidence's confirmational import*. After that, group discussion produces collective agreements or disagreements about whether a particular answer to whether *p* should stand as the group's view. The two kinds of actions just distinguished—private evidence filtering and shared evidence assessment—take place during group discussion and, for this reason, don't occur in the kinds of voting cases we're considering (i.e., voting cases in which communication among group members is absent).

As it happens with any kind of action, the four actions distinguished (i.e., evidence acquisition, private evidence assessment, private evidence filtering, and shared evidence assessment) can be competently or incompetently done. Given that, the key idea to keep in mind, in what follows, is this: when an action is incompetently carried out but successful nevertheless, it is typically successful due to luck.

For example, if a student answers the questions of a multiple-choice exam randomly and yet gets the highest grade, it is by luck (not by competence) that the student has gotten that grade.

With that idea in place, the *argument* for thinking that deliberation is more vulnerable to knowledge-undermining luck than voting starts with the premise that (1) the degree of exposure to knowledge-undermining luck of a method of collective belief formation is to a large extent premised on the degree of exposure to luck of the relevant actions that are typically performed when implementing that method and which are conducive or contribute to collective belief formation. It continues with the idea that (2) all the relevant actions that are conducive or contribute to group belief formation in voting cases are also present in deliberation cases, namely evidence acquisition and private evidence assessment.²⁰ However, (3) deliberation features two further types of action that are not present in the kind of voting cases we're considering, namely filtering of the private evidence and assessment of the shared evidence.

Given (1)–(3), the argument for the greater vulnerability to knowledge-undermining luck of deliberation over voting is straightforward. It only needs one further premise: (4) the more actions you do (in number and kind), the higher the chance of performing some of them incompetently, and, if successful, the greater exposure to luck. Deliberation does involve more belief-forming actions (in number and kind) than voting, so, given (1)–(4), the risk that luck will prevent true collective beliefs from qualifying as knowledge is greater in deliberation than in voting cases.

This argument is compatible with the possibility that all members of a deliberative group are competent at doing the relevant belief-forming actions (e.g., competent at filtering out their bad private evidence from group discussion or at assessing the confirmational import of the evidence shared by others correctly). When that happens, chances are that the group won't fall prey to knowledge-undermining luck (quite the contrary). In the same vein, the argument is compatible with the possibility that if all members of a voting group are incompetent, the resulting group view will, if correct, fall short of collective knowledge due to collective epistemic luck. Instead, what the argument aims to show is this: *all other things being equal*—namely when the relevant belief-forming actions that deliberation and voting have in common are performed *as* competently or incompetently—the risk that luck will prevent true collective beliefs from qualifying as knowledge is greater in deliberation than in voting cases, precisely because the former involves further actions (both in number and kind) that the latter doesn't.

Why is it that deliberation increases the liability to knowledge-undermining collective luck, more specifically? The main reason is that the pernicious epistemic effects of individual incompetence can more easily be passed on to the rest of the group when communication is present than when it is absent. To see this more

²⁰ In Barba and Broncano-Berrocal (n.d.), we argue that (i) lucky evidence acquisition is compatible with group knowledge, whereas (ii) lucky assessment of the private evidence can potentially undermine it.

clearly, let's compare non-deliberative anonymous majority voting with deliberation using an example.

Deliberation, as we have seen, typically features two further kinds of actions that voting doesn't: (i) filtering of the private evidence and (ii) assessment of the shared evidence. Let's see what happens when such actions are done in an incompetent but successful manner (hence luckily). Consider the following version of *Hikers*. Suppose that hikers 1, 2, 3, and 4 have, respectively, walked paths A, B, C, and D *twice*. The first time they walked their corresponding path, they used a reliable stopwatch to measure the walking time; the second time, they used an unreliable stopwatch. Yet, in neither case, they know whether the devices they've used are reliable or unreliable. As a consequence, each hiker has two conflicting bodies of evidence (i.e., two conflicting measurements for each path) and no way to decide which body of evidence is the good one (i.e., which measurement comes from the reliable stopwatch). Suppose that the only way for the group to *collectively* settle on what path is the quickest route to the top is non-deliberative anonymous majority voting. Given how indecisive their individual epistemic situation is, all hikers would reasonably cast a blank ballot, the upshot of which would be that the group does not settle on any view—this would be equivalent to *collective suspension of judgment*. Interestingly, this would still be the case if one of the hikers (e.g., hiker 1) mistakenly believed that the unreliable stopwatch she has used the first time is reliable and that the reliable stopwatch she has used the second time is unreliable. As a consequence of that double mistake, hiker 1 would think that the former stopwatch's measurement of path A's walking time is correct (and that the latter stopwatch's measurement is incorrect). Let's suppose that, on that epistemically inappropriate basis, hiker 1 votes for the view that the fastest path to the top is A, while the others still vote for "no view." In such a case, three votes for "no view" outweigh one vote for path A, so that, *even if it turns out that A is the fastest path to the top*, the group would still suspend judgment on the issue, which is the right kind of doxastic attitude the group should adopt given the hikers' impoverished epistemic situation.

Let's consider now what would have happened in the latter case if the relevant method of collective belief formation had been group deliberation instead of voting. Hiker 1 clearly overestimates the confirmational import of her evidence, in that, from the false belief that she has used a reliable stopwatch to measure path A's walking time (the first time) and from the belief in the resulting measurement, she infers the belief that path A is the quickest way to the top. Even if, *by coincidence*, it turns out that path A is the quickest—i.e., even if that measurement is, due to luck, correct (despite the unreliability of the stopwatch)—such an inference is epistemically illegitimate. The reason is that hiker 1 has no conclusive evidence for how long it takes to walk the other paths. Yet, during group discussion, we may suppose she *presents* her evidence for the proposition that A is the quickest path as conclusive to the others. In this kind of deliberative group setting, i.e., a case in which three group members have no conclusive evidence for that proposition and one group member shares non-conclusive evidence as being conclusive, chances are that the former will trust the word of the latter so that all will end up letting the latter's view stand as the group's view. The upshot is that the group of hikers would agree on the collective view that A is the quickest path. Now, that collective view is clearly unjustified and

does not amount to group knowledge for the previous reasons. Interestingly, even when such a view turns out to be correct—again, suppose that A is indeed the fastest path to the top and that by sheer luck hiker 1's unreliable stopwatch has produced a random measurement that coincides with the real time it takes to complete A—the group clearly lacks collective knowledge also due to the luck involved (i.e., on top of the fact that the group's view is unjustified).

At this point, the proponent of deliberation might press back and argue in reply that deliberation cannot be more vulnerable to knowledge-undermining luck than voting because the kinds of belief-forming actions that deliberation involves (and that voting doesn't) give rise to epistemic gains that are not available in voting cases. In particular, the sharing of the deliberators' views and the private evidence upon which such views are based (as well as the mutual scrutiny of everyone's views and evidence) provide groups with a powerful tool to implement checks and balances that can offset the possible pernicious epistemic effects of individual incompetence. Along these lines, Sunstein (2006) has suggested that one epistemic benefit of deliberation is that it makes it possible that *groups equal their best members*, something that cannot be achieved in groups that form their collective views by non-deliberative majority voting. Indeed, in some deliberation cases, a group may count with an expert who, thanks to the full disclosure of the evidence and the public exchange of reasons, can be identified and trusted as such when it comes to settling on a collective view. In voting cases, by contrast, such a possibility is not available, at least if the relevant voting rule is standard (i.e., not weighted) majority voting. Besides, in some deliberation cases where group members are not especially competent, group deliberation might *aggregate* the average individual contributions thus enhancing group competence. Finally, deliberation also has the potential, through the exchange of reasons, that a group discovers *new* good reasons that neither of the members had before discussion.

As a rejoinder, from the existence of such potential epistemic benefits, it doesn't necessarily follow that deliberation is more reliable than voting *qua* method of collective belief formation, as it might have other drawbacks that bear negatively on its reliability.²¹ At least from a formal (i.e., to some extent idealized) point of view, deliberation is truth-conducive in a similar way as majority voting. For example, Hartmann and Rafiee Rad (2018) offer a Bayesian model of deliberation that establishes precisely this. However, in less idealized group settings, real-life deliberation falls prey to several epistemically inappropriate group dynamics that can lower its reliability *qua* method of collective belief formation and that are not present in voting cases: dynamics such as *shared information bias*, the tendency to discuss shared evidence in detriment of discussing potentially relevant private evidence or evidence only shared by a few group members (cf., Stasser & Titus, 1985), or *group polarization*, the tendency of like-minded groups to adopt more extreme views after

²¹ See Broncano-Berrocal & Carter (2020) for further discussion.

group discussion than those held before deliberation (e.g., Stoner, 1961; Burnstein & Vinokur, 1977).²²

The previous rejoinder makes the prospects of a pro-deliberation argument of the following form dim (the kind of argument that seems to underline the reply of the defender of deliberation): (i) the more reliable a method for doing a certain kind of action is, the less vulnerable to luck such a method is when one uses that method and succeeds in performing such kind of action; (ii) deliberation is more reliable at forming true collective views than voting; therefore, (iii) deliberation is less vulnerable to luck than voting when groups use those methods and succeed in forming true collective beliefs.

Our argument still stands: all other things being equal, deliberation is more vulnerable to knowledge-undermining luck than voting because, for deliberation to produce collective views, more belief-forming actions (both in number and kind) are required than in the case of voting. Unless one shows that the further belief-forming actions that deliberation involves turn deliberation into a more reliable method for forming true collective views than voting, the argument is sound. If that's the case, the greater exposure to knowledge-undermining luck that deliberation gives rise to provides *a* reason (not a conclusive one, though) for thinking that non-deliberative methods of group belief formation—in particular, voting—have greater epistemic value.

6 Concluding Remarks

Let's take stock. We've started by establishing the existence of collective epistemic luck by parity of reasoning with the individual case, and we've illustrated the phenomenon with deliberative and non-deliberative voting cases in which luck undermines collective knowledge. Not only that, we've also demonstrated the existence of collective Gettier cases. One of the main points in our argumentation has been that salient differences between individual and collective belief formation open the door for a lack of control approach to collective epistemic luck, instead of modeling it in modal terms, as per usual in the case of individual epistemic luck.

We've argued that the modal view faces a critical problem: that of explaining modally robust coincidences (the coincidence problem). Importantly, we've shown that this view encounters such a problem in three different contexts: luck, individual epistemic luck, and collective epistemic luck. The reason, as we have argued, is that modal accounts of luck (of any kind) are solely concerned with the *modal profile* of the bearers of luck and ascribe luck whenever the occurrence in the case of events or the correctness in the case of beliefs are modally fragile. However, as we've shown, coincidental events and beliefs can have *any* kind of modal profile and still be lucky. This explains why counterexamples involving *modally robust coincidences* are to be expected in any context where the modal view finds application.

²² For a book-length treatment of the epistemological dimensions of group polarization, see Broncano-Berrocal and Carter (2021).

We have then argued that the lack of control view does not have this problem, because it obviates the modal profile of events and beliefs when explaining why they occur or are true by luck. The implication of this is that the lack of control view must be spelled out in different, specific ways depending on the phenomenon it aims to account for. For it is not the same to explain in what sense an event's occurrence is beyond one's control than the correctness of an individual or a collective belief. For this reason, counterexamples or objections to one lack of control account need not replicate as counterexamples to other lack of control views. This diversification constitutes a theoretical advantage over the modal account, especially because it makes it possible to give a comprehensive lack of control account of collective epistemic luck, thus avoiding the reasonable objection that one cannot have direct voluntary control of individual belief formation. In the collective case, we've argued, belief formation *is* voluntary.

Finally, we've offered a general argument for the following claim: groups that form their beliefs via internal deliberation (i.e., by exchanging views and evidence among group members) are more vulnerable to knowledge-undermining collective luck than groups that form their beliefs via non-deliberative methods (viz., non-deliberative voting). The bottom line is that the greater exposure to knowledge-undermining collective luck that deliberation gives rise to provides *a* reason (not a conclusive one, though) for thinking that non-deliberative methods of group belief formation—in particular, voting—have greater epistemic value.²³

Funding Open Access funding provided thanks to the CRUE-CSIC agreement with Springer Nature.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Ballantyne, N. (2012). Luck and interests. *Synthese*, 185, 319–334.
- Barba, M. & Broncano-Berrocal, F. (*manuscript*). Group deliberation and epistemic luck.
- Bird, A. (2010). Social knowing: The social sense of 'scientific knowledge.' *Philosophical Perspectives*, 24, 23–56.
- Bird, A. (2019). Group belief. In M. Fricker, P. J. Graham, D. Henderson, & N. Pedersen (eds.) *The Routledge handbook of social Epistemology*. Routledge

²³ Work on this paper has been supported by (i) the research project *New Perspectives on Epistemic Risk* (PGC2018-098,805-B-100) funded by the Spanish Ministry of Science and Innovation and (ii) a 2019 Leonardo Grant for Researchers and Cultural Creators, BBVA Foundation. The BBVA Foundation accepts no responsibility for the opinions, statements, and contents included in this paper, which are entirely the authors' responsibility.

- Broncano-Berrocal, F. (2015). Luck as risk and the lack of control account of luck. *Metaphilosophy*, 46, 1–25.
- Broncano-Berrocal, F. (2016). Luck. *The Internet Encyclopedia of Philosophy*. <URL: <https://iep.utm.edu/luck/>> Accessed 13 October 2021.
- Broncano-Berrocal, F., & Carter, J. A. (2020). Deliberation and group disagreement. In F. Broncano-Berrocal & J. A. Carter (Eds.), *The epistemology of group disagreement* (pp. 9–45). Routledge.
- Broncano-Berrocal, F. & Carter, J. A. (2021) *The philosophy of group polarization. Epistemology, metaphysics, psychology*. Routledge.
- Burnstein, E. & Vinokur, A. (1977). Persuasive argumentation and social comparison as determinants of attitude polarization. *Journal of Experimental Social Psychology*, 13(4), 315–332.
- Carter, J. A. (2015). Group knowledge and epistemic defeat. *Ergo: An Open Access Journal of Philosophy*, 2, 28.
- Chisholm, R. (1966). *Theory of knowledge*. Prentice-Hall International Inc.
- Gettier, E. (1963). Is justified true belief knowledge? *Analysis*, 23, 121–123.
- Gilbert, M. (1987). Modelling collective belief. *Synthese*, 73, 185–204.
- Gilbert, M. (2002). Belief and acceptance as features of groups. *ProtoSociology*, 16, 35–69.
- Gilbert, M. (2013). *Joint commitment*. Oxford University Press.
- Goldman, A. (1976). Discrimination and perceptual knowledge. *Journal of Philosophy*, 73, 771–791.
- Goldman, A. (2014). Social process reliabilism: Solving justification problems in collective epistemology. In J. Lackey (Ed.), *Essays in collective epistemology* (pp. 11–42). Oxford University Press.
- Hakli, R. (2006). Group beliefs and the distinction between belief and acceptances. *Cognitive Systems Research*, 7, 286–297.
- Hartmann, S., & Rafiee Rad, S. (2018). Voting, deliberation and truth. *Synthese*, 195(3), 1273–1293.
- Lackey, J. (2008). What luck is not. *Australasian Journal of Philosophy*, 86, 255–267.
- Lackey, J. (2014). Socially Extended Knowledge. *Philosophical Issues*, 24, 282–298.
- Lackey, J. (2016). What is justified group belief? *Philosophical Review*, 125, 341–396.
- Mathiesen, K. (2006). The epistemic features of group belief. *Episteme*, 2, 161–175.
- Owens, D. (1992). *Causes and coincidences*. Cambridge University Press.
- Pritchard, D. (2005). *Epistemic luck*. Oxford University Press.
- Pritchard, D. (2014). The modal account of luck. *Metaphilosophy*, 45, 594–619.
- Pritchard, D., Millar, A., & Haddock, A. (2010). *The nature and value of knowledge*. Oxford University Press.
- Russell, B. (1948/2009). *Human knowledge: Its scope and limits*. Routledge.
- Stasser, G. & Titus, W. (1985). Pooling of unshared information in group decision making: Biased information sampling during discussion. *Journal of Personality and Social Psychology*, 48, 1467–1478.
- Stoner, J. A. F. (1961). A Comparison of Individual and Group Decision Involving Risk. Unpublished master's thesis, Massachusetts Institute of Technology
- Sunstein, C. R. (2006). *Infotopia*. Oxford University Press.
- Tollefsen, D. (2002). Organizations as true believers. *Journal of Social Philosophy*, 33, 395–410.
- Tollefsen, D. (2003). Rejecting rejectionism. *ProtoSociology*, 18, 389–405.
- Tollefsen, D. (2006). From extended mind to collective mind. *Cognitive Systems Research*, 7, 140–150.
- Tuomela, R. (1992). Group beliefs. *Synthese*, 91, 285–318.
- Tuomela, R. (1995). *The importance of us: A philosophical study of basic social notions*. Stanford University Press.
- Tuomela, R. (2000). Belief versus acceptance. *Philosophical Explorations*, 3, 122–137.
- Wray, K. B. (2001). Collective Belief and Acceptance. *Synthese*, 129, 319–333.
- Wray, K. B. (2007). Who has scientific knowledge? *Social Epistemology*, 2, 337–347.