King's Research Portal

DOI:

[10.1111/phis.12085](https://doi.org/10.1111/phis.12085)

Document Version

Peer reviewed version

[Link to publication record in King's Research Portal](#)

Citation for published version (APA):

Dutant, J. (2016). How to be an Infallibilist. *Philosophical Issues*, 26(1), 148-171.
<https://doi.org/10.1111/phis.12085>

Citing this paper

Please note that where the full-text provided on King's Research Portal is the Author Accepted Manuscript or Post-Print version this may differ from the final Published version. If citing, it is advised that you check and use the publisher's definitive version for pagination, volume/issue, and date of publication details. And where the final published version is provided on the Research Portal, if citing you are again advised to check the publisher's website for any subsequent corrections.

General rights

Copyright and moral rights for the publications made accessible in the Research Portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognize and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the Research Portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the Research Portal

Take down policy

If you believe that this document breaches copyright please contact librarypure@kcl.ac.uk providing details, and we will remove access to the work immediately and investigate your claim.

How to be an Infallibilist*

Julien Dutant

August 2016

Two central topics in epistemology are the *nature of knowledge* and the *norm of belief*. The first concerns what knowledge is, the second what we should believe. Infallibilism is a distinctive answer to both. On the first, it claims that you know p just if you infallibly believe p . On the second, it claims that you should believe p with certainty just if you know p . Below we will have much more to say about what “infallibly believe” means. But hopefully it is already clear that the two answers naturally fit together. Believing infallibly is, roughly, having a belief that could not be mistaken. Believing with certainty is, roughly, taking something for granted, without doubt or reservation, and disregarding alternative possibilities. It is natural to think that having a belief in p that could not be mistaken—and nothing short of that—warrants taking p for granted and disregarding the possibility that p is not so. So it is natural to think that if knowledge is infallible belief, then knowledge, and knowledge only, warrants certainty.

Infallibilism has a bad press these days. It is said to be *factually* wrong because there are few things that we infallibly believe and few things that we should be certain of. It is said to be *conceptually* wrong because knowledge does not require infallible belief and is not enough for justified certainty. Budding epistemologists are warned: if you commit both errors, you will be excessively dogmatic like Descartes; if you commit the conceptual error alone, you will be excessively sceptic like Hume. Thankfully, the story goes, Reid, Peirce, Popper and others have led the way out of this madness and we are all past it now. Indeed, most epistemologists accept an outlook on which knowledge is in some important sense compatible with the possibility of error and on which there are few things we should be absolutely certain of. As Stewart Cohen (1988) put it, “the acceptance of fallibilism in epistemology is virtually universal.”

The bad press is undeserved. There *is* a mistaken infallibilist view that we can ascribe to several past philosophers that leads to a false dilemma between excessive dogmatism and excessive

*Thanks to Bob Beddor, Davide Fassio, Sandy Goldberg, Mikkel Gerken, Chris Kelp, Arturs Logins, Mona Simion, Levi Spectre, Jacques Vollet and the audience at the 2016 European Epistemology Network in Paris for useful feedback on the ideas of this paper. Special thanks to Charity Anderson for many fruitful discussions of this topic. This paper owes much to the work of Timothy Williamson and John Hawthorne though that is not to say that they endorse the views defended here.

scepticism. But what is mistaken about the view is *internalism*, not infallibilism. When spelled out properly infallibilism is a viable and even attractive view. Because it has long been summarily dismissed, however, we need a guide on how to properly spell it out. The guide has to fulfil four tasks. The first two concern the nature of knowledge: to argue that infallible belief is necessary, and that it is sufficient, for knowledge. The other two concern the norm of belief: to argue that knowledge is necessary, and that it is sufficient, for justified certainty. With such a guide in hand infallibilism can be evaluated on its own merits. The most controversial parts are the first and fourth. The idea that knowledge requires infallible belief is thought to be excessively sceptical. The idea that knowledge warrants certainty is thought to be excessively dogmatic.

The four tasks are more than a single paper can chew. The present one focuses on the first: arguing that knowledge requires infallible belief. Section 1 motivates and clarifies the claim. The motivation comes from what we may broadly call the *Gettier problem*—a range of cases and considerations that suggest that knowing requires being in some sense immune from error. Section 2 addresses three main objections. They arise from, respectively, the *Old Evil Demon problem*, *Widespread Chance* and *Specific Counterexamples*. Each provides putative cases of knowledge that is not in the relevant sense immune from error. In addressing the first we also set aside the mistaken, internalist kind of infallibilism that is responsible for the view's bad press. Section 3 asks whether the view reached here can properly be called infallibilist. In particular, it addresses *Ecumenist* views who accept the infallibilist upshot of the Gettier problem but nevertheless claim to be "fallibilist" in some other sense.

In the remainder, unless otherwise specified, "Infallibilism" stands for the claim that knowledge requires infallible belief.

1 Motivation

1.1 An argument for safety

Consider the following cases:

Sight. I look at a pen in front of me. I believe, correctly, that it is less than 2 inches long. However, if it had been a tenth of an inch longer I would believe the same and be mistaken. (based on Williamson, 2000, chap. 5)

Fakes. I look at a display of 50 identical-looking pens. They all have blue caps, but only one of them is a blue-ink pen. I happen to look at the blue-inked pen and believe that it is a blue-ink pen. If I had looked at any other pen I would have believed the same and be mistaken. (based on Goldman, 1976, due to Ginet)

Prime. I have a foolish method to determine whether a number is prime. For those below 20 I match them with a correct list of primes. For those above, I sum the digits and if the result is below 20 and prime, I deem the original number prime. If the result is above 20 I sum the digits again, and so on. Using that method I believe—correctly—that 47 is prime. I likewise believe, mistakenly, that 49 is prime. (Dutant, 2010, chap. 2)

In each of these cases we have a true belief and an actual or counterfactual false belief. The two beliefs are alike in respects that matter for the attribution of knowledge: broadly put, they are formed in similar circumstances, by similarly constituted agents, on similar bases, in similar ways, about similar subject-matters, etc. We may say that these beliefs are *alike in epistemically relevant respects*, or *epistemically alike*. I call them *peers*.¹

Cases like the above motivate the idea that one knows only if one's belief is sufficiently unlike false beliefs. Many would judge that I fail to know, and would judge so in the light of the relevantly similar mistaken belief. It is not crucial that we agree on these particular cases, however. They can be amended so that the mistaken peers are even more similar to my true beliefs. Somewhere along the line we are bound to reach a point where it is clear enough to all that I do not know. Moreover, further putative cases of true belief that is not knowledge can be gathered from the so-called Gettier literature and can be used to make the same point. Together they suggest the following, where S stands for some agent, p some proposition, t a time and w a world:

Safety. S knows that p at t in w only if S 's belief that p at t , w is *sufficiently epistemically unlike* any false belief.

Call this the *generic safety* condition.²

Three things are worth stressing. First, generic safety entails truth. Since every belief is maximally like itself in all respects, every belief is sufficiently like itself in epistemically relevant respects, so no false belief is safe. Second, beliefs in logical and necessary truths are not trivially safe. As the prime number case illustrates, a belief in a necessary truth may be sufficiently like a belief in a

¹Manley (2007, 406) uses “counterpart thought” and Williamson (2009a) uses “epistemic counterparts” for the same idea.

²Compare Williamson (2000, 100): “One’s belief that p in a case α is safe iff one avoids false belief cases sufficiently similar to α ”. Note that, like the condition above and unlike other “safety” conditions, Williamson’s statement is not framed in terms of “close worlds” or “what could easily have been the case” but of similarity. Note also that it does not require mistaken peers to involve the same proposition, nor even the same subject—since “one” is used to stand for the agent as a centre of a case, whoever that is (Williamson, 2000, 94). There is a glitch in Williamson’s statement, however. Williamson’s “cases” are centered worlds: worlds with a subject and time singled out. So if I have currently some false belief, *all* my other beliefs are such that one has a false belief in sufficiently similar case, namely my present case itself. It is obvious that Williamson intends the clause to be restricted in some way: either to a narrower reading of “cases”—a centered world with a belief singled out—or to *relevantly similar* beliefs. The condition above is the simplest way to capture the idea while avoiding the glitch. See Szabó Gendler and Hawthorne (2005, 333) and Manley (2007, 406) for similar amendments.

different, false proposition. Third, generic safety, unlike some other “safety” conditions on knowledge, is not framed in terms of counterfactuals, “close worlds”, or “what could easily have been the case”.³

The underlying metaphysics can be spelled out in various ways. One is to allow *time- and world-bound token beliefs*: my present belief that meerkats are mammals is a distinct belief from the belief that meerkats are mammals I had five seconds ago, and distinct from the belief that meerkats are mammals I would have if my nose was a tiny bit longer, and so on. Peerhood is a relation between these belief tokens and safety says that one token constitutes knowledge only if its peers are all true. Alternatively, we may dispense with belief tokens and replace them by specifications of the belief relation: there is a way in which one believes p at t, w and a way in which one believes p' at t', w' . Peerhood is a relation among these relations: the way in which one believes p at t, w is like the way in which one believes p' at t', w' . Alternatively, we may even dispense with the specific belief-relations and simply have a eight-place peerhood relation of S believing p at t, w like S believes p' at t', w' . Since the latter relation can be defined on all options the third is the most metaphysically neutral. The first is the least verbally cumbersome, however, as English has much smoother ways of talking about individuals than about relations. So we stick with the first and talk in terms of belief tokens.⁴

To be fully honest I do not think that peerhood should be restricted to beliefs. For any given subject, time, world and proposition, the agent is related to that proposition in various ways. These ways need not involve attitudes on the subject’s part: if I have never heard of Bob then my relations to the proposition that Bob is creepy are that I am not aware of it, I do not believe it nor disbelieve it, and so on and so forth. The totality of one’s relations to a proposition at a given time and world may be called *one’s standing towards that proposition* at that time and world. Peerhood can be generalised to (token) standings towards propositions. Safety then says that one knows p only if one’s standing towards p is epistemically unlike anyone’s standing to any falsehood. We may still recover the claim that knowledge entails belief if we say that all non-belief-involving standings are epistemically like each other. Since some falsehoods are not believed, every non-belief-involving standing would come out as unsafe. However, a major advantage of generalised safety is to allow

³For counterfactuals see Sosa (1996, 274–276) (“one tracks the truth, outright, in believing that p IFF one would believe that p iff it were so that p ”), Sosa (1999, 142) (“a belief by S that p is “safe” iff: S would not believe that p without it being the case that p ”). For “could easily have been the case” see Sainsbury (1997) (“If you know, you couldn’t easily have been wrong”), Sosa (1999, 142) (“call a belief by S that p “safe” iff: [...] not easily would S believe that p without it being the case that p ”), Williamson (2000, 147) (“If one knows, one could not easily have been wrong in a similar case”) and Pritchard (2012, 253) (“If S knows that p then S ’s true belief that p could not have easily been false”). For close-worlds formulations see notably Pritchard (2005, 71, 156, 162) (“if an agent knows a contingent proposition ϕ , then, in most nearby possible worlds, that agent only believes ϕ that when is ϕ true”).

⁴Phenomena of referential opacity may require us to allow distinct beliefs in the same proposition even for a single subject, world and time: for instance, the belief that Hesperus shines and the belief that Phosphorus shines or Oedipus’ father that someone killed his father and his belief that he himself killed his father. The picture above can be amended in several standard ways to accomodate the idea.

knowledge without belief or high confidence, or, alternatively, to cash out a notion of *being in position to know* that is factive but does not require belief.⁵ That seems to me crucial in order for either notion to play a central normative role. For, I would think, what one ought to believe or intend should by and large not be hostage to the vagaries of one's psychology. Be that as it may, the normative role of knowledge is not our main concern here. So I leave the generalised version of safety aside.

1.2 A non-reductive account of safety

What, then, makes two beliefs peers? We can gesture at various factors, as I did above, such as same or similar circumstances, subject-matter, agent, way of being formed, bases and so on. In fact many attempts at formulating necessary or sufficient conditions for knowledge in the literature, in particular safety conditions, can be recast as generic safety supplemented with a specific account of what makes two beliefs peers. For instance, Hawthorne (2004, 56n17) suggests that one fails to know if “[one] forms [one's] belief in a way that could very easily have delivered error”. That is equivalent to generic safety supplemented with the idea that a belief's peer are beliefs formed in the same way in circumstances that could very easily have occurred. There is little hope, however, to provide a clean set of necessary and sufficient conditions for peerhood expressed in terms that are readily understood and easily applicable to various test cases. Notions like those of “bases”, “ways of forming beliefs”, “circumstances”, “environment”, “nearby worlds” are not ones on which we have a firm grip. Counterfactuals are better grasped, but I doubt that they provide plausible safety conditions.⁶ Thus these notions are not clearly better primitives than an unspecified relation of epistemic similarity. We should not treat conditions expressed with these notions as anything more than rough and ready guides to the relevant similarity relation. In fact, I suspect some of them, like the requirement that peers occur in circumstances that “could have easily happened” are plainly wrong.

A better approach is to use facts about knowledge to spell out *what the relevant relation must be like, if generic safety is necessary for knowledge*. A good guide are what we may call *margin for error conditionals*: conditionals of the form ‘if this belief is false then that belief is not knowledge’. Consider for instance the two beliefs in the *Sight* case. It is pretty clear that *if the belief formed looking at a slightly longer pen is false then the belief in the actual scenario does not constitute knowledge*. So the beliefs are peers. By contrast, consider a pair of beliefs that your boss is in the building, one acquired by sight alone and the other by hearing alone. There is no temptation to

⁵For defences of the idea that knowledge does not entail belief, see notably Woozley (1953), Radford (1966), Lewis (1996), Murray et al. (2013), Myers-Schulz and Schwitzgebel (2013)—though see Rose and Schaffer (2013) and Buckwalter et al. (2015) for further discussion. For the notion of being a position to know see Williamson (2000, 95).

⁶See section 2.3 for an argument that mistakes that “could easily have happened” are not peers.

judge that if one is mistaken, the other fails to constitute knowledge. Thus they do not count as peers.⁷

Does the approach trivialise the safety condition? No. First, note that an account of knowledge framed in terms of a notion understood in terms of knowledge itself need not be trivial. That is well illustrated by standard models for epistemic logic. Roughly put, these models say that one knows p just if p holds at all *accessible* worlds. But what is accessibility? One answer explains it in terms of knowledge: the accessible worlds just are those compatible with everything you know (Hintikka, 1962, 45). The resulting account is far from being trivial. Indeed, it is arguably false, as it entails that one knows all the logical consequences of what one knows. Similarly, the generic safety account is not trivial. It entails that knowledge requires truth, for instance. It entails that it is possible not to know all logical truths. Suitably extended, it has implications for the question whether knowledge is preserved under competent deduction (Williamson, 2009a). Paired with the use of margin for error conditionals as a guide, it has further significant implications to which we return shortly. Second, even though the margin for error conditionals guide does shield the account from counterexamples—faced with any putative case of knowledge despite a sufficiently similar mistaken belief, the guide recommends deeming the latter not relevantly similar—that does not mean that the account is unfalsifiable. Rather, it is falsified indirectly: if the only way to square facts about knowledge with generic safety is to adopt a trivial or excessively gerrymandered notion of peerhood then the account is falsified. For instance, we may confidently judge that, insofar as any natural or relevant respects of similarity are concerned, a certain belief a is *at least as similar* to some belief b as some belief c is to some belief d . If the peerhood relation we end up with violates that constraints, counting d as a peer of c but not a has a peer of b , then the account is disconfirmed. Conversely, if the account yields a fairly natural peerhood relation and if that relation can in turn be put to further theoretical use, then the account is confirmed.

Ultimately we will not understand safety in terms of (epistemic) *similarity* but rather in terms of (epistemic) *normalcy* and (epistemic) *differentiation*. But that does not alter the methodological picture just sketched.

The margin for error conditionals guide yields two significant and perhaps surprising lessons. The first is that *truth is not a sufficient difference*. In cases like *Sight*, it is pretty clear that if one belief is false then the other is not knowledge. That is so even if the latter is true. Hence the mere fact that one belief is true and the other is false is not enough to make them epistemically unlike each other. Consequently, *knowledge is not true belief*. Since some true belief have mistaken peers,

⁷Of course the conditionals are not a hard-and-fast rule. There may be all sorts of indirect links between the fact that one belief is false and the fact that another fails to constitute knowledge. For instance, it may be that if today is not Ali's birthday then Ben's belief that there is a party tonight is mistaken. It follows that: if Ali's belief that today is his birthday is mistaken, then Ben's belief that there is a party tonight is mistaken. However, that does not indicate in the least that Ali's belief that today is his birthday is epistemically like Ben's belief that there is a party tonight.

not all true beliefs constitute knowledge. Moreover, peerhood is not identity: some beliefs are like other beliefs. In fact it is hard to imagine any belief with no peer other than itself.

The second is that *knowledge is not a sufficient difference*. Consider a range of cases in which I judge by sight whether a pen is less than 2 inches long. In the first case then pen is 0.1 inch long, in the last 2 inches long. Its size increases by tiny steps from one case to then next, everything being otherwise as equal as can be. Clearly I do know at the first but not at the last—given my unexceptional eyesight, when the pen is very close to 2 inches long I cannot know whether it is at least 2 inches long. So some case of knowledge is next to a case of non-knowledge. Moreover, every case is arguably a peer of the next—if my belief in one case is mistaken then I do not know in the previous case. So some knowledgeable belief has a non-knowledgeable belief as peer.⁸ Thus the mere fact that two beliefs differ in knowledge status is not enough to make them epistemically unlike each other. That rules out a range of fairly trivial accounts of peerhood. For instance, one may propose that all knowledgeable beliefs are peers and all non-knowledgeable beliefs are peers. The account is guaranteed to be extensionally adequate: since all knowledgeable beliefs are true, it ensures that all knowledgeable beliefs are safe, and since some non-knowledgeable beliefs are false, it ensures that all non-knowledgeable beliefs are unsafe. The margin for error conditionals guide excludes such an account. In doing so it ensures that our peerhood relation is more interesting. It also leaves it open whether a satisfactory relation can be found. But these are two sides of the same coin.

One may worry that no property makes a sufficient difference. That is not so. *Being exclusively based on different sensory modalities* is a good candidate, for instance. As we hinted above, it is hard to imagine a pair of a belief exclusively based on sight and the other on hearing such that if the first is false then the other cannot be knowledge.⁹

1.3 A diagnosis of the Gettier problem

We can now state an important application of safety, namely a *diagnosis of the Gettier problem*. We can show that knowledge is not the conjunction of a non-factive condition and truth.¹⁰ Let *J* stand for any such condition and *K* stand for knowledge. We claim:

⁸I use “knowledgeable belief” as a shorthand for “belief that constitutes knowledge”.

⁹Unless the case is set up so that there is an irrelevant link between the two facts. See note 7 above.

¹⁰There is a trivial sense in which knowledge is *logically equivalent* to a conjunction of non-factive condition and truth. Say that *Xp* holds iff *p* is false or one knows *p*. Given that knowledge entails truth, knowing is trivially equivalent to *X* and truth (compare Howard-Snyder et al., 2003, Huemer, 2005). But just as, intuitively, *being a mammal or a planet* is too gerrymandered and irrelevant to be part of an account or explanation of what *being a cat* is, *X* is too gerrymandered and irrelevant to be part of an account or explanation of what knowledge is (compare Bailey, 2010). It is in this stronger, more substantial—and admittedly hard to characterise—sense that knowledge is not the conjunction of a non-factive condition and truth. That constraint is captured below the requirement that *J* be epistemically relevant, which underpins premise (3).

1. For some p there is a case c such that: $Jp \wedge \neg p$.
2. There is a case c' and a proposition q such that:
 - a. q holds in c' , and
 - b. one's belief that $p \vee q$ is a peer of one's belief that p in c .
3. $J(p \vee q)$ holds in c' if Jp holds in c .

It follows from safety that knowledge is not the conjunction of J and truth. By (1) and (3), $J(p \vee q)$ holds in c' . By (2a), $p \vee q$ holds in c' . By (1), (2b) and safety, $K(p \vee q)$ fails in c' . So in c' : $J(p \vee q) \wedge (p \vee q)$ holds but $K(p \vee q)$ does not. Hence knowledge is not J and truth. Generalising over J : knowledge is not the conjunction of a non-factive condition with truth.

Premise (1) follows from the non-factivity of J . Premise (2) is the substantial one. We claim that it is possible to build a case in which one stands to $p \vee q$ just like one stands to p in the original case in all epistemically relevant respects. An easy way to do so—though by no means the only one—is to let q be a proposition that is neutral to the original subject, such as the number of stars is even, and make our new subject form beliefs in the weaker disjunction instead than in p . Note that, since truth is not a sufficient difference, the fact that p is false but $p \vee q$ is true does not prevent them to be peers. Note also that for any natural J , there will be more than one case in which one has J towards a falsehood; we are free to pick ones for which building the alternative case is easiest. Gettier (1963) in effect builds the alternative case by adding a step of deductive inference to the original case; we can now see that that is inessential. Finally, premise (3) is motivated by premise (2) and the fact that J is a candidate condition for knowledge. If J is a candidate condition for knowledge, then it is one of the epistemically relevant respects that are preserved across c and c' . Thus one should bear J to $p \vee q$ in c' if one bears J to p in c .

The diagnosis is not new. Versions of it have been put forward by (Sosa, 1985, 239–240), Sturgeon (1993), Zagzebski (1994) and Merricks (1995).¹¹ What is significant here is that we recover the diagnosis from the safety condition. The diagnosis is independently plausible, as it partly explains the history of post-Gettier debates: philosophers typically put forward conditions for knowledge that fell short of entailing truth and repeatedly found them insufficient.¹² That provides inductive support for the safety condition.

The diagnosis is not banal either. A popular alternative diagnosis is that a belief constitutes knowledge only if it is not an accident or a piece of luck that it is true (Unger, 1968; Pritchard, 2005, chap. 5). The non-conjunctive diagnosis is arguably more precise. If I believe that a ticket in a big, fair lottery will lose, then there is a clear sense in which it is not a piece of luck or an accident

¹¹ See Moon (2012) and the references in the previous footnote for further discussion.

¹² To be clear, I do not think that the diagnosis is the whole story of the post-Gettier literature. A number of conditions failed by failing to be necessary instead, others by simply being too obscure to be evaluated. Another important part of the story is that the reductive project was ill-conceived (Williamson, 2000, chap. 1).

that my belief is true, if it is. Without further specification the “anti-luck” diagnosis does not rule out the proposal that knowledge is true belief that is highly likely to be true. By contrast, since being highly likely to be true does not entail truth, the non-conjunction diagnosis straightforwardly rules it out.

Most importantly for our purposes, the diagnosis is an Infallibilist one, as the authors cited above all note. If we call properties of beliefs “fallible” if false beliefs can have them, the diagnosis is that knowledge is not the conjunction of a “fallible” property and truth. Since it is not truth alone either, knowledge requires a truth-entailing property stronger than truth itself. In that sense, knowledge requires an “infallible” belief.¹³

Here we state the Infallibilist upshot in terms of safety. The safety condition says that knowledge requires a belief that *is epistemically unlike any false belief*. Since, moreover, truth alone is not an epistemically sufficient difference, knowledge requires a belief that is *epistemically unlike any false belief in respects stronger than simply being true*. Put otherwise, knowledge requires a belief of a kind such that no belief of that kind could be false. There is a clear sense in which such a belief is infallible. Hence knowledge requires infallible belief.

Let us sum up the dialectic so far. A range of cases, including but not limited to Gettier-style cases, motivate a safety condition on knowledge. The safety condition is further supported by the fact that it yields an independently plausible diagnosis of the Gettier problem. But the safety condition is just the claim that knowledge requires infallible belief. So we have good reason to think that knowledge requires infallible belief.

2 Objections

2.1 The Old Evil Demon problem

The first challenge to Infallibilism comes from the consideration of sceptical scenarios. For most of things we know there are possible situations we believe as we do, things appear as they do and yet we are mistaken. Descartes’ Evil Demon hypothesis provides one such situation for pretty much every “external world” belief we have. When presented with these scenarios a tempting reaction is to say that our beliefs “could be false”. But what is meant by that? Obviously, it is neither necessary nor sufficient that *their content* could be false. For it is necessary that Hesperus is Phosphorus and

¹³It is customary to call “factive” an attitude that one can only have to truths (Williamson, 2000). By extension we may call “factive” a truth-entailing property, that is a property only true beliefs can have. Truth is the weakest factive property. Being true and reasonable is a conjunctive factive property. Here we call a belief “infallible” if it has a factive property that is *stronger* than truth and not merely the conjunction of some epistemically relevant property and truth. The idea is that the belief has a property that non-trivially entails truth and that no belief of its kind is false. Some may feel that the notion of “infallible” belief should be something stronger than that. But it is unclear to me what that could be. See section 3 for further discussion. Thanks to Chris Kelp here.

it is contingent that I am currently in pain and yet a belief in the first is presumably fallible in the relevant sense and a belief in the second is not obviously fallible in the relevant sense. Rather, the idea seems to be that one could have the *same state of belief*, or a *state of belief sufficiently like it*, that would be false.

Now we may grant that there is a sense of “internal” under which the ordinary person’s belief and the Evil Demon victim’s counterpart belief are “internally alike”. Still, the fact that two beliefs are alike in some respect and one of them is false is not enough to say that the first one “could have been false” in a some interesting sense. Call two (token) beliefs “daymates” iff they are held on the same day. The fact that some belief has a false daymate does not warrant saying that it “could have been false” in a sense that has any relevance to epistemology. Similarly, noting that some belief has a false internal counterpart does not by itself warrant saying that that belief “could have been false” in an epistemically interesting sense. There must be something more to the idea.

One simple suggestion is this: internally alike beliefs are *epistemically alike*, that is, alike in respects that matter to knowledge. That is:

Peerhood internalism Two beliefs that are internally alike are peers.

For instance, some may be tempted to say that an Evil Demon victim’s belief have “the same basis” as that of the ordinary person and that a belief’s basis is all that matters to whether one knows. Since for most of our beliefs we can devise a mistaken Evil Demon-style counterpart, most of our beliefs have internally alike false counterparts. So Peerhood Internalism and Generic Safety together entail that we know next to nothing. Hence Internalist Infallibilism are condemned to either implausibly deny that Evil Demon scenarios are possible or to embrace widespread scepticism. The dilemma is one traditional philosophers like Descartes and Hume appeared to face and there are good reasons to think that they did assume a Internalist Infallibilist conception of knowledge.¹⁴

The dilemma is largely responsible for Infallibilism’s bad press. Fortunately, however, Peerhood Internalism is wrong. An ordinary subject’s belief differs in drastic ways from those of an Evil Demon victim: they involve different circumstances, cognitive processes, causal history and so on.¹⁵ Once we reject Peerhood Internalism, we can both claim that knowledge requires infallible belief and acknowledge the possibility of Evil Demon scenarios without embracing scepticism.

That goes some way towards showing that *Internalism*, not Infallibilism, is what is wrong with the traditional view of Descartes, Hume and others. Not all the way, though. For it is still possible that the dilemma arises from Infallibilism paired with *any* reasonable conception of peerhood, internalist or otherwise.¹⁶ The strongest way to press that objection is to consider phenomena involving chance. We turn to it in the next section.

¹⁴Dutant (2015) argues that virtually all Western philosophers up to the XXth Century did.

¹⁵It is debatable whether one forms beliefs in one’s dreams (Sosa, 2007, chap. 1). But if one does, these beliefs also differ in drastic ways from the ones we form when awake.

¹⁶Thanks to Mikkel Gerken here.

2.2 The Widespread Chance Objection

Consider the following pair of cases:

Lottery. You have a ticket in a fair lottery with a million forty thousand tickets and one winner. I believe that your ticket will not win. I would believe likewise if your ticket was the winner.

Coins. You are throwing twenty fair coins. I believe that they will not all come up heads. I would believe likewise if they were about to.¹⁷

The trick, of course, is that the chances are roughly the same: there is a one in $2^{20} = 1040576$ chance that twenty fair coins will come up heads. Yet there is some pressure to say that I cannot know in *Lottery* but I can know in *Coins*. These are somewhat intuitive verdicts. Picturing the two situations may help. For the first, imagine that each inhabitant in a city of the size of Austin holds a ticket. You may be inclined to say that we cannot know of anyone of them that they will not win. For the second, imagine that two people hold a coin on the tip of each of their fingers and that they simultaneously toss all of them. You may be inclined to say that we can know that not every single coin will fall on heads. Of course, the intuitive judgements are not firm here. There are also some specific reasons to doubt them: the first involves large numbers over which we have little intuitive grasp and the second may trigger a gambler fallacy.

Intuitions aside, however, there are well-known considerations in favour of the verdicts suggested. On the *Lottery*: if we say that I know, then since my standing to every other ticket appears relevantly the same, we seem forced that I know of every losing ticket that it will lose. But I then know enough to deduce which is the winner. Since I obviously cannot know which it is, we would have to say that deduction does not preserve knowledge or that I do not stand in equal ways to each ticket. The result is unpalatable and gives some support to the verdict that I cannot know.¹⁸ On *Coins*: many ordinary processes appear to work essentially like *Coins*. Pouring milk in tea, for instance, involves many small collisions such that, if they all unfolded in some specific way, the milk could stay lumped together in the middle of the liquid instead of spreading more or less evenly. And every small collision has some chance of going in the requisite way. Granted, there are differences. The milk case involves a vastly greater number of chancy events and over time they are not wholly independent in the way the coin tosses are.¹⁹ Still the central features are the same: a number of fairly independent chancy events which, if they all unfolded in some specific way, would result in a surprising outcome.²⁰ Saying that in all such cases one cannot know that the surprising

¹⁷Compare Vogel (1999, 165) and Bacon (2014).

¹⁸See Hawthorne (2004) and the references therein for an extensive discussion.

¹⁹Thanks to Levi Spectre here.

²⁰See Hawthorne (2004, 4–5).

outcome will not occur has sweeping sceptical implications. A wide range of phenomena are as chancy as pouring milk into tea. Knowledge of the past will be just as affected as knowledge of the future: learning that some milk has been poured in tea a couple of days ago would not be enough to know that it has mixed. Straightforward perceptual knowledge would arguably be affected too. The reflection of light on a river, for instance, also involves a great number of chancy events. If they unfolded in some very specific way they would reflect light just as a fish swimming in the clear water does. So we would be pressed to conclude that we cannot ever know that a fish is swimming in water either. The sweeping sceptical consequences are wildly implausible. This suggests that for some variant of *Coins* I can know that not all coins will fall head. Numbers matter, as I obviously cannot know that in a two-coins version of the case. But at some suitable size it must be possible to know that not all coins will fall head. Yet for a correspondingly-sized *Lottery*, the considerations above would still support the conclusion that I do not know. Put together, these considerations suggest that in some *Lottery/Coins* pair the contrasted verdicts are correct.²¹

The contrasted verdicts are a puzzle for everyone.²² Think of the *Lottery* as involving a giant, million-sided coin. The coin is perfectly symmetric and may fall on any of its sides with equal likelihood. The verdict is that I cannot know that it will not fall on any particular side. The *Coins* case involves twenty coins; there are about a million possible combined outcomes. Each combined outcome may occur and each is equally likely. The verdict is that in some situations I can know that some outcome will not occur, namely, that they will not all fall heads. Why the difference?

The *Lottery* verdict is grist to the Infallibilist's mill, of course. It is easy to claim that the beliefs I would form about each ticket are peers. Since one of them is false, Generic Safety delivers the desired result that none of them would constitute knowledge. The *Coins* verdict is the troublesome one. It would be natural to treat the beliefs I would form in each possible outcome as peers. But that pushes Infallibilists into a sceptical dilemma: either they implausibly claim that the chancy outcomes are impossible, or they embrace widespread scepticism. The dilemma cannot be blamed on Internalist assumptions.

Infallibilists have only one way out of the dilemma. They must say that the *Coins* beliefs across different outcomes are not alike.²³ Let p be the proposition that not all coins will fall heads. Suppose in the actual outcome the frequency of heads is exactly half. Then the belief in p I actually form is most similar to the beliefs in p I would form if the outcome had a frequency of heads of 1 above or below half. It is slightly less similar to the beliefs I would form if the outcome had a frequency of 2 above or below half. It is much more dissimilar to the beliefs I would form if all or none were heads. On that picture the outcome plays a role analogue to that of the "circumstances" or "environment" in the Gettier literature: the all-heads scenario belief differs from the half-heads scenario one like a

²¹ See Hawthorne and Lasonen-Aarnio (2009).

²² The puzzle is at the heart of Hawthorne (2004).

²³ That is the line taken by Williamson (2009b).

fake-barns scenario belief differs from an ordinary-barns scenario one. By contrast, in *Lottery* the belief I form about your ticket is equally similar to any belief I would form about another ticket. Thus my lottery belief has a mistaken peer while my coins belief does not.

To be clear, the suggestion is not that we are allowed to treat low-chance outcomes as dissimilar. *Every* particular outcome is as unlikely as any other. The suggestion is that the *Coins* case has a structure that the *Lottery* lacks and that that structure affects peerhood.

But wait! For all we have said the lottery winner is decided *by tossing twenty coins!* If is do we implausibly say that the contrast disappears? No. Remember that peerhood relates *beliefs*, not worlds. Imagine two subjects looking at the same lottery mechanism from different sides. From Lotta's side one can only see a winning number being displayed. From Colin's side one can see twenty fair coins whose outcome gets translated in a digit winning number. Lotta holds ticket #169706 which actually translates the all-heads outcome. The actual outcome is HTTTH-HHHTH-HHTTH-TTHTH which is 1 above half heads and translates to ticket #604294 winning. On the suggested picture, Lotta's believing that her ticket will lose would be relevantly like her mistakenly believing that #604294 will lose, but Colin's belief that not all heads will come is unlike him believing that the outcome HTTTH-HHHTH-HHTTH-TTHTH will not occur. So the fact that twenty coins generate a richer structure of possible outcomes than a single million-sided coins is only half of the story. The other half is that the structure is 'picked up' by Colin's possible beliefs, in the sense that it creates epistemically relevant differences among them, while Lotta's beliefs are 'screened off' the underlying structure. Colin's belief that not all heads will come is marked out as a special belief; Lotta's belief that #169706 will not come out is not.

To be clear the above is only meant to *claim* that Colin's beliefs, but not Lotta's, pick up the underlying structure. It is not meant to *explain* why they do so. That is a further question and one that Infallibilists need not answer straight away.

The problem with granting knowledge in *Lottery* was that symmetry and closure seemed to force us to countenance implausible knowledge of unlikely outcome. Do we not face the same problem here?²⁴ Not necessarily. Consider a few variants.

Colours. I have colour stickers of a million different shades. I put the stickers on the coins so that each possible outcome consists in all the coins coming out the same shade, for some shade.

Do we have to say that for any colour *C*, a true belief that not all coins come out *C* constitutes knowledge? No. For now a belief that not all coins come out heads is like a belief that not all coins come out red #162434, a belief that that not all coins come out blue #78977, and so on and so forth for every colour. Since one of these beliefs is mistaken my belief is not safe.

²⁴The objection is pressed by Hawthorne and Lasonen-Aarnio (2009) and Sharon and Spectre (2013).

Testimony. As in *Colours* except that each colour is visible to one and only one of a million subjects. The subjects share a common numbering of the outcomes but otherwise each is only aware of the coins having a neutral and a coloured side. Each subject believes that not all coins will fall on (what they see as) the coloured side. They try to pool their knowledge together by telling each others the outcome numbers they think will not come out.

Each subject is very much like our original subject in the simple heads / tails scenario. Yet if we say that each knows and that they can pull their knowledge together, they could deduce from what they know highly unlikely claims. Two diagnoses may be offered here. One says that the presence of many subjects affects peerhood. Each subject's beliefs that not all coins will be "coloured" is like each other's. One of them is mistaken so none is safe. The other is to say that subjects cannot pool knowledge together in that situation. The true belief that one acquires from testimony from a subject who is right is relevantly like the mistaken belief that one acquires from a subject that is mistaken. Either option blocks the possibility that subjects acquire and pool together enough knowledge to deduce highly unlikely claims about the outcomes. We need not adjudicate between the two diagnoses here.

Repeated runs. I toss twenty coins repeatedly for many times. At each time I believe that not all heads will come up.

Do we have to say that I can know of each toss that it will not result in all heads? With enough runs that would leave me knowing enough to deduce some highly unlikely claims. But again, we should say that the repeated runs affect peerhood. For now the relevant peers of my belief in the first run are those involved in similar *series* of runs. If the actual series is typical and big enough, it will be similar to series where some run is all-heads. Hence my belief will have a mistaken peer. Thus when my belief is part of a broader series, I cannot know that the present run will not be all heads, and I cannot accumulate knowledge over the series to be in position to deduce some highly unlikely claim. It follows that whether my belief in a first run constitutes knowledge depends on whether I will later engage in similar runs. But that is just one more way in which whether my belief is knowledge depends on the future. There is a residual puzzle, however: what if the actual run turns out to only have outcomes with a frequency of about half heads? Can I then know of each run that it will not come out all heads? No—but details must await the next section.

Let us take stock. A wide array of things we seem to know are in effect the outcome of chancy processes. Denying that we can know them forces us into an implausibly sceptical position. Accepting that we can know them raises the difficulty faced by those who grant knowledge in lottery cases: namely, that we know enough to deduce highly unlikely claims. We have sketched a Infallibilist path between both evils. Ordinary chancy process involve enough structure for epistemically

relevant differences among beliefs to arise from differences among their outcomes. The differences allow well-positioned subjects to know that some chancy outcomes will not occur. However, there are holistic limits on leveraging these differences for acquiring or pooling knowledge of things that have a high chance of not occurring. The picture sketched leaves several questions unanswered and requires further exploration. But it is viable so far and addresses what is challenge for everyone.

2.3 Specific Counterexamples

Several counterexamples have been raised against Safety. To diagnose them it is best to draw one further lesson from the *Coins* case. Consider this variant:

Outlier. I toss twenty fair coins. Even though I know that they are all fair I believe that most will fall on heads. As it happens, all of them do.

In the original case we argued that my beliefs that not all coins will fall on heads held in half-frequency and in all-heads scenarios would be unlike each other. By parallel reasoning we could argue that believing that most coins will fall on heads in the all-heads scenario is unlike believing it in half-frequency or most-tails scenarios. If we do so it will seem that my belief in *Outlier* is safe. So can I know that most fair coins will fall on heads? It would seem not. Can we get that verdict?

One option is to bring in further considerations. Since I know that the coins are fair, my belief that most will fall heads seems to be unreasonable or defeated. If knowledge requires reasonable or undefeated belief in addition to safety, then I do not know. That line of argument is not open to Infallibilists who take infallible belief to be sufficient for knowledge, however.²⁵ A more congenial one is to argue that the belief has mistaken peers other than a all-heads belief in half-frequency or most-tails scenario. For instance, since I am symmetrically situated with respect to heads and tails outcome, one could argue that my actual belief is like a mistaken belief that most will fall on *tails* formed in the same circumstances. The solution is defensible but I think a more general diagnosis is available.

The suggestion I want to pursue instead is that *peerhood is not symmetric*. Beliefs formed in middling scenarios, where the frequency roughly matches the coin bias, do not have beliefs formed in outlier scenarios as peers. But beliefs formed in outlier scenarios do have beliefs formed in middling scenarios as peers. Thus it is open to say that in the original *Coins*, my middling scenario belief does not have outlier peers while in *Outlier*, my outlier scenario belief has middling peers. The suggestion has to be motivated, however. Moreover, similarity is a symmetric relation. So the suggestion must replace the similarity picture of peerhood. Two alternative pictures can be offered. They are not exclusive of each other.

²⁵It also assumes that safe belief is not enough for being undefeated or reasonable, which some Infallibilists will want to avoid.

The first is the *normalized-variant* picture. The peers of a belief are variants of it, but only variants that are *at least as normal* as that belief.²⁶ In the *Coins* scenario, there is a well-defined sense in which the middling outcomes are normal, namely, they are *random*. Random outcomes are not more likely than non-random ones. But they are more *typical* in the sense that they have many properties that are highly likely. For instance, if we toss twenty fair coins, the following are highly likely: we get at least one tail; we get two heads next to each other; we get between one and two thirds of tails; and so on. The all-heads outcome has few of these properties, the middling outcome given in the previous section has many. An outcome were heads and tails alternate regularly (HTHTH...) has a half-head frequency but it does not instantiate many highly likely properties—that there are two heads in a row, that there are two tails in a row, that there are three heads in a row, and so on—and is less random. The suggestion, then, is that beliefs have peers in more random outcomes. Since being more random is asymmetric, peerhood is asymmetric.

Note that randomness is relative to how outcomes are described—recall the *Colours* case: every outcome is random when described in some colour scheme and non-random when described in others. As before, that is accommodated by treating comparative randomness relation in question here holds among beliefs, not outcomes themselves. And as before, the above is not meant to answer the question *why* a given belief's peerhood picks up one description scheme over another.

The second is the *distinction* picture. In academia we practice evaluation "by peers". What is meant by that is not that we are evaluated by academics of *equal* standing, but by academics of *equal or more distinguished* standing. The suggestion is to think of belief peerhood in an analogous manner. We think of epistemically relevant features of beliefs as *marks of distinction*. Initially, so to speak, one's standing to some proposition *p* is undifferentiated. It does not stand out from how one stands to any other proposition. But as one sees that *p*, for instance, one's standing to *p* acquires a mark of distinction. It nows stands apart from any belief that is not visually based. As one also hears that *p*, one's standing acquires a further mark of distinction, and is now also unlike any belief that lacks an auditory basis. And so on. To be clear, the idea is not that the more marks of distinction a belief has, the better. Rather, the idea is that the conjunction of the marks a belief puts it in an epistemically relevant kind, and if that kind is one that includes only true beliefs, the belief is safe. On that picture, if *A* and *B* are marks of distinction, a belief with mark *A* only has a belief with marks *A* and *B* among its peers but not the opposite. So peerhood is asymmetric. Applying that picture to the *Coins* and *Outlier* variants, the suggestion is that occurring in middling outcomes is a mark of distinction while occurring in outlier ones is not.²⁷

The asymmetric account solves the remaining puzzle with *Repeated runs*. A long series of close-to-half outcomes is abnormal. Hence beliefs formed in such a series have peers formed in

²⁶See Dutant (2009) and Goodman and Salow (ming) for related suggestions and Smith (2010, 2016) for a use of a notion of normalcy in an account of justification rather than knowledge.

²⁷Thanks to Mona Simion here.

series that include mostly-heads outcomes. That is why even in such a series one cannot know that no all-heads run will occur.

On to the proposed counterexamples to safety. They have a common structure: something abnormal almost happened; it if had, one would be mistaken; but it has not and it seems that one knows. Here are summary versions:

Water. You believe that the glass you see contains pure water. If the person nearby had not just won the lottery, they would have poisoned the glass. (Neta and Rohrbaugh, 2004)

Halloween party. Your friend truthfully points you towards the party. However, you almost pick up a disguise that would have made him lie to you. (Comesaña, 2005)

Unstopped clock. You look at a functioning clock at 8:22. However, had you come by the clock at a different time, your arch-nemesis would have made it display 8:22 anyway. (Kelp, 2009)

Atomic clock. You look at a functioning clock at 8:22. The clock stands next to a radioactive isotope that can decay at any moment and make the clock display 8:20 instead. (Bogardus, 2014)

It is not wholly clear that one knows in those particular cases. But let us assume that one knows in some case with a similar structure. Targeting a widespread notion of Safety, these authors argue that in one case one knows something even though one could easily have believed something false on the same basis. Targeting Generic Safety, one could argue that one's belief in such cases is relevantly like the mistaken belief they would have in the abnormal circumstances that almost occurred.

It should now be clear how Infallibilists can handle the cases. Insofar as they want to uphold the verdict that one knows, they should say that the mistaken belief you would have in the "close" error scenario is not a peer of your actual belief. One way to defend that conclusion is to extend the asymmetric diagnosis of *Outlier* to such cases. For instance, one may say that a belief's peers are only its normalised variants and the close error scenarios involve more abnormal beliefs than the actual one.

A final word of caution. The non-reductive lessons of section 1.2 stand. We should not go around applying an untutored notion of normality to various cases in order to test the account. Rather, we should call our peerhood relation *epistemic comparative normality* or *epistemic differentiation* and figure out what it must be like for Generic Safety to succeed. In some highly constrained cases like *Coins* we find a well-established notion—randomness—that provides a well-defined and robust notion of normality. We can use it as a guide, hoping for some fruitful convergence between our account of knowledge and other areas of inquiry. In some other cases we have very firm judgments about knowledge that we can use to probe what a suitable normality-based peerhood relation

would be like. In further cases, we should rather use what look like natural extensions of our relation and see whether they generate promising verdicts. Some of the counterexamples above fall in the third category. For instance, as Bogardus spells his case out, the radioactive isotope has been exceptionally inactive in the last minute. That is structurally like a huge number of fair coins landing on heads. It would be natural to extend our account of *Outlier* here by saying that a situation in which the isotope has emitted radiation is more normal. We would thus conclude that in that scenario you fail to know. Similarly, obvious parallels between *Lottery* and *Water* may be used to argue that you fail to know in the *Water* scenario. *Unstopped clock* is harder to adjudicate: on the one hand, the subject's picking up of a time to pass by the clock is lottery-like, but on the other hand, a nemesis' tampering with the clock would make for circumstances less normal than the actual ones. In *Halloween party* there is less obvious obstacle to treating the error possibilities along the lines of a all-heads outcome.²⁸

3 Infallibilism or Ecumenism?

Infallibilists claim that knowledge requires infallible belief. We have defended the claim by explaining a range of cases and diagnosing the Gettier problem. Evil Demon scenarios, widespread chance and some counterexamples all suggest that the requirement is too strong. We have shown that the Infallibilist view can steer away from these consequences. Thus suitably understood the Infallibilist claim is a viable and even attractive option.

The claim is in fact so appealing that it is *de facto* endorsed by many epistemologists. Still even most of those who do would describe themselves as "fallibilists" in some important sense. They are *Ecumenists*: that combine the Infallibilist requirement on knowledge with some significant fallibilist claim elsewhere. Let us conclude by briefly considering such views.

Most epistemologists who call themselves "fallibilists" explain the label in either of three ways. One appeals to *ruling out*: one may know *p* even though one has not "ruled out" every possibility that not-*p*. Another appeals to *modals* or *possibility*: roughly, one may know *p* even though, for some *q* incompatible with *p* (not-*p*, for instance), it might be that *q* or it is possible for one that *q*. Another appeals to *some feature of belief*: one may know *p* even though the *evidence*, *reasons*, *basis* or *justification* has for one's belief are ones a false belief (or a belief that otherwise falls short of knowledge) could have.²⁹ We may leave the first aside: it is ambiguous and its best versions boil

²⁸Szabó Gendler and Hawthorne (2005) discuss a range of puzzling cases for a safety account of knowledge. None of the cases are clear-cut counterexamples to generic safety. They provide a good field for further probing the peerhood relation. We leave that task to further work.

²⁹For the first see Goldman (1976), Stine (1976), Dretske (1981), Lewis (1996), Holliday (2015). For the second see notably Dougherty and Rysiew (2009), Dodd (2011), Fantl and McGrath (2009), Anderson (2014), Worsnip (2015). For the third see notably Cohen (1988), Reed (2002), Pryor (2000).

down to either the second or third.³⁰ The third may be paired with Internalism, if one claims that internally equivalent beliefs share their evidence, reasons, basis or justification. But it need not to. Each of the statements appeals to a central notion. Let *epistemic possibility* and *epistemic belief-feature* be generic labels for the notions appealed to in the possibility and belief-feature statements, respectively.

The first task for Ecumenists is to say *what the notion denotes*. On one simple view one's evidence and one's reasons are just what one knows. Belief-feature Ecumenists who appeal to evidence or reasons cannot understand these notions that way. On one simple view the basis and justification of a belief are whatever factors determine whether that belief constitutes knowledge—what we have called its epistemically relevant features. Belief-feature Ecumenists who appeal to bases or justification cannot understand these notions that way. On one simple account, epistemic possibilities are just possibilities compatible with everything one knows. Epistemic-possibility Ecumenists cannot appeal to such a notion. On another account, epistemic possibilities are possibilities compatible with one's evidence. Epistemic-possibility Ecumenists who appeal to that notion are committed to Belief-feature Fallibilism framed in terms of evidence and face the latter's first task.

The first task can be tackled in a variety of ways—some of the methodological lessons of our discussion of peerhood apply here too. The notion can be treated as primitive, grasped through paradigm cases or heuristics or theorised more abstractly. Suitably framed it is not particularly troublesome. On the demanding way of framing it, we think of the notion—say, evidence—as being antecedently given. The task is then to give a suitable account for it. For instance, an Ecumenist may suggest that one's evidence consists in one's internal mental states and we can debate whether that is a correct theory of evidence. On the undemanding way of framing it, the Ecumenist is simply delineating a property of epistemic interest. If we insist that their notion is not evidence proper they can simply relabel it "internal evidence", say, and leave the matter at that. For present purposes the undemanding version is all Ecumenists need.

The second task for Ecumenists is to say *why the property is of epistemic interest*. That is the demanding one. To illustrate, let the "internal basis" of a belief be the set of experiences (internally construed) and background beliefs that are in some suitable way causally involved in the formation or support of that belief. Call two beliefs *internal-peers* if they have the same—or relevantly similar—internal bases. It is clear that many beliefs that appear to constitute knowledge have mistaken internal-peers. The question is, what does that tell us? Recall the notion of "daymates"

³⁰"Ruling out" a possibility is sometimes understood as *knowing that it does not obtain* (Dretske, 1981). On that reading, the statement is equivalent to the denial of single premise closure, which I take to be a marginal way of understanding the label "fallibilism". Alternatively, "ruling out" can be either taken as a primitive, or analysed. The standard analysis is that a possibility is ruled out if it is incompatible with one's evidence—or, on Lewis' (1996) version, incompatible with one's having the evidence one has. That makes the statement equivalent to the third option. If taken as primitive, let one's *possibility set* be the set of possibilities one does not rule out. The statement is then a version of the second option.

(section 2.1): many beliefs that appear to constitute knowledge have mistaken daymates but that is of little interest to epistemologists.

As I see it, to address the second task Ecumenists can take two main routes. One is to have their notion play a role in an account of knowledge. The second is to have it play a role in the norms for belief. The two are not exclusive of each other.

On the knowledge path the idea is to use epistemic possibility or the epistemic belief-feature in an account of knowledge. For instance, one may suggest that one knows iff all internal peers of one's belief *that occur in "relevant" or "close" scenarios* are true.³¹ The account amounts to factoring peerhood into two components: internal-peerhood and closeness. Whether it succeeds requires detailed discussion. If internal-peerhood requires the *same* experience, for instance, then it follows that any true belief about one's experience constitutes knowledge; the consequence is implausible. The broader challenge for Ecumenists who follow that path, then, is to articulate and defend an account of knowledge framed in terms of their favoured notion.

On the norm of belief path the idea is that epistemic possibility or the epistemic belief-feature play a central role in norms for belief. For instance, Belief-feature Ecumenist may suggest that whether a belief is justified supervenes on the epistemic feature: one's reasons, evidence, basis or justification for a belief are what determines whether a belief is justified. Possibility Ecumenists may argue that epistemic possibilities are those that matter to what it is rational for one to be certain of or to do.

On the knowledge path, I would suggest that the burden of proof now lies with Ecumenists. To defend that claim, however, it would be necessary to argue that current Ecumenists accounts of knowledge are unsatisfactory. That is well beyond the scope of this paper. My aim was to make Infallibilism palatable, not to refute Fallibilism. The normative path takes us to the other side of Infallibilism: the debate over norms of belief. Here the burden of proof is still on the Infallibilist side. But that is a topic for another day.

To sum up, I have argued that knowledge requires infallible belief. Ecumenists concede the point, but argue that most of our knowledge involve belief that is an important, but different, sense fallible. Their task is to delineate the sense in question and to explain why it is important. I have sketch two main paths for doing so: the knowledge path and the norm of belief path. The first one challenges Ecumenists to put forward an account of knowledge that matches the straight Infallibilist one. The second moves the discussion to the topic of Infallibilism about norm of belief. These are topics for the next chapters in an Infallibilist's guide.

³¹That is the kind of account found in Goldman (1976) and Lewis (1996).

References

- Anderson, C. (2014). Fallibilism and the flexibility of epistemic modals. *Philosophical Studies*, 167(3):597–606.
- Bacon, A. (2014). Giving your knowledge half a chance. *Philosophical Studies*, (2):1–25.
- Bailey, A. (2010). Warrant is unique. *Philosophical Studies*, 149(3):297–304.
- Bogardus, T. (2014). Knowledge under threat. *Philosophy and Phenomenological Research*, 88(2):289–313.
- Buckwalter, W., Rose, D., and Turri, J. (2015). Belief through thick and thin. *Noûs*, 49(4):748–775.
- Cohen, S. (1988). How to be a fallibilist. *Philosophical Perspectives*, 2:91–123.
- Comesaña, J. (2005). Unsafe knowledge. *Synthese*, 146(3):395–404.
- Dodd, D. (2011). Against fallibilism. *Australasian Journal of Philosophy*, 89(4):665–685.
- Dougherty, T. and Rysiew, P. (2009). Fallibilism, epistemic possibility, and concessive knowledge attributions. *Philosophy and Phenomenological Research*, 78(1):123–132.
- Dretske, F. (1981). The pragmatic dimension of knowledge. *Philosophical Studies*, 40(3):363–378.
- Dutant, J. (2009). Two notions of safety. *Swiss Philosophical Preprints*, 87.
- Dutant, J. (2010). *Knowledge, Method and the Impossibility of Error*. PhD thesis, University of Geneva.
- Dutant, J. (2015). The legend of the justified true belief analysis. *Philosophical Perspectives*, 29(1):95–145.
- Fantl, J. and McGrath, M. (2009). *Knowledge in an Uncertain World*. Oxford University Press.
- Gettier, E. L. (1963). Is justified true belief knowledge? *Analysis*, 23(6):121–123.
- Goldman, A. I. (1976). Discrimination and perceptual knowledge. *Journal of Philosophy*, 73(20):771–791.
- Goodman, J. and Salow, B. (forthcoming). Taking a chance on kk. *Philosophical Studies*.
- Hawthorne, J. (2004). *Knowledge and Lotteries*. Oxford University Press.

- Hawthorne, J. and Lasonen-Aarnio, M. (2009). Knowledge and objective chance. In Greenough, P. and Pritchard, D., editors, *Williamson on Knowledge*, pages 92–108. Oxford University Press.
- Hintikka, J. (1962). *Knowledge and Belief*. Cornell University Press.
- Holliday, W. H. (2015). Fallibilism and multiple paths to knowledge. *Oxford Studies in Epistemology*, pages 97–144.
- Howard-Snyder, D., Howard-Snyder, F., and Feit, N. (2003). Infallibilism and gettier’s legacy. *Philosophy and Phenomenological Research*, 66(2):304–327.
- Huemer, M. (2005). Logical properties of warrant. *Philosophical Studies*, 122:171–182.
- Kelp, C. (2009). Knowledge and safety. *Journal of Philosophical Research*, 34:21–31.
- Lewis, D. (1996). Elusive knowledge. *Australasian Journal of Philosophy*, 74:549–567.
- Manley, D. (2007). Safety, content, apriority and self-knowledge. *Journal of Philosophy*, pages 403–23.
- Merricks, T. (1995). Warrant entails truth. *Philosophy and Phenomenological Research*, 55(4):841–855.
- Moon, A. (2012). Warrant does entail truth. *Synthese*, 184(3):287–297.
- Murray, D., Sytsma, J., and Livengood, J. (2013). God knows (but does god believe?). *Philosophical Studies*, 166(1):83–107.
- Myers-Schulz, B. and Schwitzgebel, E. (2013). Knowing that p without believing that p. *Noûs*, 47(2):371–384.
- Neta, R. and Rohrbaugh, G. (2004). Luminosity and the safety of knowledge. *Pacific Philosophical Quarterly*, 85(4):396–406.
- Pritchard, D. (2005). *Epistemic luck*. Oxford University Press.
- Pritchard, D. (2012). Anti-luck virtue epistemology. *Journal of Philosophy*, 109(3):247–279.
- Pryor, J. (2000). The skeptic and the dogmatist. *Noûs*, 34(4):517–549.
- Radford (1966). Knowledge — by examples. *Analysis*, pages 1–11.
- Reed, B. (2002). How to think about fallibilism. *Philosophical Studies*, 107(2):143–157.

- Rose, D. and Schaffer, J. (2013). Knowledge entails dispositional belief. *Philosophical Studies*, 166(1):19–50.
- Sainsbury, R. M. (1997). Easy possibilities. *Philosophy and Phenomenological Research*, 57(4):907–919.
- Sharon, A. and Spectre, L. (2013). Epistemic closure under deductive inference: What is it and can we afford it? *Synthese*, 190:2731–2748.
- Smith, M. (2010). What else justification could be. *Noûs*, 44:10–31.
- Smith, M. (2016). *Between Probability and Certainty: What Justifies Belief*. Oxford University Press.
- Sosa, E. (1985). Knowledge and intellectual virtue. *The Monist*, 68:226–245.
- Sosa, E. (1996). Postscript to “proper functionalism and virtue epistemology”. In Kvanvig, J. L., editor, *Warrant in Contemporary Epistemology*. Rowman & Littlefield, Lanham, Md.
- Sosa, E. (1999). How to defeat opposition to moore. *Philosophical Perspectives*, 13:141–53.
- Sosa, E. (2007). *A Virtue Epistemology: Apt Belief and Reflective Knowledge*, volume 1. Oxford University Press.
- Stine, G. C. (1976). Skepticism, relevant alternatives, and deductive closure. *Philosophical Studies*, 29(4):249–261.
- Sturgeon, S. (1993). The gettier problem. *Analysis*, 53(3):156–164.
- Szabó Gendler, T. and Hawthorne, J. (2005). The real guide to fake barns: A catalogue of gifts for your epistemic enemies. *Philosophical Studies*, 124:331–352.
- Unger, P. (1968). An analysis of factual knowledge. *Journal of Philosophy*, 65(6):157–170.
- Vogel, J. (1999). The new relevant alternatives theory. *Noûs*, 33 (s13):155–80.
- Williamson, T. (2000). *Knowledge and its Limits*. Oxford University Press.
- Williamson, T. (2009a). Probability and danger. *The Amherst Lecture in Philosophy*, pages 1–35.
- Williamson, T. (2009b). Reply to john hawthorne and maria lasonen-aarnio. In Greenough, P. and Pritchard, D., editors, *Williamson on Knowledge*, pages 313–329. Oxford University Press.

- Woozley, A. (1953). Knowing and not knowing. *Proceedings of the Aristotelian Society*, 53:151–72.
- Worsnip, A. (2015). Possibly false knowledge. *Journal of Philosophy*, 112(5):225–246.
- Zagzebski, L. (1994). The inescapability of gettier problems. *The Philosophical Quarterly*, 44(174):65–73.