How Infallibilists Can Have It All

Nevin Climenhaga*

ABSTRACT

I advance a novel argument for an infallibilist theory of knowledge, according to which we know all and only those propositions that are certain for us. I argue that this theory lets us reconcile major extant theories of knowledge, in the following sense: for any of these theories, if we require that its central condition (evidential support, reliability, safety, etc.) obtains to a maximal degree, we get a theory of knowledge extensionally equivalent to infallibilism. As such, the infallibilist can affirm that, when their conditions are suitably interpreted, most post-Gettier theories of knowledge offer necessary and sufficient conditions for knowledge. The infallibilist can thus reconcile major theories of knowledge, and is in a better position to explain the intuitive appeal of these theories than the fallibilist who only accepts one of them, and rejects the rest.

In the *Meditations* (1641/1996), Descartes famously claims that knowledge (*cognitio*) consists in "clear and distinct perception" that precludes the possibility of error. In his *Essay* (1689/1996, IV, iii.14), Locke equates 'true knowledge' with 'certain knowledge', and says that in empirical matters this does not extend beyond what we directly observe:

> [T]hough we see the yellow color, and upon trial find the weight, malleableness, fusibility, and fixedness, that are united in a piece of gold; yet because no one of these *ideas* has an evident *dependence*, or necessary connection with the other, we cannot certainly know, that where any four of these are, the fifth will be there also, how highly probable soever it may be: because the highest probability, amounts not to certainty; without which there can be no true knowledge.

Several early twentieth-century philosophers express or assume similar views:

> The highest degree of rational belief, which is termed *certain* rational belief, corresponds to *knowledge*. We may be said to know a thing when we have a certain rational belief in it, and *vice versa*. (Keynes 1921, II.2)

> [I]t is quite certain that [(2)] other human beings also have known corresponding propositions: that is to say (2) also *is* true, and *I* know it to be true. (Moore 1925/1959a, 43)

*Dianoia Institute of Philosophy, Australian Catholic University, AU

> A statement is certain, i.e. is an expression of knowledge, only in one or other of two cases....
> (Ross 1930, ch. 2)

A common idea in these quotations is that knowledge and certainty are coextensive. This view is unpopular nowadays. My goal in this essay is to rehabilitate it. In section 1, I develop a theory of knowledge, INFALLIBILISM, based around this idea.[1] In section 2, I examine several familiar contemporary theories of knowledge. These theories all mention necessary conditions on knowledge that can be interpreted in weaker or stronger ways. I argue that if we take these otherwise independent conditions and interpret them in the strongest way possible, each theory ends up extensionally equivalent to INFALLIBILISM. In section 3, I argue that this gives us reason to favor INFALLIBILISM over fallibilist theories of knowledge, because INFALLIBILISM is better able to explain the intuitive appeal of various proposed conditions on knowledge. To borrow Parfit's (2011, 419) evocative metaphor, INFALLIBILISM implies that evidentialists, reliabilists, sensitivity theorists, and so on are all "climbing the same mountain on different sides."

1.   INFALLIBILISM

INFALLIBILISM comprises several theses. The first is that knowledge and certainty are coextensive:

(1) S knows that P if and only if P is certain for S.

The certainty here is *epistemic* certainty, as distinct from *psychological* certainty. These are technical terms, but they distinguish pretheoretic concepts. In ordinary language, I would use the locution "I am certain that P" to say that I am psychologically certain that P, and the locution "P is certain" to say that P is epistemically certain for me (Audi 2003, 224; Stanley 2008, 36–37). The former locution ascribes a *psychological* property to me (relative to P). The latter ascribes an *epistemic* property to P (relative to me).

What does it mean to say that a proposition is certain? While I do not have a reductive analysis of the concept of epistemic certainty to offer, a few comments may help clarify it and distinguish it from other concepts. First, unlike psychological certainty, epistemic certainty is factive: if P is certain for S, P is true. Second, epistemic certainty is related to epistemic probability in that, if P is certain for S, P has epistemic probability 1 for S. (The converse is not true, as I discuss in section 2.1.) Third, epistemic certainty is related to psychological certainty in that if P is epistemically certain for S, it is rational for S to be psychologically certain that P.[2] Finally, although the epistemic certainty of P for S is not dependent on S's doxastic states, it is dependent on S's mental states more generally. If it is certain for me that I am in pain, this is because of my current mental states. This proposition is not certain for other people with different mental states. Or, to take an a priori example, suppose I consider whether 5 is prime, and, holding the integers 2, 3, and 4 before my mind, see that none of these can be multiplied together to yield 5. Then that 5 is prime is certain for me, while it is not certain for someone who has not even considered the proposition.

In addition, it is plausible that when it is certain for me that P, this is because of some mental relation I stand in towards *P* in particular. Now, one might think that there need be no unique mental relation that plays this role. For example, it could be certain for me that I am in pain because of my introspecting that I am in pain, and certain for me that 5 is prime because of my intuiting that 5 is prime. But plausibly, these specific mental states of introspection and intuition are instances of a more general mental state. For consider the way certainty can be extended through deductive inference. If it is certain for me that I am in pain, and that 5 is prime, and I

competently deduce that (I am in pain)&(5 is prime), this conjunction becomes certain for me as well. But I neither introspect nor intuit that (I am in pain)&(5 is prime). So we need to posit some other mental relation that I stand in to this proposition to explain why it is certain for me.

What is this mental relation? According to (1), knowledge and certainty are coextensive. A natural suggestion is then that the mental relation that grounds certainty is knowing.[3] Introspection, intuition, and so on are kinds of knowledge; however, knowledge is not just the disjunction of these, but a primitive mental state itself. It is the mental state that makes the proposition known certain for the knower.[4] This is INFALLIBILISM's second thesis:

(2) When P is certain for S, this is because S stands in a primitive mental relation to P, namely, knowing P.

The conception of knowledge developed so far is close to the view that Antognazza (2015; 2020) defends, ascribing it to the historical Western tradition. According to Antognazza, prior to the mid-twentieth century most Western philosophers (including Descartes and Locke, mentioned above) did not accept the now common idea that knowledge is a kind of belief, and instead thought of knowledge and belief as fundamentally different kinds of cognitions:

> Knowledge is a primitive perception or an irreducible mental 'seeing' what is the case; knowledge is a primitive presence of a fact to the mind (or to the senses) in which there is no 'gap' between knower and known. Belief, on the contrary, is a mental state or a cognitive mode in which precisely the perception or presence which characterizes knowledge is lacking, and assent to the object of cognition is given (rightly or wrongly) on grounds external to the object itself. (Antognazza 2015, 169)

I similarly think of knowledge as a cognition different in kind from belief. However, contra Antognazza, I also think that one can know and believe the same proposition, and that while the common post-Gettier assumption that knowledge is a special kind of belief is mistaken, the assumption that belief always accompanies knowledge is correct. The basic idea is this. My knowledge that I am in pain consists in an immediate and pre-doxastic awareness that I am in pain. But this knowledge, in turn, compels me to believe that I am in pain.[5] This gives us INFALLIBILISM's final thesis:

(3) Whenever S knows that P, S's knowing P immediately compels S to believe that P.

This means that knowledge entails belief.[6]

Combining theses (1)–(3) above, we get:

INFALLIBILISM: Knowledge that P is a primitive mental state that grounds, and is coextensive with, P's being certain, and compels belief that P.

INFALLIBILISM is not an analysis of the concept of knowledge, in that it does not reduce this concept to more primitive concepts. Instead, INFALLIBILISM is a (nonreductive) account of what knowledge is, and how it relates to certainty, as understood above. In particular, INFALLIBILISM is a theory on which knowledge that P is a primitive mental relation with P, one that makes P certain and compels belief that P. For dialectical purposes, it will be helpful to have a neutral term for this mental state that does not presuppose that it is identical with knowledge. In what follows I shall call this mental state *apprehension*.[7] INFALLIBILISM is then the theory that knowledge is apprehension.

'Apprehension' here is a technical term: it picks out the mental relation an agent S has to a fact P that makes P certain for S. I do not take a stand here on precisely what facts we can apprehend.

This will largely depend on the outcomes of debates in the philosophy of mind and perception on which I remain neutral. A more conservative, internalist view might equate apprehension with what Russell (1910/1917) called "direct acquaintance," leading to the skeptical conclusion that the only facts we know are either a priori or describe our mental lives. On this view, I can apprehend that I am in pain because I am directly aware of myself, my pain, and the relation between them; I cannot apprehend that you are in pain because I am not directly aware of you or your pain. On the other hand, a direct realist about perception could claim that I can apprehend facts about the external world (such as that you are in pain).

Fallibilists, who hold that epistemic certainty is not necessary for knowledge, typically argue that infallibilism has unacceptably skeptical consequences for what we know (see, e.g., DeRose [1999], 202–203]; Hawthorne [2004, 126–31]). The proponent of INFALLIBILISM might try to meet this challenge by adopting a liberal, externalist view of apprehension, on which we can know a variety of facts about the external world through sensory perception, testimony, etc. But again, I don't wish to assume such a view here. If the internalist view of apprehension is right, then the skeptical consequences of INFALLIBILISM are a prima facie problem for the theory. But, as I argue in Climenhaga 2021, it is wrong to summarily dismiss infallibilist theories on account of their skeptical consequences. First, there are potential infallibilist explanations of our ordinary claims to knowledge that do not presuppose their truth: for example, "that we often wrongly attribute knowledge of propositions that are not really certain because we wrongly *think* that they are certain" and "that we attribute knowledge in cases where we recognize that it is not present because we are engaging in *loose talk*" (Climenhaga 2021, 73–74). In order to determine how bad the skeptical consequences of INFALLIBILISM are for the theory, we would have to evaluate the plausibility of these explanations.

Second, even if the skeptical consequences of INFALLIBILISM end up being a serious strike against the theory, theories still have to be evaluated relative to our total evidence. This is a truism about scientific theories, and it is no less true of philosophical theories. I do not claim here that INFALLIBILISM's reconciliation of major theories of knowledge outweighs the skeptical costs of the theory on its own. But it may do so in combination with other advantages of INFALLIBILISM, as part of a cumulative-case argument. I present such a cumulative case in Climenhaga (2021), arguing that infallibilism can explain more easily than fallibilism why knowledge seems to be "(1) qualitatively different from non-knowledge, (2) uniquely valuable, . . . (3) not possessed by subjects in Gettier cases . . . (4) evidence, (5) a basis for epistemic modals, (6) an inquiry stopper, (7) a basis for action, and (8) closed under competent deduction" (Climenhaga 2021, 73). This paper adds a ninth explanatory advantage to infallibilism, developing an infallibilist theory that can explain the intuitive appeal of the major necessary conditions on knowledge that have been proposed in the literature.[8]

2. HAVING IT ALL

The most influential theories of knowledge advanced in the past half-century give a different condition or conditions, in addition to belief and truth, for S's knowing that P. Here are several familiar conditions from the post-Gettier (1963) literature:[9]

S knows that P only if

  (a) P is true,
  (b) S believes that P,

and

(i)   P is sufficiently probable on S's evidence. (Evidentialism)
(ii) If P were false, S would not believe that P. (Sensitivity)
(iii) If S were to believe that P, P would not be false. (Safety)
(iv) S's belief is produced by a reliable cognitive process. (Reliabilism)
(v)   There are no true defeaters for S's justification for belief. (No Defeaters)
(vi) S's belief is caused by the fact that P. (Causal)

While some theories of knowledge combine more than one of (i)–(vi), these necessary conditions are usually seen as competitors, with each condition (modified as necessary to avoid counterexamples) holding pride of place as the central nonbelief, nontruth condition in rival theories of knowledge. Sometimes this pride of place is signified by presenting the condition in question as an analysis of 'justification'. One could accordingly view my argument as showing that different analyses of justification are equivalent when interpreted in an infallibilist-friendly way.

In this section, I will argue for two claims. First, INFALLIBILISM lets us accept all the above conditions as necessary for knowledge. Second, after independently motivated revisions to avoid familiar counterexamples to theories of knowledge based on these conditions, INFALLIBILISM lets us accept these conditions as sufficient for knowledge too. As such, INFALLIBILISM implies that, suitably revised, the above conditions are equivalent.

Conditions (i)–(iv) all mention properties—probability, sensitivity, safety, and reliability—that come in degrees. As such, a view on which one of these is necessary for knowledge faces a threshold problem—how much of that property is necessary (BonJour 2010)? For example, views on which S's knowing that P requires that P's probability on S's evidence be above some threshold must specify how probable is probable enough. Ordinary interpretations of these conditions make the required threshold nonmaximal—e.g., a probability of .95. Conditions (v) and (vi) do not mention degreed properties, but (as we will see in sections 2.3 and 2.4) they are also amenable to weaker or stronger interpretations, with fallibilists adopting weaker interpretations.

Under their typical nonmaximal or weaker interpretations, conditions (i)–(vi) can come apart from each other. For example, suppose the sensitivity condition is explicated as "in the nearest world in which P is false, S does not believe P" and the evidentialist condition as "the epistemic probability of P for S is above .95." Then S's (true) belief that a lottery ticket in a fair 100-ticket lottery will lose is sufficiently probable but not sufficiently sensitive if the nearest world in which that ticket wins is the mundane one in which the person picking out the winning ticket moves her hand a little to the left. Conversely, if S's paranormal precognition makes S's beliefs strongly correlate with what ticket will actually be picked, but S has strong (misleading) evidence that he possesses no such paranormal abilities, his (true) belief that his ticket will win is sensitive but not probable.

Matters change, however, if we adopt maximalist interpretations of the conditions. I shall now argue that—after independently motivated revisions to avoid familiar problems—conditions (i)–(vi) are equivalent when interpreted in the strongest way possible. For conditions (i)–(iv), this means requiring that the property mentioned in the condition obtain to a maximal degree. For condition (v), this means not restricting the range of defeaters that must be absent, and for condition (vi), it means requiring that the causation be direct. These conditions are then equivalent because each condition is met iff S apprehends that P.

2.1   Probability

First, consider the relation of condition (i) with certainty.[10] Condition (i) says that P is sufficiently probable on S's evidence.[11] To make this precise, we need to both say what S's evidence

includes and what the threshold for sufficient probability is. The strongest interpretation of this requirement is that the probability of P has to be 1 conditional only on those propositions that are certain for S. Here we only include in S's evidence those propositions that S is in a maximally good epistemic position to draw inferences from, and we require that the probability of P on those propositions be maximal.[12]

On this maximal interpretation, that P is certain for S implies that condition (i) is met—i.e., that P has probability 1 conditional on what is certain for S. For suppose I apprehend that P. Then P is part of my evidence, so my evidence (trivially) entails P, and the probability of P on my evidence is 1. So, certainty implies epistemic probability 1.

The converse is not true. Although logical consequences of facts known with certainty have probability 1 given those facts, these consequences are often not certain.[13] For example, it is not certain for most people that 7921 is composite (can be factored into whole numbers other than itself and 1). This is so even assuming that the belief condition (b) is met—if I believe that 7921 is composite, not on the basis of factoring it, but on the basis of wishing that this proposition were true (suppose I've been offered a prize if I correctly guess whether it's prime, and have guessed that it isn't), then this proposition is still not certain for me, even though it has probability 1 given my evidence. In order for this proposition to be certain for me, it is not enough that it be entailed by propositions I apprehend; I have to apprehend it specifically.

Intuitively, however, if P is made highly probable by my evidence but I believe it for unrelated bad reasons, I do not know that P. In order to know that P, my belief that P needs to be based on my evidence, in a way that respects or tracks P's being sufficiently probable on my evidence. Plausible versions of evidentialism thus add to (i) a *basing* requirement that precludes knowledge in cases like the above. Once we add this requirement, it is not enough to know P that I believe it and that it is entailed by my evidence. My belief must also be based on my evidence.

While the proper analysis of basing is controversial (see Korcz [2019] for discussion), it is plausible that once we add this basing requirement to the requirement that P have probability 1 on what I apprehend, P cannot fail to be certain. For example, suppose I consider the proposition that 6 is composite. I apprehend that it can be factored into 2 and 3, and on this basis conclude that 6 is composite. Here it is plausible that—if I really do apprehend that 6 can be factored into 2 and 3—I come to apprehend that it can factored into whole numbers other than itself and 1. More generally, if I directly and competently infer P from facts I apprehend, and these facts entail P, I will come to apprehend that P, and P will become certain for me.[14]

If we add the basing requirement to (i), is it still plausible that P's being certain for S implies that *both* these conditions are met? Yes. For we are assuming that apprehension that P compels belief that P. If this thesis is true, it is because apprehension that P causes belief that P directly, not because it causes it via some indirect causal chain. A plausible analysis of basing will imply that if S is compelled to believe that P by his apprehension that P in this direct way, S's belief that P is based on that apprehension. For example, causal accounts of basing standardly include the requirement that the causation of a belief by a reason be nondeviant. An easy way to make causation nondeviant is to make it direct, so that there is no room for irrelevant causal influences to enter in (McCain 2012, 371–74; cf. section 2.4 below).

2.2 Safety, sensitivity, and reliability

Second, I will argue that conditions (ii)–(iv) are identical when 'sensitive', 'safe', and 'reliable' are interpreted as maximal, and suitably reformulated to make them plausible as conditions for knowledge. A belief that P is sensitive just in case, were P not true, S would not believe P (Nozick 1981, 172–76; Sosa 1999, 141). Safety is the contrapositive of sensitivity: a belief that P is safe just in case, were S to believe P, P would be true (Sosa 1999, 146). Traditionally, the former condition is seen as satisfied if, in the nearest world(s) in which P is not true, S does not believe it; and the latter

is seen as satisfied if, in all nearby worlds in which S believes P, P is true. These can come apart if, e.g., S believes P in all worlds, and P is true in all nearby worlds but false everywhere else. Then S's belief is true in all nearby worlds, but false everywhere else; it is safe but not sensitive.

However, a belief that P is *maximally* sensitive just in case S does not believe P in any world in which P is not true. And it is maximally safe just in case P is true in every world in which S believes it (not merely close ones). These are equivalent.

Safety and sensitivity run into well-known problems as necessary conditions for knowledge. Suppose I believe I am in pain on the basis of introspection. If I am very suggestible and would have believed I am in pain had you told me so, and you would have told me so had I not actually been in pain (and I could have easily failed to be in pain), then my belief is not sensitive or safe. But it still seems that I know that I am in pain—I can *feel* that I am, after all! In light of this difficulty, epistemologists who propose safety and sensitivity as necessary conditions for knowledge have tended to require that S use the same "method" for determining whether P in all worlds under consideration (Nozick 1981, 179–85).

Incorporating belief-forming methods into these conditions brings them closer to (process) reliabilism, according to which S knows that P only if S's belief that P is produced by a cognitive process that produces a high proportion of true to false beliefs. To make this requirement more precise, we must answer two questions: how high a proportion, and in what contexts (Goldman 1979, 11)? If the cognitive process is to be *maximally* reliable, then the answer to the first question is, "100%," and the answer to the second question is, "all contexts." For example, suppose that S forms the true belief that P with a one-off cognitive process that he has never used before and never uses again, but this process would have issued in false beliefs in most counterfactual scenarios. Then S's cognitive process has a 100% success rate in the actual world, but is not counterfactually reliable. S's cognitive process is only maximally reliable in the relevant sense if it has a 100% success rate in all possible worlds in which it is used.

This gloss of reliabilism differs from safety/sensitivity in that it requires that the cognitive process leading to S's belief that P be *generally* reliable, and not just reliable when it issues in the belief that P (the proposition that S believes in the actual world). However, while our initial formulations of safety/sensitivity only considered cases where S believes that P, the safety/sensitivity theorist arguably ought to revise her account in this direction anyway. This is because of the problem of necessary truths (cf. Williamson 2000, 181–82). If I believe that 7919 is prime, my belief is trivially safe/sensitive, because 7919 is prime in all possible worlds. This is true, moreover, for whatever method I use. However, it does not seem that I know that 7919 is prime whenever I believe it. The most natural response to this difficulty is to look at the safety/sensitivity of beliefs in other propositions formed by the same method. (For example, if I believe that 7919 is prime because you told me, under what conditions in general do I believe your testimony, and how safe/sensitive are my other such beliefs?)

If we revise safety and sensitivity in this way, then we have the following maximal formulations of all these conditions on S's knowing that P:

> (Maximal Sensitivity) For all worlds W and propositions Q: if Q is not true in W, then S does not believe that Q in W via method M.
> (Maximal Safety) For all worlds W and propositions Q: if S believes that Q via method M in W, Q is true in W.
> (Maximal Reliability) For all worlds W and propositions Q: if S believes that Q via process P in W, Q is true in W.

So formulated, Maximal Sensitivity is logically equivalent to Maximal Safety. And the only difference between these and Maximal Reliability is that the former refer to belief-forming methods, and the latter to belief-forming processes.

Are these conditions coextensive with certainty? This depends on what qualifies as a belief-forming method or process. A method or process for forming beliefs issues in beliefs in response to some stimulus. Not just any stimulus will do, though. Goldman mentions "confused reasoning, wishful thinking, reliance on emotional attachment, mere hunch or guesswork, and hasty generalization" as unreliable cognitive processes, and "standard perceptual processes, remembering, good reasoning, and introspection" as reliable processes (Goldman 1979, 9–10). In canonical examples like these, the individuating stimuli for the belief-forming process (further beliefs, hopes, hunches, perceptual experiences, memories, etc.) are mentally accessible in some way.[15] Similarly, a method for forming beliefs needs to appeal to mentally accessible bases. We cannot follow a method like "believe the true propositions."

Different views in the philosophy of mind and perception will differ on what kinds of factors are mentally accessible in this way. In section 1, I developed the concept of apprehension in a way that remains neutral on these debates. So if, for example, externalism about the mind is true, and we can have direct access to external world facts, then we can apprehend those facts. As such, it appears that the factors that can be appealed to in a method/process are whatever we can apprehend, and what these include will depend on the outcome of internalist/externalist debates in the philosophy of mind and perception.

Whatever the outcome of these debates, apprehension that P entails that one's belief that P is formed by a maximally safe method. (Similar remarks go for belief-forming processes, here and below.) Since apprehension that P entails that P, if one believes that P *via* the method "believe those propositions the truth of which I apprehend," one cannot go wrong—one's belief will always be true. Hence, one's belief is maximally safe.

Are there any maximally safe methods besides apprehension-based belief? For a method to be maximally safe, there must be no possible worlds in which it outputs a false belief. How might this be the case? Either there needs to be some intrinsic connection between the basis of the method and the truth of its output—as there is in apprehension-based belief—or there needs to be some external factor—some kind of epistemic "guardian angel"—that guarantees that these never come apart.

As an example of the latter possibility, suppose that God exists in all possible worlds, and it is necessarily part of God's plan (though God has not told me this) that I be able to answer any question by tossing a coin, believing P when the coin lands heads and ~P when the coin lands tails. (Assume that I can apprehend that I have asked whether P, and that the coin has landed heads or that the coin has landed tails.) In every possible world in which I toss a coin and form a belief in the above manner, God ensures that the coin gives me the right answer. Then it is not metaphysically possible for this belief-forming method to output a false belief. So my coin-based belief that, say, the number of stars is even is formed by a maximally reliable process. And yet, I do not apprehend that the number of stars is even. This is true even on a liberal, externalist view of apprehension on which I can come to apprehend facts through perception and testimony: I do not perceive the number of stars, and have not been told that the number of stars is even.[16]

This kind of case is one in which, for deep metaphysical reasons, states of affairs that appear internally coherent nevertheless end up metaphysically impossible. Such cases are familiar from the philosophy of religion. For example, Leibniz (1710/1985) held that God created the best of all possible worlds, and his principles apparently commit him to the claim that it is necessarily true that God creates the best (Rowe 2004, ch. 1). It follows that our world could not have been even slightly different. But this is very counterintuitive. It seems clearly possible for there to have been one more star in the world, for example. There is nothing obviously incoherent about this state of affairs. Similarly, even if in the above scenario, it is in fact impossible that my coin lands heads and yet my belief that P is false, it *seems* possible in the sense that there is nothing internally inconsistent about this state of affairs.

Plausibly, in whatever sense ordinary possibility talk makes sense in a world with a necessarily existent God, in that same sense my beliefs in cases involving miraculous protection from false belief are not maximally safe. For example, perhaps we can make sense of ordinary possibility talk by looking at counterpossible but internally coherent worlds. Then on our ordinary modal conception of reliability, my coin-based beliefs will not be formed by a maximally reliable process, because they are false in some of the counterpossible but internally coherent worlds. By contrast, holding a false belief on the basis of apprehension is *analytically* impossible, given the functional definition of apprehension as that which makes a proposition certain for a subject. So in this sense, believing P iff my coin lands head is a less reliable method for forming beliefs than believing what I apprehend. And the same goes for any other method on which there is only an extrinsic connection between the basis of my belief and its truth.

Turning now to the former possibility: are there any other belief-forming methods, besides believing what one apprehends, on which there is a similar intrinsic connection between the basis of the belief and its truth? The only other kind of example I can think of is a method that tells you to believe those propositions that are *entailed* by propositions the truth of which you apprehend. However, in order to *follow* this method, you need to first see which propositions are entailed by what you apprehend. And if you apprehend that Q and that Q∨R follows from this, then you come to apprehend that Q∨R. So this method is not actually distinct from the method of believing what one apprehends. The only maximally reliable, safe, or sensitive method is then apprehension-based belief, and apprehension that P is both necessary and sufficient for Maximal Sensitivity, Maximal Safety, and Maximal Reliability.

2.3   No defeaters

So far I have considered proposed conditions (i)–(iv) on knowledge, all of which require that S's belief possess some property that comes in degrees. Our two remaining conditions, (v) and (vi), are not like this. The first, (v), requires that there be no true defeaters for S's justification for believing P. The no-defeaters condition is not usually taken to be sufficient for knowledge, and is more commonly combined with a justification condition to wield off counterexamples to the sufficiency of that condition, together with true belief, for knowledge.

A simple way to understand "S's justification for believing P" is as S's total evidence. If we label this evidence "K," then we can say that D is a defeater for that justification iff $P(P|D\&K) < P(P|K)$ (compare Lehrer and Paxson [1969, 227–28] and Shope [1983, 46]). According to this analysis, D is a defeater for S's justification for believing P just in case D lowers the probability of P relative to this evidence. We might also require that $P(P|D\&K)$ be less than .5, or some other threshold.

If P is certain for S, and facts one apprehends are part of one's evidence, then there is no other proposition S could learn (while continuing to apprehend that P) that would defeat S's evidence for P. For the probability of P on S's evidence will remain 1 so long as S apprehends that P, no matter what background K is also part of S's evidence. This is because $P(P|D\&P\&K) = 1$ for any defeater D and background K.

Is the no-defeaters condition sufficient for certainty? There is one trivial case in which the condition is met and P is not certain: when P has probability 0 on S's evidence. If S's evidence includes ~P, then $P(P|D\&{\sim}P\&K) = P(P|{\sim}P\&K) = 0$ for all D, and so nothing defeats S's justification for believing P.

This limiting case arises when S has, in a natural sense, *no* justification for believing P. That condition (v) is not sufficient for certainty here is not too interesting, for everyone will already agree that we cannot know propositions inconsistent with our evidence. If we add to (v) the very weak requirement that the probability of P on S's evidence is nonzero (so that S has some justification to potentially be defeated), then there is a case to be made that this condition does entail

that P is certain for S. For given the analysis of defeat as probability-lowering, if the epistemic probability of P for S is between 0 and 1, there will plausibly always be *some* defeaters for S's evidence. Simply take some false proposition Q such that $P(P|Q\&K) = P(P|K)$. Since Q is false, ~Q is true, and so ~Q∨~P is true. And necessarily, $P(P|[\sim Q \lor \sim P]\&K) = P(P|\sim[P\&Q]\&K) < P(P|K)$. Hence, ~Q∨~P is a defeater for S's evidence for P given the first analysis above. Hence, there is a defeater for any evidence for P provided there is some false proposition that is independent from P relative to that evidence. And there is plausibly always such a proposition. Just take an event causally isolated from P, such as the outcome of the next fair coin flip in China—if the coin lands heads, then the proposition that the coin lands tails is a false proposition independent from P.

What if we require that the defeater lower P's probability below .5—can we always find such a defeater then? For any true proposition R such that $P(P|K) < P(\sim R|P\&K)$, $P(P|[R \lor \sim P]\&K) <$ .5.[17] That is, R∨~P will be a true defeater that lowers P's probability below .5 provided that R is true and that R is more likely to be false, given P&K, than P is to be true, given K. It again seems plausible that there is always such a proposition (cf. BonJour 2010, 80 n.12). For example, suppose there is an infinite future, and let R be a complete description of the world from t_n to t_{n+m}. Then pick a distant enough starting point *n* that whether or not P is true gives us practically no information about what the world is like at that time—so that $P(\sim R|P\&K) \approx P(\sim R|K)$—and a sufficient amount of time *m* such that the probability of that complete description is closer to 0 than the probability of P is to 1—so that $P(P|K) < P(\sim R|K)$. If the future is infinite, we should always be able to find such *n* and *m*. If the future is not infinite, we can follow a structurally similar procedure provided we can find some other infinite domain to assign nonextreme probabilities to, such as an infinite past, an infinitely extended or infinitely divisible physical space, an infinite number of other universes, or an infinite God.

If probability 1 is not sufficient for a believed proposition to be certain, as argued in section 2.1 above, then my evidence for a believed proposition could be undefeatable without that proposition being certain. (For example, my evidence that 7919 is prime cannot be defeated on the above analysis of defeat.) However, if we add a basing requirement to (v), as we did to (i), then undefeatability is sufficient for certainty.

Partly because a broad conception of defeaters appeared to make (v) such a restrictive condition on knowledge, early advocates of (v) added ad hoc clauses restricting the range of defeaters that need to be avoided (see Shope [1983, chapter 2.1] for an overview). But proponents of INFALLIBILISM can endorse (v) in its simple, straightforward, and intuitive form: S knows that P only if (S has some justification for believing P and) there are no defeaters for S's justification for believing P. Moreover, I have argued that when we add a basing requirement to (v), then except when P has probability 0, INFALLIBILISM makes this condition sufficient for knowledge too.

2.4   The causal condition

This leaves us, finally, with (vi), the causal condition on knowledge, most famously suggested in Goldman's (1967) early post-Gettier paper. According to the simplest formulation of this condition, S knows that P only if S's belief that P is caused by the fact that P.

If S believes that P on the basis of apprehending that P, must S's belief be caused by the fact that P? There are two kinds of cases that might lead one to question this. The first is knowledge of a priori domains like mathematics. One might think that abstract objects cannot be causally efficacious, so that my apprehension that $1 + 1 = 2$ cannot be caused by the fact that $1 + 1 = 2$.

I am not persuaded by this suggestion. While abstract objects cannot directly causally influence nonmental concrete things, it seems to me that rational intuition acquaints us with abstracta in a way that lets those abstracta causally influence our state of mind. Even though the abstracta do not change, our attention does; and when the abstract and mental causes combine

in the right way, we get a mental effect we would not have otherwise. However, even if I am wrong about this, I take it that everyone should agree that if I apprehend that $1 + 1 = 2$, the fact that $1 + 1 = 2$ still *explains* my belief (together with other factors, such as my considering the proposition). Since presumably the causal theorist should allow for mathematical knowledge, we could then revise condition (vi) to say that S's belief that P is *explained* by the fact that P.

The second case is knowledge by deduction. Suppose I apprehend that $1 + 1 = 2$. This leads me to infer, and apprehend, that either $(1 + 1 = 2) \lor (4$ is prime$)$. One might think that in this case, my inference is sufficient to explain my new knowledge, so that the fact that $(1 + 1 = 2) \lor (4$ is prime$)$ does not help cause my knowledge (and so belief) that $(1 + 1 = 2) \lor (4$ is prime$)$.

This conclusion is not inevitable, however. We can think of this case as one in which my apprehension that $1 + 1 = 2$ puts me in a position to apprehend something else I was not previously in a position to apprehend, namely that $(1 + 1 = 2) \lor (4$ is prime$)$. But my apprehension of this latter fact is still caused by this fact in the same way my apprehension of the former fact is caused by that fact. To speak metaphorically, a light has been shone on the fact that $(1 + 1 = 2) \lor (4$ is prime$)$, so that I can now see something that was previously shrouded in darkness. The inference does not directly—without the fact that $(1 + 1 = 2) \lor (4$ is prime$)$ playing any role—lead me to know $(1 + 1 = 2) \lor (4$ is prime$)$; rather, what it does is put the fact that $(1 + 1 = 2) \lor (4$ is prime$)$ before my mind in a way that it was not before, allowing that fact to imprint itself upon my mind.

So it is plausible that if S believes that P on the basis of apprehending that P, S's belief is caused (or explained) by the fact that P. Not all belief caused by facts is based on apprehension, however. For example, S could have a belief that he has fallen down the stairs as an indirect result of having fallen down the stairs and hit his head, causing his brain to get scrambled and giving him all sorts of random beliefs, including that he has fallen down the stairs (Feldman 2003, 83).

Even proponents of a causal account of knowing would not want to count S as knowing he has fallen down the stairs in this case. Hence, in order to make the causal condition sufficient for knowledge, causal theorists need to revise their account to avoid deviant causal chains such as this one, as Goldman (1967, 358–59) recognized when he first proposed the theory. Such revisions face the problem of specifying what kind of causation is "appropriate."

One infallibilist-friendly way of spelling out appropriate causation is *direct* causation. Now, we don't want to say that the fact that P must directly cause belief that P, because there might (and plausibly always will) be intervening mental states, such as perceptual experience. But we could hold that the mental state causing belief must be directly caused by P.

Let's look at belief based on apprehension as our model. Here our causal chain looks like this:

Fig 1.

There will be other causal factors, together with P, that lead S to apprehend that P—e.g., S considering P, or S drawing an inference from facts that entail P. But there is a nondeviant chain from P to S's belief here because P is (together with these other factors) a direct cause of a mental state that entails belief that P. So we could spell out the appropriateness requirement as saying that P appropriately causes S's belief that P only if P directly causes a mental state that entails that S believes that P.

Different views in the philosophy of mind and perception will differ on when a fact P directly causes a mental state, and, in particular, whether this ever happens when P is a fact about the external world. In section 1, I developed the concept of apprehension in a way that remains neutral on these debates. So if, for example, direct realism about perception is true, and the glass on the table is a direct cause of my perception, then I will apprehend that there is a glass on the table.[18] As such, on this explication of appropriate causation, whenever P appropriately causes S's belief that P, S's belief that P will be based on apprehension that P. So, on this understanding of (vi), this condition will be met when, and only when, S apprehends that P.

3.   CLIMBING THE EPISTEMIC MOUNTAIN

In section 2, I argued that if P is certain for S, each of conditions (i)–(vi) is met. I further argued that if any of conditions (i)–(vi) is (fully) met, P is certain for S. For conditions (i)–(iv), once we revise them in light of counterexamples, they all end up coextensive with certainty (and so coextensive with each other) if the properties they mention obtain to a maximal degree. Counterexamples to both the necessity (my belief that I am in pain) and sufficiency (my belief that 7919 is prime) of conditions (ii) and (iii) moved us to reformulate these conditions in the direction of condition (iv), and when maximally interpreted, these three conditions are then coextensive with certainty. Counterexamples to the sufficiency of (i) for knowledge motivated us to add a basing requirement to it; and together with this basing requirement, this condition is also coextensive with certainty. As for condition (v), except in the edge case in which P has probability 0, this is sufficient for probability 1 (condition [i]) on its simplest interpretation, and so together with a basing requirement sufficient for certainty. And condition (vi) is sufficient for certainty if we require that the causation be suitably direct.

So, once we suitably reformulate conditions (i)–(vi), and interpret them in an infallibilist-friendly way, we find that S's belief that P meets any one of these conditions when, and only when, P is certain for S. If knowledge and certainty are coextensive, then each of conditions (i)–(vi) is, together with (a) and (b), necessary and sufficient for knowledge. INFALLIBILISM thus implies that versions of most standard post-Gettier accounts of knowledge are extensionally correct.

If, like me, you find each of conditions (i)–(vi) appealing, you will see this result as an advantage of INFALLIBILISM. But perhaps you are not like me. Perhaps you are a convinced reliabilist, and think that evidentialism is a nonstarter. Even so, you may still see the appeal of *some* of the other conditions above (say, the safety, sensitivity, and causal conditions)—and in this case, you should regard INFALLIBILISM's reconciliation of these conditions as a point in its favor.

Moreover, whatever your own intuitions about conditions (i)–(vi), each of them has been endorsed and defended by many philosophers. Infallibilists can explain this sociological fact more easily than fallibilists who accept one of these conditions and reject the rest. Many philosophers have been attracted to theories centered on each of these conditions because, suitably interpreted, every one of these theories is true.[19]

In my experience receiving feedback on this argument from various audiences (including referees at this and other journals), fallibilists tend to object to it in three ways. The first objection is that it may be possible to formulate a version of fallibilism that also reconciles multiple proposed conditions on knowledge. The second objection is that INFALLIBILISM cannot really explain the intuitive appeal of widely proposed theories of knowledge, because the version of those theories INFALLIBILISM ends up vindicating is too different from the fallibilist versions that have actually been proposed. The third objection is that even if INFALLIBILISM enjoys an advantage over fallibilism in rendering versions of standard theories of knowledge extensionally correct, this is swamped by its inability to render true so many of our ordinary claims to knowledge. INFALLIBILISM does not let us "have it all" because it lets us have hardly any actual

knowledge, and surely—the objection goes—this is more significant than its compatibility with popularly proposed conditions on knowledge.

I have already addressed the last objection in section 1, and will not dwell on it here except to say that my use of the phrase "have it all" is largely rhetorical. I do not think that INFALLIBILISM is a superior theory to fallibilist theories of knowledge on every possible point of comparison. For the sake of argument, I grant the critic of INFALLIBILISM that INFALLIBILISM implies that most of our ordinary claims to know things are false, so that fallibilists are indeed able to explain these claims more easily than proponents of INFALLIBILISM. The proponent of INFALLIBILISM is able to "have it all" not in vindicating our ordinary claims about what we know but in accommodating the key ideas of major philosophical theories of what knowledge is. How much relative weight to give these two considerations is too big a question to answer here.

Whether or not the advantages of INFALLIBILISM end up outweighing its disadvantages, the central contention of the present paper is simply that INFALLIBILISM's capacity to reconcile different theories of knowledge is a significant advantage of the theory. The second objection above challenges this claim. According to this objection, epistemologists have endorsed the theories of knowledge considered here only because of their first-order implications for what we know—and once we reformulate the theories in an infallibilist direction, so that they imply that we know very little, they lose their philosophical appeal. Hence INFALLIBILISM cannot explain the appeal of the versions of conditions (i)–(vi) that have actually been defended in the literature.[20]

It's true that many familiar arguments for the above theories of knowledge rely on their implying that we know particular uncertain propositions (although some others rely on the theories' implying that we don't know particular uncertain propositions, like propositions in Gettier cases). But other arguments for and against proposed theories of knowledge focus on the theoretical plausibility of their central conditions—e.g., the extent to which those conditions can explain how knowledge is valuable, or guides action, or lets us identify reliable informants. And discussions of rival conditions often characterize them as attempts to explicate how knowledge is "connected to the truth" in a way that mere true belief is not (e.g., Alston 1985; Ichikawa and Steup 2018; Pritchard et al. 2018). The most popular conditions are popular partly because they are seen as plausible ways of explicating this truth-connection.

I take it, then, that most epistemologists hold that their preferred condition on knowledge is *theoretically appealing* in addition to delivering the right (fallibilist) verdicts about concrete cases. In accepting maximal versions of these conditions, INFALLIBILISM inherits this theoretical appeal, even if it has different consequences for what we know than the fallibilist versions of these theories. According to INFALLIBILISM, these conditions—these different ways of being connected to the truth—are different sides of the same mountain. Knowledge is the peak of this mountain, where all these conditions meet.

Put otherwise, fallibilists have an easier time than proponents of INFALLIBILISM explaining why it is *fallibilist versions* of conditions (i)–(vi) that have been popular—versions that imply that we know lots of things. But fallibilists can't as easily explain why all of (i)–(vi) are theoretically appealing in the first place. And the former explanatory advantage of fallibilism is just its antiskeptical consequences noted above, which is already widely recognized as an advantage of fallibilism. What I have argued here is that (a version of) infallibilism has a hitherto-unrecognized advantage, namely explaining the theoretical appeal of different popular conceptions of justification more easily than versions of fallibilism that only accept one of these conceptions and reject the rest.

This brings us to the first objection above, which is that it may be possible to develop a fallibilist theory that reconciles the above conditions. The basic argument in section 2 was that conditions (i)–(vi) converge when (i)–(iv) are taken to their maximum, and (v) and (vi) (which don't have a similarly straightforward "maximum") are interpreted in an infallibilist-friendly

way. Otherwise, these conditions come apart. For conditions (i)–(iv), this is because they require that some quantity (probability, sensitivity, safety, reliability) be present to a sufficient degree, and at nonextreme points, these quantities can be varied independently of each other—e.g., we can increase probability without increasing reliability, or decrease safety without decreasing sensitivity.

For example, if we set the threshold for sufficient probability at .5, this will not be equivalent to a reliability threshold of 50%. Suppose that S sees that the urn in front of him contains one black ball and one white ball. He closes his eyes, sticks his hand in and draws a ball out. Let P be the proposition that S draws a black ball. If S randomly believes P or ~P, then it's plausible that both probability and reliability here are at 50%. But by changing the details of the case we can vary either of these quantities without affecting the other. For example, let S also see in the urn what looks like a red ball but is really just a reflection of the light on the ceiling. But leave S's psychological tendencies completely unchanged by this fact. Then the reliability of S's belief-forming process remains at 50%, while the probability of P drops below 50%. Or add in the detail that there is a magnet hidden inside the white ball that will control S's brain if S draws it and compel him to believe P. Then reliability drops to 25% while probability remains at 50%.

I lack the space to explicitly discuss all other pairs of conditions, but I think inspection shows that they diverge in similar ways for nonextreme thresholds—as assumed in standard fallibilist epistemology, in which one condition is defended to the exclusion of the others. The one potential exception is that I noted in section 2.2 that there are principled reasons to revise conditions (ii) and (iii) to talk about the sensitivity/safety of beliefs in other propositions formed by the same methods as the present belief that P. The safety condition will then end up equivalent to the reliability condition if we evaluate safety and reliability relative to the same beliefs in the same contexts. (The sensitivity condition will remain distinct from the reliability condition, because for nonmaximal sensitivity and reliability, a belief-forming method/process can only be made more sensitive by resulting in fewer false beliefs, while it can also be made more reliable by resulting in more true beliefs.)

With that said, there is one other *extreme* threshold for justification with convergence properties worth investigating, namely the minimal threshold. Consider Sartwell's (1991, 1992) theory that S knows that P iff (a) P is true and (b) S believes P. Sartwell observes that the theory has an advantage similar to that enjoyed by INFALLIBILISM:

> [O]n [this theory], we are under no apparent pressure to *choose* between a broadly externalist and a broadly internalist account of justification. The pressure to choose between these views arises largely because proponents of each argue that their account is an account of the sense of justification that is *logically required for knowledge*. But justification is in *no* sense logically required for knowledge. (Sartwell 1991, 163)

Sartwell says that his theory implies that no kind of justification is required for knowledge. But this is not quite right. For conditions (a) and (b) already imply *minimal* versions of conditions (i), (iii), and (iv). P is not epistemically impossible (because it is true), S's belief is safe at least to the extent that it is true in this world, and S's belief-forming process is not maximally unreliable (in that it issues in a true belief in this world). We could maybe even salvage condition (v) if we say that D defeats S's justification for believing P in the relevant sense only if D entails ~P. But if D entails ~P, then P entails ~D, so if P is true there is no true defeater D for S's justification for believing P. Then a weak reading of the no-defeaters condition is satisfied.

Sartwell's theory can then unite conditions (i), (iii), (iv), and (v), in that given conditions (a) and (b), the minimal thresholds for probability, safety, reliability, and indefeasibility are equivalent (because already met). Of course, stronger fallibilist theories of knowledge also imply that

minimal versions of these conditions are met, but because these theories tether themselves to some particular conception of epistemic justification, it's harder for them to explain why (say) probability only needs to be present to a minimal degree while reliability needs to be present to a greater degree. The mere true belief theory of knowledge unites these conceptions of justification by holding that they need to be present to the same degree, namely a minimal degree.

However, conditions (a) and (b) do not imply minimal versions of conditions (ii) or (vi). S truly believes P in the actual world, but S may also believe P in all possible worlds where P is false, in which case S's belief is not at all sensitive. And S's belief that P may be completely causally independent of P, so that (vi) is not satisfied even on a very weak construal. So the true belief theory of knowledge is not wholly neutral between different conceptions of justification.

This discussion suggests that the best a fallibilist can do is set a *minimal* threshold for justification that unites all but two of conditions (i)–(vi): namely sensitivity and the causal condition. And this feat is only achievable by the true belief theory of knowledge, which most fallibilists will consider much too weak. More standard fallibilist theories will be unable to accommodate at least four extensionally distinct conditions. Only the infallibilist can have it all.

4. CONCLUSION

In this essay I have argued that INFALLIBILISM—the view that knowledge that P is a primitive mental state that grounds, and is coextensive with, P's being certain, and compels belief that P—allows us to reconcile major post-Gettier theories of knowledge, in that, if we interpret key conditions in these theories in maximalist ways, these theories end up extensionally equivalent to INFALLIBILISM. I have further argued that this reconciliation is reason to accept INFALLIBILISM: for the proponent of INFALLIBILISM is in a better position to explain the intuitive appeal of these various theories of knowledge than the fallibilist who only accepts one of them, and rejects the rest. Infallibilists can have it all, but fallibilists cannot.

NOTES

1. The term 'infallibilism' has been used to refer to a number of different theses (see Brown 2018, ch. 1 for an overview). I discuss the relationship between these theses and my own theory in Climenhaga (ms.).
2. INFALLIBILISM thus entails Unger's (1975, ch. 3.2) thesis that if S knows that P, it is permissible for S to be psychologically certain that P, and Dodd's (2011, 665) thesis (which he also calls 'infallibilism') that knowledge requires epistemic probability 1.
3. Other infallibilist theories of the explanatory relationship between knowledge and certainty are possible. For example, one might accept that knowledge and certainty are coextensive but think that certainty grounds knowledge, or that certainty and knowledge have a common ground. However, both of these possibilities require the existence of three distinct coextensive things: certainty, knowledge, and the mental state that grounds certainty. If we say that knowledge grounds certainty, then we only have to posit two coextensive things: certainty and knowledge. So this is arguably the simplest way to develop an infallibilist theory of knowledge.
4. Williamson (2000) is the best-known contemporary defender of the view that knowledge is a mental state. Williamson does not call himself an infallibilist, however, and although he thinks that knowledge implies probability 1, denies that it implies "absolute certainty" (213). My own understanding of the relevant mental state is closer to Antognazza's (2015, 2020), as noted below. For reasons of space I focus here on contrasting my theory with traditional fallibilist theories of knowledge. I discuss Williamson's and Antognazza's theories (and the extent to which they can reconcile different proposed conditions on knowledge) further in Climenhaga (ms.).
5. Antognazza's "traditional" account of knowledge contains a similar idea, which is that knowledge and belief both involve *assent*, but in knowledge (unlike belief) this assent is compelled by the object of one's knowledge. The differences between these ideas are partly verbal, but not wholly so. See Climenhaga (ms.) for further discussion.

6. Again, other infallibilist views about the relation of knowledge and belief are possible. The challenge facing a form of infallibilism that denies (3) is that the conditions discussed in section 2 imply that knowledge entails belief, and so no version of these conditions can be extensionally correct if knowledge does not entail belief. So an infallibilist theory that denies that knowledge entails belief will have to explain the plausibility of these conditions in a way that does not require their being extensionally correct. See Climenhaga (ms.) for further discussion.

7. On similar uses of this term in early 20th-century philosophy, see Antognazza (2020, 297–300).

8. See Climenhaga (ms.) for a more comprehensive examination of the evidences for and against INFALLIBILISM and assessment of their relative weights.

9. For an overview of some of the most common accounts of knowledge, see Shope (1983), Feldman (2003), and Ichikawa and Steup (2018). I reference advocates of the particular conditions discussed here below.

10. For a theory of knowledge in which (i) holds pride of place, see McGrew and McGrew (2007).

11. Some evidentialists hold that evidential support cannot be measured by probability. But they still tend to agree that evidential support comes in different degrees. For example, Conee and Feldman (2008, 98) write: "One's justification for a proposition can be of various strengths. Correspondingly, one's evidence varies in the strength of its support of different propositions. According to our explanatory coherence view of evidential support, this variation in strength of support derives from differences in how well the supported propositions explanatorily cohere with one's evidence." While I will assume that evidential support is measured by probabilities, even if it is not, my argument will still go through provided that evidential support comes in varying degrees, and there is a maximal degree of evidential support. This is because what I say about probability 1 will still be plausible for maximal evidential support.

12. Some evidentialists hold that evidence is nonpropositional. However, I take it that such views will get the same substantive results about how probable my evidence makes particular propositions—e.g., if my evidence consists of my apprehensions themselves, rather than their propositional objects, then the presence of an apprehension will ensure the truth of its object, and so the apprehension of P will "entail" P in a looser sense of 'entail'.

13. Apparent counterexamples to the sufficiency of probability 1 for certainty also arise in cases involving infinite sample spaces. Suppose I throw an infinitely fine dart at a square dartboard the sides of which range from 0 to 1. The probability that I will hit point (0.2, 0.4) is 0. So, the probability that I will not hit point (0.2, 0.4) is 1. And yet, it is not certain for me that I will not hit this point. Arguably, however, this kind of case just shows that epistemic probabilities come in finer degrees than the real numbers, and that we should allow for probabilities to take on nonreal values. We can then assign a probability infinitesimally lower than 1 to the proposition that I will not hit point (0.2, 0.4).

14. There are two related objections one could make to this principle. The first is that competent deduction does not require infallibility, so that in drawing a deductive inference one takes on some degree of risk that one's conclusion is false, even if one's premises are certain (cf. Lasonen-Aarnio 2008). The second is that you might be uncertain about whether you have inferred competently, and so be uncertain about whether the conclusion of your inference is true (cf. Schechter 2013). The second objection fails because apprehension need not be luminous; and what determines whether P is certain for you is whether you apprehend P, not whether you apprehend that you apprehend that P. As for the first objection, while I agree that competent deduction does not imply that one never makes mistakes in *other* deductions, I do not agree that it is compatible with uncertainty in *this* deduction. Long chains of deductions do not preserve certainty because one cannot hold all the steps in one's mind at once, so that by the time one infers the ultimate conclusion, one no longer apprehends the initial premises and so no longer apprehends the intermediate conclusion either. But competent direct inferences where one apprehends the premises all at once are certainty-preserving.

15. One might worry that reliabilism is an externalist theory, and that epistemic externalists think the factors that justify belief are not all mentally accessible. However, reliabilism is standardly understood to be an externalist theory because it allows for the *reliability* of the cognitive process to be inaccessible (see, e.g., Goldman and Beddor [2016, sec. 2]). This is compatible with requiring that the stimuli in response to which a belief-forming process operates are accessible.

16. This case is similar to common counterexamples to the sufficiency of safety for knowledge based on miraculous protection from false belief in similar situations (e.g., Goldman 2009, 76–77). The main difference is that the present case involves protection from false belief in *all* possible situations, not merely similar ones.

17. Proof:

Suppose: $P(P|K) < P(\sim R|P\&K)$

Then: $\frac{P(P\,|\,[R\vee\sim P]\,\&K)}{P(\sim P\,|\,[R\vee\sim P]\,\&K)} = \frac{P(P\,|\,K)}{P(\sim P\,|\,K)} \times \frac{P([R\vee\sim P]\,|\,P\,\&K)}{P([R\vee\sim P]\,|\,\sim P\,\&K)} = \frac{P(P\,|\,K)}{P(\sim P\,|\,K)} \times \frac{P(R\,|\,P\&K)}{1}$

$\qquad < \frac{P(P|K)}{P(\sim P|K)} \times \frac{P(\sim P|K)}{1} = P(P|K) < 1$

So: $\frac{P\,(P\,|\,[R\vee\sim P]\,\&K)}{P(\sim P\,|\,[R\vee\sim P]\,\&K)} < 1$

So: $P\,(P\,|\,[R\vee\sim P]\,\&\,K) < P(\sim P\,|\,[R\vee\sim P]\,\&\,K)$

So: $P\,(P\,|\,[R\vee\sim P]\,\&\,K) < .5$

18. That the glass on the table is a direct cause of my perception does not imply that there is no indirect causal chain from the glass on the table to my state of mind. It only implies that there is not *only* an indirect causal chain from the glass to my mind. A plausible version of direct realism will say something like the following: since my state of mind in apprehending that there is a glass on the table is partly constituted by the glass itself, there is a direct connection between the glass and my mind in addition to the indirect connection mediated by the light traveling to my eyes, etc. Here the causal network looks like this:
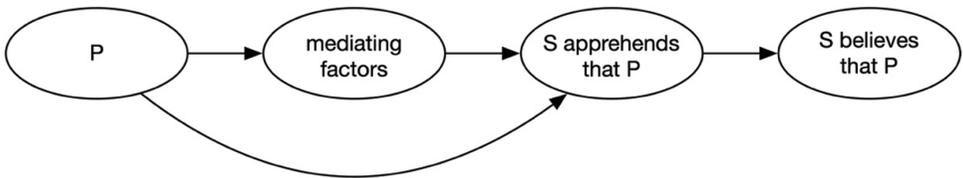
Fig 2.

19. In addition, while my focus here has been on post-Gettier theories of knowledge, I noted in the introduction that Descartes, Locke, Keynes, Moore, and Ross all saw knowledge and certainty as coextensive. And Dutant (2015) argues that most philosophers prior to the mid-20th century were infallibilists of some sort, while Antognazza (2015, 2020) ascribes to the historical tradition a view of knowledge very similar to that defended here. While a detailed examination of historical views of knowledge is beyond the scope of this essay, this suggests that INFALLIBILISM can reconcile contemporary theories of knowledge not only with each other, but also with historical theories of knowledge. Fallibilists, by contrast, face the burden of explaining why infallibilist theories were so popular historically, if in fact knowledge does not require certainty. (For some relevant discussion along these lines, see Pasnau [2017].)

20. Echeverri (2023, 63–65) advances a similar objection to the argument of Climenhaga (2021).

REFERENCES

Alston, William 1985. "Concepts of Epistemic Justification," *The Monist* 68: 57–89.

Antognazza, Maria Rosa 2015. "The Benefit to Philosophy of the Study of its History," *British Journal for the History of Philosophy* 23: 161–84.

———. 2020. "XII—The Distinction in Kind Between Knowledge and Belief," *Proceedings of the Aristotelian Society* 120: 277–308.

Audi, Robert 2003. *Epistemology: A Contemporary Introduction to the Theory of Knowledge*, Second Edition, Routledge.

BonJour, Laurence 2010. "The Myth of Knowledge," *Philosophical Perspectives* 24: 57–83.

Brown, Jessica 2018. *Fallibilism: Evidence and Knowledge*, Oxford: Oxford University Press.

Climenhaga, Nevin 2021. "A Cumulative Case Argument for Infallibilism," in *Skeptical Invariantism Reconsidered*, ed. Christos Kyriacou and Kevin Wallbridge, Routledge.

———. ms. *Knowledge and Certainty: A Defense of Infallibilist Epistemology*.

Conee, Earl and Richard Feldman 2008. "Evidence," in *Epistemology: New Essays*, ed. Quentin Smith, Oxford: Oxford University Press.

DeRose, Keith 1999. "Contextualism: An Explanation and Defense," in *The Blackwell Guide to Epistemology*, J. Greco and E. Sosa, eds., Malden, MA: Blackwell, 185–203.

Descartes, René 1641/1996. *Meditations on First Philosophy*, trans. John Cottingham, Cambridge: Cambridge University Press.

Dodd, Dylan 2011. "Against Fallibilism," *Australasian Journal of Philosophy* 89: 665–85.

Dutant, Julien 2015. "The Legend of the Justified True Belief Analysis," *Philosophical Perspectives* 29: 95–145.

Echeverri, Santiago 2023. "Critical Notice: Christos Kyriacou and Kevin Wallbridge's *Skeptical Invariantism Reconsidered*," *International Journal for the Study of Skepticism* 13: 61–78.

Feldman, Richard 2003. *Epistemology*, New York: Prentice Hall.

Gettier, Edmund L. 1963. "Is Justified True Belief Knowledge?" *Analysis* 23: 121–23.

Goldman, Alvin 1967. "A Causal Theory of Knowing," *Journal of Philosophy* 64: 357–72.

———. 1979. "What is Justified Belief?" in *Justification and Knowledge*, ed. G.S. Pappas, Dordrecht: D. Reidel, 1–23.

———. 2009. "Williamson on Knowledge and Evidence," in *Williamson on Knowledge*, ed. Patrick Greenough and Duncan Pritchard, Oxford: Oxford University Press, 73–91.

Goldman, Alvin and Bob Beddor 2016. "Reliabilist Epistemology," *The Stanford Encyclopedia of Philosophy*, ed. E.N. Zalta. https://plato.stanford.edu/archives/win2016/entries/reliabilism/.

Hawthorne, John 2004. *Knowledge and Lotteries*, Oxford: Oxford University Press.

Ichikawa, Jonathan Jenkins and Matthias Steup 2018. "The Analysis of Knowledge," *The Stanford Encyclopedia of Philosophy*, ed. E.N. Zalta. https://plato.stanford.edu/archives/sum2018/entries/knowledge-analysis/.

Keynes, John Maynard 1921. *A Treatise on Probability*, London: Macmillan and Co.

Korcz, Keith Allen 2019. "The Epistemic Basing Relation," *The Stanford Encyclopedia of Philosophy*, ed. E.N. Zalta. https://plato.stanford.edu/archives/fall2019/entries/basing-epistemic/.

Lasonen-Aarnio, Maria 2008. "Single Premise Deduction and Risk," *Philosophical Studies* 141: 157–73.

Lehrer, Keith and Thomas Paxson 1969. "Knowledge: Undefeated Justified True Belief," *The Journal of Philosophy* 66: 225–37.

Leibniz, Gottfried Wilhelm 1710/1985. *Theodicy*, transl. E.M. Huggard, ed. Austin Farrar, London: Routledge and Kegan Paul.

Locke, John 1689/1996. *An Essay Concerning Human Understanding*, ed. Kenneth P. Winkler, Indianapolis, IN: Hackett.

McCain, Kevin 2012. "The Interventionist Account of Causation and the Basing Relation," *Philosophical Studies* 159: 357–82.

McGrew, Timothy and Lydia McGrew 2007. *Internalism and Epistemology: The Architecture of Reason*, London: Routledge.

Moore, G.E. 1925/1959a. "A Defence of Common Sense," in *Contemporary British Philosophy* (2nd series), ed. J.H. Muirhead; reprinted in Moore's *Philosophical Papers*, Collier Books, 32–59.

Nozick, Robert 1981. *Philosophical Explanations*, Cambridge, MA: Harvard University Press.

Parfit, Derek 2011. *On What Matters: Volume 1*, Oxford: Clarendon Press.

Pasnau, Robert 2017. *After Certainty: A History of Our Epistemic Ideals and Illusions*, Oxford: Oxford University Press.

Pritchard, Duncan, John Turri, and J. Adam Carter 2018. "The Value of Knowledge," *The Stanford Encyclopedia of Philosophy*, ed. E.N. Zalta. https://plato.stanford.edu/archives/spr2018/entries/knowledge-value/.

Ross, William David 1930. *The Right and the Good*, Oxford: Oxford University Press.

Rowe, William L. 2004. *Can God Be Free?* Oxford: Oxford University Press.

Russell, Bertrand 1910/1917. "Knowledge by Description and Knowledge by Acquaintance," in *Mysticism and Logic*, New York: Barnes and Noble, 152–67.

Sartwell, Crispin 1991. "Knowledge Is Merely True Belief," *American Philosophical Quarterly* 28(2): 157–65.

———. 1992. "Why Knowledge Is Merely True Belief," *Journal of Philosophy* 89: 167–80.

Schechter, Joshua 2013. "Rational Self-Doubt and the Failure of Closure," *Philosophical Studies* 163: 429–52.

Shope, Robert 1983. *The Analysis of Knowing*, Princeton, NJ: Princeton University Press.

Sosa, Ernest 1999. "How to Defeat Opposition to Moore," *Philosophical Perspectives* 13: 141–53.

Stanley, Jason 2008. "Knowledge and Certainty," *Philosophical Issues* 18: 35–37.

Unger, Peter 1975. *Ignorance: A Case for Skepticism*, Oxford: Oxford University Press.

Williamson, Timothy 2000. *Knowledge and Its Limits*, Oxford: Oxford University Press.