4   A Cumulative Case Argument for Infallibilism

Nevin Climenhaga

This essay is a defense of the following thesis about knowledge:

Infallibilism   S knows that P only if P is certain for S.[1]

The certainty here is *epistemic* certainty, as distinct from *psychological* certainty. In ordinary language, we 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), while the latter ascribes an epistemic property to P (relative to me). The two are connected, though: if P is epistemically certain for S, then S *ought* to be psychologically certain that P.[2]

In the next eight sections, I present eight intuitive claims about knowledge. Each claim is plausible given infallibilism, but in tension with fallibilism. This essay is "big-picture," and I do not try to demonstrate that, in each case, there is *no* way for the fallibilist to accept the intuitive claim in question. Rather, I aim to show that each of these claims is easier to reconcile with infallibilism than fallibilism, and that consequently, each intuition is a *prima facie* consideration in favor of infallibilism.[3] While individually, these considerations might not be conclusive, together they provide strong support for infallibilism.

Explanatory arguments like that advanced here are fundamentally comparative: my claim is that infallibilism explains these intuitions *better than* rival fallibilist theories. For the sake of space, I am unable to discuss all rival theories here. In particular, I limit my targets to fallibilist theories on which knowledge does not require epistemic probability 1. Williamson's "knowledge-first epistemology," on which knowledge implies epistemic probability 1 but not certainty (see Williamson 2000: chapter 10.2),[4] deserves separate discussion. I suspect that Williamson's theory can account for some of the intuitive data I go on to adduce and not others, but the issue is too complicated to get into here. Instead, I only argue here that infallibilism is superior to theories on which knowledge is compatible with epistemic probability less than 1; and when I refer to 'fallibilism' in what follows, this is the view I have in mind.

I do, however, devote special attention to two non-knowledge-first theories specially tailored to account for some of the intuitive data I present: contextualism and interest-relativism. In presenting my eight claims in Sections 1–8, I initially assume that if infallibilism is false, traditional (non-contextualist, non-interest-relativist) fallibilism is true. Then, in Section 9, I consider what resources contextualists and interest-relativists have for explaining the intuitions I put forward.

A complete defense of infallibilism must respond to the main objection contemporary epistemologists level against the view: that it has unduly skeptical consequences that conflict with our ordinary intuitions and ascriptions of knowledge (see, e.g., DeRose 1999: 202–203; Hawthorne 2004: 126–131). After all, there seems to be very little we should be *maximally confident* in. Even if it's highly probable that I have hands, it seems epistemically possible that I am a handless brain in a vat. One way to respond to this objection is to deny that infallibilism has such restrictive consequences about the scope of our knowledge (see, e.g., Dutant 2016). For the sake of argument, however, I assume in this essay that the critics are right that if infallibilism is true, we know very little. In particular, I assume that the only facts certain for us are introspectable facts about our mental lives and *a priori* facts that we can know immediately through intuition.[5] I argue in Section 10 that, while the skeptical consequences of infallibilism are then a cost of the theory, this cost is outweighed by infallibilism's capacity to explain such a large variety of intuitive data about the nature and theoretical roles of knowledge.

On to the eight claims. The claims I will present in Sections 1–3 are as follows:

1   There is a qualitative difference between knowledge and non-knowledge.
2   Knowledge is valuable in a way non-knowledge is not.
3   Subjects in Gettier cases do not have knowledge.

The claims in Sections 4–8 have to do with *theoretical roles* that knowledge apparently plays—the ways in which it appears to be related to rational inference, epistemic modality, rational action, rational inquiry, and deduction:

4   If S knows that P, P is part of S's evidence.
5   If S knows that P, ~P is epistemically impossible for S.
6   If S knows that P, S can rationally act as if P.
7   If S knows that P, S can rationally stop inquiring whether P.
8   If S knows each of $\{P_1, P_2, \ldots P_n\}$, and competently deduces Q from these propositions, S knows that Q.

I will argue that each of these claims is plausible if infallibilism is true, but inconsistent with or in strong tension with fallibilism.

1 Knowledge Is Qualitatively Different from Non-knowledge

The most influential theories of knowledge advanced in the past century defend some condition or conditions as, together with belief and truth, necessary and sufficient for knowledge. Often these conditions are presented as analyses or varieties of "epistemic justification." Here are several proposed necessary conditions on S's knowing that P[6]:

  i S's belief that P is sufficiently probable on S's evidence.
  ii If P were false, S would not believe that P.
  iii If S were to believe that P, P would not be false.
  iv S's belief is produced by a reliable cognitive process.
  v S can rule out all relevant alternatives to P.

Most of these conditions contain threshold terms: "sufficiently probable," "reliable," "relevant alternatives," etc. The others, while not containing threshold terms, are still open to more or less strict interpretations: for example, one could interpret sensitivity as requiring that in *all* worlds in which P is false, S does not believe P, or as simply requiring this for some particular worlds.

It appears that there is a qualitative difference between knowing a proposition and having a belief that is highly justified, but not highly enough for knowledge. Infallibilism explains this, because if infallibilism is true, then I know P only when P is certain for me. In this case I can rule out *all* alternatives to P; and my belief is maximally probable, formed by a maximally reliable process, etc. (Climenhaga forthcoming). This is a qualitatively different situation from having a belief that is highly, but not maximally, justified in one of the above ways. So, infallibilism makes it plausible that

1 There is a qualitative difference between knowledge and non-knowledge.

By contrast, fallibilist theories based around conditions that come in degrees have difficulty explaining how knowledge could be qualitatively different from almost-but-not-quite knowledge. For the move from 0.89 to 0.9 probability, or truth in the five nearest possible worlds to truth in the six nearest worlds, and so on, is a quantitative one, and not a qualitative one.

2 Knowledge Is Uniquely Valuable

BonJour (2010) observes that the apparent qualitative difference between knowledge and non-knowledge accompanies an apparent qualitative

difference in *value*. Knowledge appears to be valuable in a way that non-knowledge is not. Infallibilism explains this, because it is uniquely valuable to have epistemic certainty, in that this completely protects one's belief from error (cf. Brown 2018: 7). Hence, infallibilism makes it plausible that

2　Knowledge is valuable in a way non-knowledge is not.

By contrast, fallibilism has a hard time explaining (2). Most of the necessary conditions for knowledge proposed by epistemologists are valuable in some way. So a theory of knowledge based on them can arguably explain why knowledge is valuable. However, the conditions mentioned above involve possession of some quantity that comes in degrees. And slight increases in this quantity, so long as it remains sub-maximal, constitute only quantitative differences in value. As BonJour puts it, finding "conclusive justification" for a proposition

> would be the best situation of all. But the claim of the weak [fallibilist] conception is that there is some specific level of justification that is less than conclusive but nonetheless transforms your cognitive situation in a much more radical way than did increases in justification up to that point (or further increases above it). Before this level is attained, you merely have a belief that is more and more likely or probable, but at that point you suddenly have knowledge. But why does achieving this specific level of justification make such a difference and what exactly is this difference supposed to amount to? … It is hard to see why such further increases are not valuable in exactly the same way, to precisely the same extent, as those that came earlier, before the supposed "magic" level was reached.
> (BonJour 2010: 61)

Hence, fallibilism cannot easily explain why the value of knowledge is *qualitatively* different from the value of non-knowledge.

One explanation for why fallible knowledge is qualitatively more valuable than near-knowledge is that knowing has extrinsic *benefits* that almost-knowing does not. If you pass your class only if you score 80% on the final exam, scoring 80% is better than scoring 79% in a way that scoring 81% is not better than scoring 80%, even though intrinsically both are mere quantitative increases. Hannon (2014) responds to BonJour's argument by appeal to benefits of this sort. According to Hannon, we see S as knowing that P at the point when we think that S can legitimately close inquiry as to whether P: "This level of justification is cognitively valuable because it satisfies one of the platitudes about the functional role of knowledge ascription: it signals the point of legitimate inquiry closure" (Hannon 2014: 1132; cf. Kappel 2010: 191–192).

This response, however, is only plausible if coupled with interest-relativism. For as I argue in Section 7, if knowledge is not interest-relative, then we can only *always* close inquiry on whether P upon coming to know P if knowledge implies probability 1. I make similar arguments for other roles that knowledge plays, including rationalizing action and serving as evidence. If these arguments are right, then the only way for a non-interest-relativist to appeal to these roles of knowledge to explain its unique value is to adopt infallibilism. If traditional fallibilism is true, then the arbitrariness of whatever the threshold for knowledge is means that it does *not* have the extrinsic benefits, such as allowing us to rationally act or close inquiry, required by this response. So, the traditional fallibilist remains unable to easily explain the unique value of knowledge.

3 Gettiered Subjects Do Not Know

BonJour mentions two other problems for fallibilism: the Gettier problem and the lottery paradox (BonJour 2010: 63–70; see also Reed 2012: 588–591). I will discuss the former in this section; I discuss the latter later, under the heading of closure.

'The Gettier problem' is a term of art, but contemporary epistemologists think of it primarily as a problem of 'lucky' knowledge (e.g., Zagzebski 1994; Howard-Snyder et al. 2003, Ichikawa and Steup 2014: Section 3): you give me a theory of the conditions under which S knows that P, and I'll give you a case in which those conditions are met but S's belief is only luckily true, and so intuitively not knowledge. But infallibilism completely eliminates the relevant kind of luck (Kirkham 1984; Kyriacou 2017: 29–30). If it is certain for you that P, it is not at all lucky that your belief turns out to be true. Indeed, in his original essay, Gettier (1963: 121) is explicit that his counterexamples to the justified true belief theory of knowledge rely on the assumption that it is possible for a person to be justified in believing a false proposition.

Infallibilism thus entails

3   Subjects in Gettier cases do not have knowledge.

The last 50 years of epistemology, by contrast, show that it is difficult for fallibilists to give a theory which other epistemologists are not able to come up with "Gettier-style" counterexamples to. Fallibilists have a harder time than infallibilists allowing for (3).

4 Knowledge Is Evidence

The term 'evidence' can be used in two related ways. First, we can say that P is evidence *for* Q: call this the "evidence-for" sense of 'evidence.' P is evidence for Q relative to some background iff P raises the probability of

Q relative to that background. The second sense is the "having-evidence" sense: we can say that S has P as evidence, or that P is part of S's evidence. If S has P as evidence, then P, together with the other propositions in S's evidence, determines the epistemic probability of any proposition *for* S.

We are concerned here with the "having-evidence" sense of 'evidence.' Here is a simple argument that knowledge implies probability 1:

4   If S knows that P, P is part of S's evidence.[7]
9   The epistemic probability of P for S is *n* iff *n* is the probability of P conditional on S's evidence.
10 For any evidence K apart from P, P(P|P&K) = 1.
11 If S knows that P, the epistemic probability of P for S is 1. [from (4), (9), (10)]

(10) is analytically true for standard probability axioms. Thus, if the fallibilist wishes to accept (4), he must deny (9). While I lack the space to consider all the alternatives to (9) that might be proposed, I will make two points in defense of this premise. First, (9) is the simplest, most natural way to spell out the relation between probability and evidence. Other theories which seek to avoid the consequence that one's evidence has probability 1 will generally be *ad hoc*. Second, the simplest theories that avoid this consequence will render obviously incorrect verdicts.[8] For example, one might say that the epistemic probability of P for S is the probability of P conditional on all of S's *other* evidence (besides P). This will wrongly deliver the result that if S has no strong evidence for P besides P itself, and P is antecedently unlikely, the epistemic probability of P for S is very low. For example, suppose that I suddenly feel a shooting pain in my foot. Relative to everything else I know at this moment, it is quite unlikely that I feel pain in my foot. Nevertheless, the epistemic probability that I feel pain in my foot is obviously not low for me.

These remarks suffice to show that it is at least difficult for fallibilists to deny (9), and so difficult to accept (4) but deny (11). Now, fallibilists who deny (4) may be inclined to doubt that this claim is intuitive. But when we look at philosophers, scientists, and statisticians writing about probability in other contexts, the language they use suggests that they tacitly presuppose that knowledge is evidence. For example, Bayesians use the phrase "background knowledge" to refer to that part of an agent's evidence which was not just now learned, or is not under explicit consideration right now (e.g., Eells and Fitelson 2000: 667–669). Here are some more specific examples (emphases mine):

> *Prob*$_f$*(h)* is to be a number representing the person's personal probability for *h*, when he *knows f*; for short, his probability given *f*.
> (Hacking 1967: 313)

[T]o speak of the probability of an event tout court, without any qualification, does not have any concrete meaning. Rather, it must be kept in mind that probability is always relative to the *state of knowledge* of the person who is making the judgement.

(de Finetti 1979/2008: 36)

When you ask yourself how much support *e* gives *h*, you are plausibly asking how much a *knowledge* of *e* would increase the credibility of *h*, which is the same thing as asking how much *e* boosts the credibility of *h* relative to what else you currently know.

(Howson 1991: 54)

In principle perhaps, non-demonstrative inference should be based on 'total evidence'… In practice, however, investigators must think about *which bits of what they know* really bear on their question.

(Lipton 2001: 113)

If P(*A* | *B*, *C*) = P(*A* | *C*), we say that *A* and *B* are conditionally independent given *C*; that is, *once we know C*, learning *B* would not change our belief in *A*.

(Pearl 2000: 3)

The [ideally reasoning] robot always takes into account *all of the evidence it has* relevant to a question. It does not arbitrarily ignore some of the information, basing its conclusions only on what remains. … The robot always represents *equivalent states of knowledge* by equivalent plausibility assignments. That is, if in two problems the robot's state of knowledge is the same…, then it must assign the same plausibilities in both.

(Jaynes 2003: 19)

These quotes suggest that, in contexts in which philosophical issues about the nature of knowledge are not at stake, both philosophers (Hacking, de Finetti, Howson, Lipton) and scientists (Jaynes, Pearl) take for granted that if S knows that P, P is part of S's evidence. To the extent that philosophers are inclined to deny that knowledge is evidence, then, it is most likely for theoretical reasons. In particular, inasmuch as philosophers recognize the truth of (9) and (10), they may realize that knowledge being evidence would imply that we know only what has probability 1 for us, and so reject this view because they reject this consequence of it.

The infallibilist is able to explain the intuitive appeal of (4) more easily than the fallibilist. For the infallibilist can accept that (4) is true, and so hold that our ordinary thought and language which presupposes (4) reflects a tacit recognition of (4). By contrast, the traditional fallibilist can only accept (4) by denying (9), and (9) is very plausible. The fallibilist could deny (4)

and offer some error theory for why (4) is presupposed in so much ordinary thought and language. But it is not obvious what shape such an error theory would take, and I know of no fallibilist attempts to formulate one.[9]

5 Not-P Is Possible for S Only If S Does Not Know That P

Dodd (2011) presents the following argument for the claim that if S knows that P, the epistemic probability of P for S is 1:

  5   If S knows that P, ~P is epistemically impossible for S.
12   ~P is epistemically impossible for S only if the epistemic probability of ~P for S is 0.
13   If the epistemic probability of ~P for S is 0, the epistemic probability of P for S is 1.
11   If S knows that P, the epistemic probability of P for S is 1. [from (5), (12), (13)]

(13), like (10), is analytic. Fallibilists looking to reject (11) must deny one of the other two premises. (12) is an extremely plausible claim about epistemic possibility. As Dodd (2011: 668) observes, the analog of (12) is obviously true for physical possibility and probability: P is physically impossible only if there is zero physical probability that it occurs. Dodd thinks, and I agree, that the most plausible fallibilist response to the argument is to deny (5). However, although it is not as obvious as (12), (5) is pretheoretically very plausible. In the last section, I noted that in contexts in which philosophers and scientists are not preoccupied with the nature of knowledge—and, consequently, not worried about the threat of skepticism—they tend to presuppose that knowledge is evidence. Dodd (2011: 669) similarly observes that "virtually all" writers on epistemic modals accept (5). When the threat of skepticism is not salient, the intuitive theoretical roles of knowledge come to the fore, and philosophers freely use it to analyze other concepts, such as epistemic possibility.

   (5) is further supported by the impropriety of *concessive knowledge attributions* (CKAs), that is, sentences of the form

     "I know that P, but maybe ~P."

CKAs are typically infelicitous to assert. For example, it is infelicitous to say "I know the Red Sox won't win, but they might." (5) gives a simple explanation of why CKAs are infelicitous: they are always false. If you know that P, then ~P is epistemically impossible for you.[10]

   The infallibilist can easily accept (5). The fallibilist can only do so by denying (12). But (12) is extremely plausible. So, the fallibilist again has a harder time accepting an intuitive datum about knowledge than the infallibilist.

6 Knowledge Lets Us Rationally Act

Here is another argument for (11):

6   If S knows that P, S can rationally act as if P.
14 If knowledge does not imply epistemic probability 1, then it is possible that S knows that P but cannot rationally act as if P.
11 If S knows that P, the epistemic probability of P for S is 1. [from (6), (14)]

(6) should be read as a necessary claim. Thus, if (14) is true, its antecedent is inconsistent with (6), and (11) follows.

(6) is plausible because there is an apparent conceptual connection between knowledge and rational action, just as there is an apparent conceptual connection between knowledge and evidence. Principles similar to (6) are endorsed by many epistemologists, often in the context of arguing for interest-relativism about knowledge. For example, Hawthorne (2004: 30) writes that "it is acceptable to use the premise that *p* in one's [practical] deliberations if one knows it and (at least in very many cases) unacceptable to use the premise that *p* in one's practical reasoning if one doesn't know it."

Similarly, Stanley (2005: 10) says that

> A standard use of knowledge attributions is to justify action. … To say that an action is only based on a belief is to criticize that action for not living up to an expected norm; to say that an action is based on knowledge is to declare that the action has met the expected norm.[11]

Interest-relativists are in a position to deny (14). This is because sometimes P is uncertain for S, but it is rational for S to act as if P because S's evidence that P makes P sufficiently probable relative to the costs of S's being wrong—and according to interest-relativism, this is just part of what it is to know that P. But the traditional (non-interest-relativist, non-contextualist) fallibilist is *not* in a position to deny (14). So (6), while usually presented as evidence for interest-relativism, is equally good evidence for infallibilism.

For example, suppose that S has a choice between φ-ing and doing nothing. If S does nothing, nothing happens. If S φ-s and P is true, S gains \$1. If S φ-s and P is false, then S loses some amount of money. S knows all this (with certainty). In this case φ-ing constitutes acting as if P: this is the course of action that has higher utility if P is true, and S knows this (with certainty).

Now suppose that S's belief that P has epistemic probability 0.9, that this exceeds any non-maximal probabilistic threshold for knowledge,

and that any other necessary conditions for fallibilist knowledge are met. Then S (fallibilistically) knows that P. However, suppose that we make the loss S will incur if P is false $10. Then the expected utility of φ-ing is (0.9)($1) + (0.1)(–$10) = $0.90 – $1 = –$0.10. Since the expected utility of not φ-ing is $0, S ought not φ.

If the threshold for knowledge is above 0.9, then we can simply increase S's loss if P is false. The only way to ensure that acting as if P maximizes expected utility is to give P an epistemic probability of 1. Hence, (14) is true. Thus, the infallibilist, but not the (traditional) fallibilist, can accept (6).

7 Knowledge Lets Us Rationally Close Inquiry

We saw that one fallibilist response to BonJour's argument that only infallibilism can explain the value of knowledge is that knowledge is valuable because it allows us to close inquiry. The fact that knowledge lets us close inquiry can itself be used to support infallibilism, *via* an argument structurally identical to the one in the last section:

7   If S knows that P, S can rationally stop inquiring whether P.
15   If knowledge does not imply epistemic probability 1, it is possible that S knows that P but cannot rationally stop inquiring whether P.
11   If S knows that P, the epistemic probability of P for S is 1. [from (7), (15)]

(7) is plausible because there is an apparent conceptual connection between knowledge and inquiry closure. If you come to know that P, then this settles the question of whether P for you. There is no more need for you to inquire whether P (cf. Kyriacou 2017: 30–31; Kappel 2010).

Together with (15), however, (7) implies that what we know has epistemic probability 1. And the traditional fallibilist is not in a position to deny (15). For suppose that S's belief that P has epistemic probability 0.9, that this exceeds the threshold for knowledge, and that any other necessary conditions for fallibilist knowledge are met. Then, S (fallibilistically) knows that P. However, if the costs of being wrong about P are high enough, then S ought not close inquiry with respect to P; instead, she ought to keep inquiring whether P. By continually increasing the stakes, we can run the same argument for any constant threshold that falls short of probability 1. Hence, (15) is true. Thus, the infallibilist, but not the (traditional) fallibilist, can accept (7).

8 Knowledge Is Closed under Competent Deduction

Epistemologists have traditionally found plausible various *closure* principles about knowledge. Closure principles capture the intuitive idea that

we can extend our knowledge by deduction. In this section, I will argue that we can only accept what is sometimes called multi-premise closure if knowledge requires probability 1:

8 If S knows each of $\{P_1, P_2, \ldots P_n\}$, and competently deduces Q from these propositions, S knows that Q.

16 If knowledge does not imply epistemic probability 1, then it is possible that S knows each of $\{P_1, P_2, \ldots P_n\}$, and competently deduces Q from these propositions, but does not know that Q.

11 If S knows that P, the epistemic probability of P for S is 1. [from (8), (16)]

The proper formulation of closure principles is controversial (see Hawthorne 2005). For example, perhaps (8) fails in cases where S loses knowledge of the premises while performing the deduction. However, my argument that fallibilists cannot accept (8) also shows that they cannot accept principles that add conditions to the antecedent to avoid these kinds of counterexamples. For if we amended (8) to avoid these kinds of counterexamples, we could similarly amend (16), and it would remain plausible for the reasons adduced below. So I will stick with (8) here for simplicity's sake. All my argument requires is that a multi-premise closure principle along these lines is plausible, even if the antecedent needs some chisholming.

One common argument for (16) is from lottery cases (e.g., Hawthorne 2004: 6–7, 182; BonJour 2010: 66–70; Reed 2012: 588–590). However, these arguments require the controversial premise that, if fallibilism is true, one can know of a losing lottery ticket that it is a loser. Some fallibilists deny this, holding that, e.g., we cannot know on the basis of purely statistical evidence (Nelkin 2000; for criticism, see Christensen 2004: 62–64; BonJour 2010: 68–69). I thus prefer to rely on a different argument for (16), based on the preface case.

Suppose you write a meticulously researched book consisting of 1000 claims: $C_1, C_2, \ldots, C_{999}$, and C_{1000}. Each of these claims has a high but non-maximal epistemic probability for you. Suppose that the probability of each claim is 0.99, and that conditional on your evidence, each claim is probabilistically independent of every other claim or conjunction of claims.[12] This means that the probability of any conjunction of them is equal to the product of the probability of the conjuncts—e.g., $P(C_1 \& C_2 | K) = 0.99^2 = 0.9801$ (where K is your evidence). It follows from this that the probability that all the claims in your book are true is equal to $0.99^{1000} \approx 0.00004$. As such, you acknowledge in your preface that, in all probability, there are some false claims in this book, and say that they are solely the result of your own error.

Since in this case the claims in your book can be about anything at all, we can assume that any conditions on knowledge of the individual

C_i are met: they are all true, your beliefs in them are reliably formed, your evidence for them is not purely statistical, and so on. It follows that you ought to be able to deduce, and come to know, that all the claims in your book are true. But clearly, you cannot know this. This follows immediately if we assume that there is any probabilistic requirement on knowledge. For example, if knowledge requires epistemic probability above 0.5, you cannot know that all the claims in your book are true. If the probabilistic threshold is lower than this, then we can always add more claims to the book to make the probability of their conjunction lower than that threshold. Likewise, if we raise the probabilistic threshold above 0.99 so as to preclude knowledge of the individual claims of the book, we can simply increase your evidence for them so that their individual probability now exceeds that threshold. The only way to avoid this result is to set the threshold to 1—i.e., to accept (11).

One might hold that there is no probabilistic requirement on knowledge. In this case, one can coherently accept that you are in a position to know that all the claims in your book are true. But this is still an absurd result. It is an absurd result even if one rejects mathematical probability as an adequate formalization of plausibility.[13] It just doesn't seem that you can know a proposition like "all the claims in my book are true," when this proposition is one you should be extremely confident is not true.[14]

Infallibilists can allow that knowledge can always be extended by deduction. But, I have argued, fallibilists cannot. Like the theoretical claims about knowledge considered in the previous several sections, the claim that knowledge is closed under competent deduction is not indubitable. If the costs of rejecting fallibilism are high enough, then we should deny it. But, just like these other theoretical claims about knowledge, multi-premise closure is very plausible. That a theory of knowledge entails that it is false is a cost of that theory.

9 Contextualism and Interest-Relativism

In Sections 1–8 of this chapter, I presented eight intuitive claims, arguing that each is easier to reconcile with infallibilism than traditional fallibilism. In this section I will consider how well contextualism and interest-relativism can explain these intuitions.

Contextualists about knowledge hold that 'knows' picks out different epistemic states in different contexts, so that the truth-conditions of 'S knows that P' depend on the context in which that sentence is uttered. Contextualists will typically hold that in contexts in which skeptical hypotheses are taken more seriously, or speakers think that getting it right is very important, 'knows' will pick out a harder-to-reach epistemic state than in more "ordinary" contexts.

Interest-relativists hold that whether S knows that P depends partly on S's practical interests vis-à-vis P: how much is at stake for S with regard

to being right about P? If it is very important for S to be right about P—if the costs of being wrong are high—then interest-relativists will typically hold that it is harder for S to know that P.

Interest-relativists can offer elegant resolutions to several of the problems facing traditional fallibilists. For example, they can say that S's belief meets the probabilistic threshold for knowledge just in case the expected utility of acting as if P is greater than the expected utility of not acting as if P. Although this threshold will vary depending on an agent's utilities, it is not arbitrary—and interest-relativists can thus explain why (1) knowledge is qualitatively different from non-knowledge. In connecting rational belief to rational action in this way, interest-relativism can explain why (6) knowledge that P lets us act on the assumption that P. In addition, if we can close inquiry about whether P when P is probable enough on our evidence for us to act on P, interest-relativism can also explain why (7) knowledge that P lets us close inquiry on whether P. Finally, a belief that one can rationally act on and close inquiry regarding is valuable in a way that other belief is not; hence, interest-relativism can also explain (2) the unique value of knowledge.

Contextualists can offer an analogous explanation of (1) that appeals to the speaker's context.[15] It is less obvious that this will explain our other intuitive data, but perhaps the contextualist could argue that the threshold for knowledge-attributions in a context varies with the threshold for what the speaker *takes* to be high enough (e.g.) probability to act upon/stop inquiring, and that in the speaker's context belief that meets this threshold is conceived as valuable in a way that belief that fails to meet it is not. For the sake of argument, let us grant that the contextualist can in this way explain the unique value of knowledge and its conceptual connection to rational action and inquiry.

We are granting, then, that interest-relativists and contextualists can explain (1), (2), (6), and (7). This leaves four other intuitive data: (3) subjects in Gettier cases lack knowledge, (4) knowledge is evidence, (5) knowledge that P makes ~P epistemically impossible, and (8) knowledge is closed under competent deduction.

Interest-relativists and contextualists continue to face the Gettier problem. The same counterexamples that plague traditional fallibilist theories plague interest-relativist and contextualist theories, when we fix the speaker's context and the knower's practical stakes. So, interest-relativists and contextualists have a harder time than infallibilists explaining (3).

Interest-relativists and contextualists cannot allow for multi-premise closure. For interest-relativists, this is because the property of being probable enough to rationally act upon is not closed under conjunction. For example, perhaps C_1 (the first claim of your book) is probable enough for you to take it for granted in your deliberations. Even so, the conjunction $C_1 \& \ldots \& C_{1000}$ is *not* probable enough for you to take it for

granted in your deliberations. Suppose that your book is about to be scrutinized by a committee who will infallibly determine whether it contains errors. If it does *not* contain errors, you will be given a $1,000,000 cash prize. Even if each individual claim in the book is probable enough that you can act on the assumption that it is true, you cannot rationally act on the assumption that you will win $1,000,000 by, e.g., quitting your job and buying a ticket for a cruise around the world. However, if (8) were true, then, if you competently deduced that your book contains no errors, you would know that your book contains no errors, and so, according to (6), you *could* rationally quit your job and buy a ticket for a cruise around the world. Inasmuch as the interest-relativist is committed to (6), he must thus reject (8).

The contextualist would like to save a metalinguistic version of closure, namely:

17 If in a context c, "S knows that P_1," "S knows that P_2," …, "S knows that P_n" all state true propositions, and S competently deduces Q from $\{P_1, P_2, …, P_n\}$, then "S knows that Q" would state a true proposition in c.

Suppose we just have two premises, P_1 and P_2, and let $Q = P_1\&P_2$. If the contextualist sets a certain threshold of probability for knowledge-attributions in a given context, that P_1 and P_2 meet that threshold does not imply that $P_1\&P_2$ meets that threshold. For example, if the speaker's threshold for knowledge-attributions is 0.8 probability, then P_1 and P_2 may each individually be at least this probable while their conjunction is less probable. Thus, their conjunction is not probable enough for the proposition that would be expressed by "S knows that $P_1\&P_2$" to be true.

The contextualist might try to rescue multi-premise closure by claiming that mention of a conjunction changes the context in such a way that the speaker can no longer truly say "S knows that P_i" for the individual propositions P_i. If we are talking about your book, I might truly say, "You know that C_1," "You know that C_2," etc., but then when you ask me "And do I know $C_1\&…\&C_{1000}$?" you have shifted the context to a high-standards one in which I can no longer truly say "You know that C_1," "You know that C_2," etc.—nor, for that matter, "You know that $C_1\&…\&C_{1000}$."

Setting aside whether this claim about the change of contexts is plausible or not, it does not save multi-premise closure. (17) is about closure *within* a context, not closure *across* contexts. Even if your mentioning the conjunction of the claims in your book shifts the context, it is still the case that in the *initial* context, "You know that C_i" would express a truth for each C_i, and that, if (17) is true, then "You know that $C_1\&…\&C_{1000}$" would express a truth in that context. But it is implausible that "You

know that $C_1\&...\&C_{1000}$" would express a truth in that context, even granting that it is a "low-standards" context. If standards are codified as probability, then this follows immediately given a sufficient number of claims and independence assumptions; the C_i could individually all be above the probability-threshold and the conjunction be below it. And even if standards are not codified in this way, it just is not plausible that there is *any* sense of 'know' (even one in which "You know that you are not a brain in a vat" expresses a truth!) such that the proposition expressed by "You know that $C_1\&...\&C_{1000}$" is true, when you should be extremely confident that $C_1\&...\&C_{1000}$ is false.

We can see from the above conclusion that interest-relativists and contextualists cannot accept that knowledge is evidence either. For it follows from the claim that knowledge is evidence that knowledge has epistemic probability 1 (given (9), which says that the epistemic probability of P for S is *n* iff *n* is the probability of P conditional on S's evidence). And if knowledge has epistemic probability 1, then the conjunction of everything one knows has epistemic probability 1. Hence, it ought to be knowable as well. But, since we just saw that an interest-relativist cannot allow that knowledge is closed under competent deduction, it follows that they cannot allow that the conjunction of everything one knows has epistemic probability 1.

We need to be careful here. Interest-relativists will endorse (6), according to which S can know that P only if it is rational for S to act as if P. If we construe "acting as if P" as acting the way that would maximize expected utility conditional on P being true, it follows that, if the pragmatic condition on knowledge is met, one can "pretend" that P is part of one's evidence and it will not make a difference for rational action— what it is rational for one to do will be the same whether one assigns P its actual probability on one's evidence or one conditionalizes on it, assigning it probability 1. One can, as it were, hypothetically add P to one's evidence for the purposes of action. It might initially seem, then, that the interest-relativist *can* endorse the rational permissibility (at least practically speaking) of assigning *everything* one knows probability 1, and hypothetically adding everything one knows to one's evidence.

This argument tacitly assumes a false closure principle of the following form: if adding P to one's evidence does not change what actions it is rational to perform, and adding Q to one's evidence does not change what actions it is rational to perform, then adding P&Q to one's evidence does not change what actions it is rational to perform. Where E is your actual evidence, the first claim in your book, C_1, may be probable enough for you that conditionalizing on $C_1\&E$ rationalizes the same actions as conditionalizing on E. For example, if you are considering a bet on C_1 that would have positive expected utility either way, you should take it either way, and if you are considering whether to quit your job and buy that cruise around the world, you should not do it either way,

because you almost certainly will not win the \$1,000,000 prize even conditional on C_1&E, and so that action has negative expected utility either way. The same may be true for each other C_i. But, conditionalizing on C_1&…&C_{1000}&E does not rationalize the same actions as just conditionalizing on E, for relative to C_1&…&C_{1000}&E, quitting your job and cruising around the world has (we can suppose) positive expected utility, whereas conditional just on E it still has negative expected utility.

So interest-relativists cannot accept that knowledge is evidence, because interest-relativist conditions on knowledge are not closed. Contextualists cannot accept a metalinguistic analogue of the claim that knowledge is evidence for similar reasons. It follows from the claim that an utterance of "S knows that P" in a context expresses a truth only if an utterance of "P is part of S's evidence" expresses a truth in that context that an utterance of "S knows that P" expresses a truth in that context only if an utterance of "P has epistemic probability 1 for S" also expresses a truth in that context. Contextualists must deny that the latter generally holds, because otherwise the truth of knowledge-ascriptions would (wrongly) be closed under competent deduction by the knower. Thus, they must deny that the former holds as well.

If knowledge (or the truth of "S knows that P") does not imply epistemic probability 1, then, as argued in Section 5 above, it is not the case that we are in a position to accept that S knows that P (or to truly say, "S knows that P") only if ~P is epistemically impossible for S. So, just like traditional fallibilists, the contextualist and interest-relativist must deny the link between knowledge and epistemic modality posited by (5).

Interest-relativists and contextualists can arguably avoid the arbitrary threshold problem, explain the value of knowledge, and preserve the conceptual links between knowledge, action, and inquiry. However, they cannot preserve the conceptual links between knowledge, evidence, and epistemic possibility, they still face the Gettier problem, and they cannot allow that knowledge is closed under competent deduction. So four of our eight intuitive data provide evidence against interest-relativist and contextualist versions of fallibilism in addition to traditional forms of fallibilism.

One might hold that this is not as bad a result for the contextualist as the interest-relativist. This is because contextualists *can* accept the non-metalinguistic statement of infallibilism, that S knows that P iff P is certain. They can hold that, in the current context, 'knows' picks out "knowing with certainty"—even if in other contexts it picks out something that does not require certainty. And contextualists can happily endorse the above connections between evidence, epistemic possibility, and *knowledge with certainty*, and they can accept that knowledge with certainty is closed under competent deduction. I think, however, that the contextualist's inability to accept the relevant metalinguistic theses connecting 'knows,' 'evidence,' 'may,' and so on remains a serious cost. When we fix our attention on ordinary, non-philosophical contexts, it

remains implausible that the proposition expressed by "I know that the bank is open, but perhaps it's not" could ever be true. Similarly, sentences like those used by authors writing about probability in Section 4, about the connection between knowledge and evidence, might easily be uttered in low-standards contexts—e.g., contexts in which we are engaging in mundane, low-stakes reasoning and are not at all concerned about Cartesian skepticism. As for closure, one frequently touted advantage of contextualism is that it can preserve a metalinguistic formulation of *single*-premise closure (see, e.g., DeRose 1995: 27–29; Lewis 1996: 563–564). If this thesis is plausible, a metalinguistic formulation of multi-premise closure should also be plausible.

10 Counting the Costs

I have presented a cumulative case argument for infallibilism, citing eight advantages infallibilism enjoys over fallibilism. Infallibilism can explain why knowledge is (1) qualitatively different from non-knowledge, (2) uniquely valuable, and (3) not possessed by subjects in Gettier cases. Traditional versions of fallibilism have a harder time allowing for these claims. In addition, infallibilism allows us to accept that knowledge is (4) evidence, (5) a basis for epistemic modals, (6) an inquiry stopper, (7) a basis for action, and (8) closed under competent deduction. Traditional fallibilists cannot accept any of these claims, or can do so only by denying plausible claims about the relation between probability and evidence, the relation between probability and possibility, and so on. Interest-relativists and contextualists can easily accept (1), (2), (6), and (7), but cannot as easily accept (3), (4), (5), or (8).

In spite of these advantages, infallibilism is widely rejected today as an unacceptably skeptical theory of knowledge. If—as I am granting here for the sake of argument—the only facts certain for us are *a priori* truths and facts about our mental lives, infallibilism implies that we cannot know anything about the external world, like what we ate for breakfast this morning or what we are wearing.

Infallibilists can try to offer error theories for why it wrongly appears to us that we know more than we do. BonJour (2010: 71–72) suggests that we often wrongly attribute knowledge of propositions that are not really certain because we wrongly *think* that they are certain:

> [A]n ordinary person may reasonably regard the justification for a belief as conclusive even where deeper philosophical insight shows, or at least seems to show, that it is not. Consider, for example, a case where an ordinary person seems to himself or herself to be perceiving a standard sort of "medium-sized" physical object at close range and under good conditions, and believes on this basis that such an object is there. Even if the person's justification in such a case is not in fact conclusive (because of subtle philosophical objections having to do

with, for example, Cartesian demons or the possibility of being a brain-in-a-vat), it is easy to see how it might nonetheless seem to him or her to be conclusive, leading to a self-attribution of knowledge (and to attributions of knowledge to others whose situation is similar).

A second error theory, suggested by Davis (2007), Fumerton (2010: 251), BonJour (2010: 73), and Kyriacou (this volume: section 2), is that we attribute knowledge in cases where we recognize that it is not present because we are engaging in *loose talk*. Just as we may say "it's 3:00" when in fact it's 3:02, and the difference between 3:00 and 3:02 is not important for conversational purposes, we may say "I know that the bank is open" when we do not in fact know that the bank is open, if we come close enough to knowing for conversational purposes.

While fully evaluating the plausibility of these error theories is beyond the scope of this essay, their possibility shows that infallibilism's implying that most of our ordinary knowledge-attributions are false is not *conclusive* reason to reject the theory: it needs to be weighed against its other advantages and disadvantages. I summarize these in Table 4.1.

Evaluating the cumulative force of multiple evidences is a difficult task. Some evidences may be weightier than others, and there may be subtle dependencies among the evidences that need to be taken into account.[16] Still, it seems to me that the skeptical costs of infallibilism are worth paying to reap all the explanatory benefits listed in Table 4.1. Infallibilism lets us account for eight intuitions that traditional fallibilism cannot explain. As for interest-relativism and contextualism, there are many independent objections to these views based on other counterintuitive consequences they have (e.g., Anderson 2015). When we take these together with the four intuitive data that even they cannot explain, it seems to

Table 4.1

	Infallibilism	Traditional fallibilist theories	Contextualism and interest-relativism
Knowledge qualitatively different	Yes	No	Yes
Knowledge uniquely valuable	Yes	No	Yes
No Gettier problem	Yes	No	No
Knowledge is evidence	Yes	No	No
Knowledge that P makes ~P epistemically impossible	Yes	No	No
Knowledge lets us act	Yes	No	Yes
Knowledge lets us end inquiry	Yes	No	Yes
Knowledge is closed under competent deduction	Yes	No	No
We know a lot	No	Yes	Yes

me that the disadvantages faced by interest-relativism and contextualism also outweigh the advantage of counting more of our ordinary knowledge-ascriptions as true.

Some philosophers maintain that skepticism about knowledge in ordinary cases is so implausible that one should never accept a theory of knowledge that has skeptical implications. In allegedly "Moorean" fashion, they assert that the premises of a skeptical argument are never more certain than the ordinary knowledge claims they seek to undermine (Lewis 1996: 549; Pryor 2000: 518). But even if this is true for particular arguments, several arguments together can have a cumulative force beyond that of each particular argument, and the *disjunction* of their premises may indeed be more plausible than the ordinary knowledge claims Mooreans cling to.

Compare: I am at a trial for my best friend, who has been accused of murdering his wife. I am almost certain that he did not kill her: I have known him my whole life and this is simply not something he would do. This is a "Moorean fact" for me. When a witness comes forward to testify that she saw him do it, I am rational in judging that it is more likely that she is lying or confused than that my friend really did it. And I am rational in making similar judgments about the credibility of any individual piece of evidence, if that is the only evidence presented. But when four more witnesses come forward, the murder weapon is found to have my friend's fingerprints on it, he is demonstrated to have motive, means, and opportunity, and he confesses to the murder, the cumulative force of *all* these evidences should lead me to abandon my faith in my friend. It is *not* more plausible that *all* these witnesses are lying or mistaken, that my friend's fingerprints got on the weapon some other way, *and* that he was pressured into confessing to protect someone else, and so on, than that I was mistaken about what my friend is capable of.

Likewise, it is not more plausible—or so it seems to me—that the Gettier problem is irresolvable, that the unique status and value of knowledge are inexplicable, and that knowledge is not evidence, a basis for epistemic modals, a basis for action, an inquiry-closer, or closed under competent deduction than it is that I do not know that I have hands. At a certain point the evidence, or the arguments, must overturn common sense.

I think that the evidences for infallibilism presented in this essay take us to that point. But even if you are not persuaded of that, I hope to have at least convinced you that infallibilism should not be summarily dismissed on account of its skeptical consequences. Its plausibility needs to be evaluated in light of all the intuitive evidence we have. This includes not only our ordinary claims to know things but also our intuitions about what kind of a thing knowledge is, and the roles that it plays in epistemology more broadly.[17]

Notes

1 In Climenhaga (2021, forthcoming), I develop a fuller theory of knowledge, which I there call (capital-I) 'Infallibilism,' that implies what I am here calling (lowercase-i) 'infallibilism.' My definition here is close to Dodd's (2011: 665), who defines infallibilism as the thesis that knowledge requires epistemic probability 1. Several other philosophers (e.g., Audi 2003: 224; Stanley 2005: 127; BonJour 2010: 57; Brown 2018) define infallibilism as the thesis that S knows that P only if S has entailing justification/evidence/grounds for P. This definition implies that S can know that P on the basis of evidence E even if E is not itself certain for P. Neta (2011: 668–669) sees this consequence as consistent with infallibilism, but I agree with Dougherty (2011: 140–141) that a theory (such as Neta's) that allows knowledge wholly on the basis of uncertain grounds should be considered fallibilist.

2 For more on the nature of epistemic certainty, see Climenhaga (2021, forthcoming).

3 In Climenhaga (2021), I consider several additional proposals for reconciling fallibilism with these intuitions that I omit here because of space considerations.

4 Many philosophers deny that probability 1 is sufficient for certainty in cases involving infinities. 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 apparently 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. Williamson's view is more radical than this: he thinks that probability 1 and certainty come apart even in cases that do not involve infinite sample spaces of this sort. I compare Williamson's theory of knowledge with the form of infallibilism I defend in Climenhaga (2021).

5 Brown (2018) argues that combining infallibilism with a less skeptical view about the scope of our knowledge leads to a variety of counterintuitive results. She also argues that non-skeptical infallibilists have just as hard a time as fallibilists accepting several of the intuitive claims I go on to discuss, including (1), (6), and (8). For the most part, though, these arguments leave open the possibility that skeptical infallibilists can accept these claims.

6 For an overview of theories of knowledge based on these and other popular conditions, see Shope (1983) and Ichikawa and Steup (2014).

7 I sometimes speak loosely of (4) as the claim that knowledge is evidence, even though, strictly speaking, what it says is that *what is known* is evidence. In this I follow the language of Williamson and others who endorse (4).

8 I am setting aside Jeffrey conditionalization, because this approach implies radical subjectivism about probabilities: see Williamson (2000: 216).

9 Brown (2018: ch. 4) argues that (4) has counterintuitive consequences when combined with a non-skeptical view about the scope of our knowledge. But this does not explain the intuitive plausibility of (4) itself; and to the extent that this remains unexplained, it is evidence for a skeptical form of infallibilism over a non-skeptical form of fallibilism.

10 For a fallibilist error theory of the infelicity of CKAs that denies (5), see Dougherty and Rysiew (2009, 2011).

11 Stanley and Hawthorne's discussions suggest that knowledge is not only sufficient but also necessary for rational action. However, while this necessity claim may be plausible for using a proposition as a premise in practical reasoning, it is not plausible for *acting as if* a proposition is true. As I am using

this phrase, φ-ing counts as "acting as if P" just in case φ-ing is the optimal course of action conditional on P being true. In cases where there are no deontological moral considerations at play, I assume that this is equivalent to φ-ing maximizing expected utility conditional on P being true. So understood, if one can rationally use P as a premise in practical reasoning, one can rationally act as if P—so if knowledge that P is sufficient for the former, it is sufficient for the latter. On the other hand, even if knowledge that P is necessary for rationally using P as a premise in practical reasoning, it is not necessary for rationally acting as if P. This is because it can be rational to act as if P when P is (as it turns out) false, and so not known. (In the example below, even if P were false, it would be rational for S to φ if the loss for P being false was reduced to $1, provided that the epistemic probability of P for S remains at 0.9.)

12 This is an unlikely stipulation to hold true in ordinary cases. In ordinary books, the claims made will be positively relevant to each other. However, even with some positive relevance, given enough claims we can still make the probability of their conjunction arbitrarily low.

13 In response to BonJour's argument that lottery cases show that fallibilists must reject closure, Pritchard and Turri (2014) suggest that rejecting probability theory as measuring justification lets fallibilists keep closure, writing, "If we pair the [fallibilist] conception of knowledge with a qualitative [i.e., non-probabilistic] model of justification…, then it's no longer clear that the resulting view will fail to respect the closure of knowledge under conjunction, because the rules of probability theory employed in BonJour's argument don't obviously apply to the qualitative categories."

14 For further counterintuitive consequences of the claim that one can know that all of the claims in one's book are true, see Christensen (2004: ch. 3.3–3.4). (Christensen's focus is on rational belief, but his remarks are equally applicable to knowledge.) For responses to other common objections to the Preface Paradox, see Christensen (2004: ch. 3.1) and Easwaran and Fitelson (2015: 65–70).

15 Strictly speaking, when I talk about whether contextualists can accept (1)–(8), I mean whether they can accept meta-linguistic analogs of these claims. For example, in saying that they can accept (1), what I mean is that they can accept that there is a qualitative difference between the state picked out by "S knows that P" in a context and the state picked out by "S justifiably and truly believes that P but the degree of S's justification is just shy of that required for knowledge" in that same context. I make these meta-linguistic formulations explicit when needed below, and at the end of this section discuss the relevance of whether contextualists can accept the non-meta-linguistic versions of (1)–(8).

16 In Climenhaga (2021), I address these complexities by presenting a Bayesian version of the cumulative case presented informally here, in which the weight of and dependencies among different evidences are formally quantified. I also consider some additional objections to infallibilism, and argue that they don't substantially add to the evidential weight of our ordinary intuitions and attributions of knowledge.

17 I am grateful to Robert Audi, Blake Roeber, Daniel Immerman, Al Hájek, Mark Satta, Greg Stoutenburg, Christos Kyriacou, and an anonymous reviewer for feedback on earlier drafts of this paper, and to audiences at Western Michigan University, Deakin University, the Society for Exact Philosophy, and Pacific APA for feedback on oral presentations of this work.

References

Anderson, Charity (2015). "On the Intimate Relation of Knowledge and Action," *Episteme* 12: 343–353.

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).

Christensen, David (2004). *Putting Logic in Its Place* (Oxford: Oxford University Press).

Climenhaga, Nevin (2021). *Knowledge and Certainty: A Defense of Infallibilist Epistemology.* Manuscript.

Climenhaga, Nevin (forthcoming). "How Infallibilists Can Have it All," *The Monist.*

Davis, Wayne (2007). "Knowledge Claims and Context: Loose Use," *Philosophical Studies* 132: 395–438.

De Finetti, Bruno (1979/2008). "Lecture VI: Bayes' Theorem," in *Philosophical Lectures on Probability: Collected, Edited, and Annotated by Alberto Mura* (Springer), pp. 31–45.

DeRose, Keith (1995). "Solving the Skeptical Problem," *The Philosophical Review* 104: 1–52.

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

Dodd, Dylan (2011). "Against Fallibilism," *Australasian Journal of Philosophy* 89: 665–685.

Dougherty, Trent (2011). "Fallibilism," in Duncan Pritchard and Sven Bernecker (eds.), *The Routledge Companion to Epistemology.*

Dougherty, Trent, and Patrick Rysiew (2009). "Fallibilism, Epistemic Possibility, and Concessive Knowledge Attributions," *Philosophy and Phenomenological Research* 78: 123–132.

Dougherty, Trent, and Patrick Rysiew (2011). "Clarity about Concessive Knowledge Attributions: Reply to Dodd," *Synthese* 181: 123–132.

Dutant, Julien (2016). "How to Be an Infallibilist," *Philosophical Issues* 26: 148–171.

Easwaran, Kenny, and Branden Fitelson (2015). "Accuracy, Coherence and Evidence," in T. Szabo Gendler and J. Hawthorne (eds.), *Oxford Studies in Epistemology*, vol. 5 (Oxford: Oxford University Press).

Eells, Ellery, and Branden Fitelson (2000). "Measuring Confirmation and Evidence," *Journal of Philosophy* 97: 663–172.

Fumerton, Richard (2010). "Fencing out Pragmatic Encroachment," *Philosophical Perspectives* 24: 243–253.

Gettier, Edmund L. (1963). "Is Justified True Belief Knowledge?" *Analysis* 23: 121–123.

Hacking, Ian (1967). "Slightly More Realistic Personal Probability," *Philosophy of Science* 34: 311–325.

Hannon, Michael (2014). "Fallibilism and the Value of Knowledge," *Synthese* 191: 1119–1146.

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

Hawthorne, John (2005). "The Case for Closure," in Matthias Steup and Ernest Sosa (eds.), *Contemporary Debates in Epistemology* (Malden, MA: Blackwell), pp. 40–55.

Howard-Snyder, Daniel, Frances Howard-Snyder, and Neil Feit (2003). "Infallibilism and Gettier's Legacy," *Philosophy and Phenomenological Research* 66: 304–327.

Howson, Colin (1991). "The 'Old Evidence' Problem," *The British Journal for the Philosophy of Science* 42: 547–555.

Ichikawa, Jonathan Jenkins, and Matthias Steup (2014). "The Analysis of Knowledge," in *The Stanford Encyclopedia of Philosophy*, <http://plato.stanford.edu/archives/spr2014/entries/knowledge-analysis/>.

Jaynes, E.T. (2003). *Probability Theory: The Logic of Science* (Cambridge: Cambridge University Press).

Kappel, Klemens (2010). "Expressivism about Knowledge and the Value of Knowledge," *Acta Analytica* 25: 175–194.

Kirkham, Richard (1984). "Does the Gettier Problem Rest on a Mistake?" *Mind* 93: 501–513.

Kyriacou, Christos (2017). "Bifurcated Sceptical Invariantism: Between Gettier Cases and Saving Epistemic Appearances," *Journal of Philosophical Research* 42: 27–44.

Kyriacou, Christos (this volume). "Moderate Pragmatic Skepticism, Moorean Invariantism and Attributions of Intellectual Virtue/Vice," in Christos Kyriacou and Kevin Wallbridge (eds.), *Skeptical Invariantism Reconsidered* (Routledge).

Lewis, David (1996). "Elusive Knowledge," *Australasian Journal of Philosophy* 74: 549–567.

Lipton, Peter (2001). "Is Explanation a Guide to Inference? A Reply to Wesley Salmon," in Giora Hon and Sam Rakover (eds.), *Explanation: Theoretical Approaches and Applications* (London: Kluwer), pp. 93–120.

Nelkin, Dana (2000). "The Lottery Paradox, Knowledge, and Rationality," *The Philosophical Review* 109: 373–409.

Neta, Ram (2011). "A Refutation of Cartesian Fallibilism," *Noûs* 45: 658–695.

Pearl, Judea (2000). *Causality: Models, Reasoning, and Inference* (Cambridge: Cambridge University Press).

Pritchard, Duncan, and John Turri (2014). "The Value of Knowledge," in *The Stanford Encyclopedia of Philosophy*, <http://plato.stanford.edu/archives/spr2014/entries/knowledge-value/>.

Pryor, James (2000). "The Skeptic and the Dogmatist," *Noûs* 34: 517–549.

Reed, Baron (2012). "Fallibilism," *Philosophy Compass* 7: 585–596.

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

Stanley, Jason (2005). *Knowledge and Practical Interests* (Oxford: Oxford University Press).

Stanley, Jason (2008). "Knowledge and Certainty," *Philosophical Issues* 18: 35–57.

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

Zagzebski, Linda (1994). "The Inescapability of Gettier Problems," *Philosophical Quarterly* 44: 65–73.