

FALLIBILISM AND THE AIM OF INQUIRY

Bob Beddor

1 Stage Setting

A murder has been committed. Arriving at the scene, Hercule Poirot uncovers clues indicating that the butler is guilty. Being a thorough detective, Poirot does not immediately wrap up the case; he continues his investigation. But how long must he continue? When is he entitled to conclude his inquiry and announce that the butler did it?

According to an increasingly prominent view, our mustachioed detective is entitled to close his investigation when he comes to *know* who committed the murder. While different philosophers put the idea in somewhat different ways, the core thought goes something like this:

K-AIM: The aim of inquiring into a question Q is to come to know the answer to Q .¹

Part of the appeal of K-Aim is that it offers to shed light on the value of knowledge. On the resulting picture, knowledge is important because it functions as an *inquiry-stopper*.

However, K-Aim stands in tension with another highly attractive epistemological doctrine: fallibilism. Here too, different epistemologists formulate this doctrine in slightly different ways. Here I'll focus on what I take to be a particularly plausible version of the core fallibilist thought:

FALLIBILISM: It's possible for a rational agent to know p without being absolutely certain that p .²

¹See Kvanvig 2009; Kappel 2010; Kelp 2011, 2014, 2018, forthcoming; Rysiew 2012; Friedman 2013, 2017.

²Fallibilism—in some form or other—has been widely taken for granted in contemporary epistemology. For example, Cohen writes, “the acceptance of fallibilism is nearly universal in epistemology” (1988: 91). For endorsements of Fallibilism, see Stanley 2005b; Rysiew 2007; Reed 2013; Worsnip 2015; Brown 2018, among many others. We will consider some recent opposition to Fallibilism in due course.

To see the source of the tension, take a case of fallible knowledge—a case where an agent knows p , even though they are not completely certain of p . Now, offer this agent the opportunity to acquire decisive evidence regarding p . Intuitively, our agent is under rational pressure to inquire further by acquiring this evidence. But this conflicts with K-Aim, which says that our agent has already attained the aim of inquiry.

After laying out the problem (§§2-3), I consider some strategies for reconciling K-Aim and Fallibilism and find them wanting (§4). So we are forced to choose between the two principles—a choice that I take up in §§5-6. I submit that once we take a closer look at the considerations motivating Fallibilism, we find that K-Aim loses much of its initial appeal. These considerations can also be used to motivate an alternative picture of the aim of inquiry, according to which inquiry aims at maximizing the epistemic value of our credal states. I go on to explore the consequences of this picture for our views on the nature of interrogative attitudes, the norm of practical reasoning, and the value of knowledge (§7).

2 Bringing out the tension: two cases

In what circumstances can we know something without being rationally certain of it? Cases of inductive knowledge offer one natural candidate. Consider:

Murine Research Mia is a scientist who forms the hypothesis h : a particular drug (Accuphine) causes hyperactivity in mice. Mia conducts a number of experiments that support h . Eventually, she conducts enough experiments to know that h is true. But she still is not completely certain of h , and rationally so: she rationally assigns at least some credence, however slight, to the possibility that all of her experimental results are attributable to confounding factors.

Presumably, if Fallibilism is true, cases of this sort are possible. To cause trouble for K-Aim, let us continue our tale:

More Murine Research One day Mia receives an email from a researcher at another university. In their email they announce that they have just completed the most comprehensive study to date on the effects of Accuphine on mice, with a specific focus on whether Accuphine causes hyperactivity. As a courtesy, they have provided all their data as an attachment.

What should Mia do? There is, I think, a strong intuition that Mia is rationally permitted to read the results of the latest study. After all, she is not absolutely certain of h , and the results of the latest study might change her mind one way or the other. Indeed, I think there are ways of fleshing out the case where intuition tells us not merely that she is permitted to read the results, but that she is rationally required to do so. However, this intuition stands in tension with K-Aim. After all, downloading and reading the attachment constitutes further inquiry into the question: *Is h true?* According to K-Aim, she has already achieved the aim of inquiry on this question. So there should be no point in inquiring further.

There are a number of ways one might seek to resolve this tension; I'll discuss some strategies in §4. For now, I want to note that nothing special hinges on the details of this particular case, or cases of inductive knowledge more generally. It's easy to concoct other examples that make the same point:

Ancient History Tess is about to take her Roman history test. She learned the material well, but it has been some time since she reviewed. She is fairly confident in r : *The Western Roman Empire fell in 476 CE*. However, she assigns some credence to the possibility that she got the dates wrong. As a matter of fact, her memory is correct.

If Fallibilism is true, then presumably Tess can know r .³ To put pressure on K-Aim, we need only add a wrinkle:

More Ancient History Before the test, Tess' teacher announces: "Since it's the last day of class, I'll be nice. One of the questions you'll be asked is, 'When did the Western Roman Empire fall?' You now have five minutes to review your materials before the test begins." As it happens, Tess has her textbook, *Ancient Roman Civilization*, in front of her. To check the date, all she would need to do is to flip it open to the relevant page and peruse the text.

What should Tess do? Intuitively, she should consult her textbook. After all, she isn't completely sure when the Roman Empire fell. And she now has the opportunity—at no cost to herself—to attain complete certainty on this matter, or at least something closer thereto. But this intuition conflicts with K-Aim. Since Tess knows the answer to this question, K-Aim says that she has attained the aim of inquiry on this topic. So it would be pointless for her to inquire further.

³This case bears some resemblance to Radford's 'unconfident examinee' (1966). But whereas Radford's examinee is fumbling and hesitant in his answers, we need not say this about Tess. All we require is that her degree of confidence in r is less than 1.

There is a simple recipe for whipping up structurally similar cases. *Step 1:* Describe an agent who is inquiring into the truth of p . *Step 2:* Stipulate that as a result of their inquiry, they come to know p , even though they are not yet rationally certain of p . *Step 3:* Give them the opportunity to acquire decisive evidence regarding p , at no cost to themselves. By K-Aim, they are under no rational pressure to look at the evidence. But this conflicts with the intuition that they rationally ought to take a look.

3 The source of the tension

Our examples reveal a tension between K-Aim and Fallibilism. But K-Aim and Fallibilism are not logically inconsistent. So whence does the tension arise? This section offers a diagnosis, drawing on the resources of epistemic decision theory.

First, some background. Traditional decision theory starts with idea that rational agents assign practical value to outcomes. Epistemic decision theory starts with the epistemic analogue of this idea: rational agents assign epistemic value to credal states.⁴

What makes a credal state epistemically valuable? A very natural thought—one which runs through much of epistemic decision theory—is that part of the answer involves accuracy. That is:

VALUABLE ACCURACY If a credence function c is more accurate than a credence function c' , then c is more epistemically valuable than c' .

This idea should be appealing to anyone who subscribes to the “veritist” idea that the ultimate epistemic goods consist in attaining truth and avoiding error. But it is important to note that Valuable Accuracy is considerably weaker than veritism. Valuable Accuracy does not claim that accuracy is the only source of epistemic value, only that it is one source. (More on this in §6.)

What makes a credal state accurate? Intuitively, the accuracy of a credal state corresponds to its “proximity” to the truth; a .9 credence in a truth is more accurate than a .2 credence in a truth. More generally:

ALETHIC PROXIMITY If c assigns higher credences to more true propositions and lower credences to more false propositions than c' , then c is more accurate than c' .

⁴For a sampling of important contributions to epistemic decision theory, see [Joyce 1998](#); [Greaves and Wallace 2006](#); [Moss 2011](#); [Pettigrew 2016](#); [Schoenfield 2016](#); among many others.

There are a number of different formal measures of accuracy that respect Alethic Proximity.⁵ For our purposes, we need not take a stand on which of these measures is correct. All we need is the minimal thesis of Alethic Proximity itself.

Having laid the groundwork, let us tie this back to inquiry. What is the point of engaging in inquiry? A natural thought is that as we inquire, we gain evidence, and we revise our credences in light of this evidence. But why revise our credences in light of new evidence? Here's another natural thought: by doing so, we improve our credences from the epistemic point of view. That is:

EV-AIM The aim of inquiring into a question Q is to make your credence in the answer to Q as epistemically valuable as possible.

Combine these ingredients (EV-Aim + Valuable Accuracy + Alethic Proximity) and we get a plausible diagnosis of why K-Aim stands in tension with Fallibilism. Return to our cases from §2. In **Ancient History**, we stipulated that Tess knows r (*The Roman Empire fell in 476*), but she doesn't know this with complete certainty. To make things concrete, let's say that her credence in r is .96. By Alethic Proximity, it follows that her credence in r has a high degree of accuracy, but not the highest degree thereof. By Valuable Accuracy, it follows that her credence in this proposition is not maximally valuable; room for improvement remains. By EV-Aim, it follows that she has not attained the aim of inquiring into the question, *When did the Roman Empire fall?* So we avoid the implausible consequence that there is no need for her to glance down at the pages of her textbook.

Now, from this it does not automatically follow that it is rational for Tess and Mia to inquire further. However, there are some plausible bridge principles that we can use to fill this gap. For example:

AIM-RATIONALITY BRIDGE If it is rational for you to pursue aim α , and it is rational for you to believe that you have not yet attained α , then it is rational for you to continue pursuing α .

Mia and Tess are aware that they are not certain of the answers to the relevant questions. So it is rational for them to believe that they have not yet attained

⁵One popular candidate is the Brier score. For some set of propositions X , the Brier score of a credence function c defined over X , as held at a world w , is given by the formula:

$$BS(c, w) = \sum_{p \in X} (c(p) - w(p))^2 \quad (1)$$

where $w(p) = 1$ if p is true at w , and 0 otherwise. This provides a measure of inaccuracy: a credence of 1 in a true proposition will get a Brier score of 0, which is the best score possible. A credence of 1 in a false proposition will get a score of 1, the worst score possible.

maximally epistemically valuable credal states with regards to the question at hand. By EV-Aim and Aim-Rationality Bridge, it is rational for them to continue their inquiries.⁶

Let's take stock. §2 argued that there is a tension between K-Aim and Fallibilism. The argument was case-based: I gave scenarios where it is rational for fallible knowers to continue their inquiries. This argument did not itself presuppose the decision theoretic machinery introduced in this section. Thus merely rejecting this machinery is not sufficient for resolving the tension. That said, the theoretical framework introduced here provides insight into the source of the tension. According to my diagnosis, the tension arises from the plausible idea that inquiry aims—at least in part—at rendering one's credences as accurate as possible, and that any non-extremal credence is guaranteed to fall short of perfect accuracy. Combined with Fallibilism, this conflicts with K-Aim.

4 Reconciliation Strategies

I now turn to consider three strategies for trying to reconcile K-Aim and Fallibilism. I argue that none of these reconciliation strategies pass muster. We are thus forced to choose between the two principles—a choice I take up in §§5-6.

4.1 Impurism

Our first reconciliation strategy is to embrace some form of 'impurism': the view that knowledge depends practical factors.⁷ While there are different ways of de-

⁶We could also arrive at this result a different way, drawing on further resources from epistemic decision theory. Oddie 1997 proves that, given certain assumptions, it always maximizes expected epistemic value to consult new evidence, provided the cost of doing so is negligible. More precisely, Oddie shows that the expected epistemic value of your current credal state is always less than or equal to the expected epistemic value of the credal state that will result from gathering new evidence and conditionalizing on it, and strictly less when there's a chance that the new evidence will affect your credences. To illustrate, consider again **Ancient History**. Since Tess is not certain of r , there's a good chance that if she consults her textbook, this will impact her credence in r . In particular, there's a chance that she will read that the Western Roman Empire did fall in 476, in which case she will presumably raise her credence in r (to, say, .99). There's also a chance—albeit a more remote one—that she will read that the Western Roman Empire fell on some other date, in which case she will lower her credence in r . What Oddie's theorem shows is that the expected value of her current credal state is strictly lower than the expected value of the credal state that will result from checking. So another way of arriving at the result that it is rational for Mia and Tess to continue their inquiries would be to appeal to Oddie's theorem, together with the assumption that it is epistemically rational to do whatever maximizes expected epistemic value. For further discussion of the applications—and limitations—of Oddie's theorem, see Beddor 2019; Campbell-Moore and Salow 2020.

⁷AKA 'pragmatic encroachment.' For defenses, see e.g., Fantl and McGrath 2002, 2009; Stanley 2005a; Weatherson 2012; Ross and Schroeder 2014. Needless to say, impurism is controversial. For

veloping this strategy, the basic idea goes like this: knowledge does not entail certainty, so Fallibilism is true. But knowledge is sensitive to pragmatic considerations. In both **Murine Research** and **Ancient History**, these pragmatic considerations deprive our protagonists of knowledge. So our protagonists have not achieved the K-Aim after all.

In order to evaluate this strategy, let us take a closer look at details. According to what is perhaps the most prominent form of impurism, knowledge depends on stakes: in high stakes situations, one needs more evidence to know a proposition than in low stakes situations. This idea is a non-starter for handling our cases, since we can stipulate that the practical stakes for our agents are quite low. (Imagine that in **Ancient History**, Tess is not taking the class for credit.) Still, the intuition that they epistemically ought to inquire remains.

Perhaps, then, we should opt for a somewhat different version of impurism. Adopting some terminology from [Anderson and Hawthorne 2019](#), say that p is *practically adequate* for you if and only if the action you actually prefer, given your epistemic position, is the same as the action you prefer conditional on p . Impurists might propose:

PRACTICAL ADEQUACY CONDITION S knows p only if p is practically adequate in S's situation.

This condition entails that when our agents are given the opportunity to receive cost-free evidence, they lose their knowledge. Take Tess in **Ancient History**. Once Tess hears her teacher's announcement, the action that Tess actually prefers given her current epistemic position is *checking the textbook*, which is different from the action she prefers conditional on r (*not bothering to check*). So r is no longer practically adequate for her; by the Practical Adequacy Condition, she loses her knowledge of r .

However, this approach faces difficulties of its own. Imagine that after hearing her teacher's announcement, Tess reflects as follows: "I'm pretty sure that I know when the Roman Empire fell. But I'm not completely certain I know it, so I might as well check." This reflection seems perfectly natural. But if the Practical Adequacy Condition is correct, it is hard to make sense of Tess' train of thought. After all, Tess is aware that r is not practically adequate. So given the Practical Adequacy Condition, she should be in a position to realize that she doesn't know when the Roman Empire Fell after all.

A second difficulty for the Practical Adequacy Condition comes from [Anderson and Hawthorne 2019](#), who note that it licenses an unwelcome epistemic instability. Imagine that as Tess starts to flip to the relevant page in her textbook,

challenges, see e.g., [Russell and Doris 2008](#); [Reed 2010](#); [Roeber 2018](#).

her teacher pipes up: “Oh, but if you do check your textbook, I’ll charge you \$10,000.” Now, r becomes practically adequate for Tess, allowing her to regain her knowledge of r . Suppose a moment later the teacher announces they were just joking about the cost of checking. By the Practical Adequacy Condition, she loses her knowledge once again. Now, it may well be that cost considerations affect whether it is rational to pursue an investigation. Still, it seems implausible that knowledge can be gained and lost so easily.

So it proves difficult to develop the impurist reconciliation strategy in a way that delivers plausible results in our cases. Furthermore, the impurist reconciliation strategy will be unwelcome to anyone who agrees with our diagnosis of the source of the tension in §3. On the diagnosis developed there, K-Aim and Fallibilism stand in tension because a fallible knower can still fall short of the epistemically optimal credal state. Nothing in this explanation presupposed a purist theory of knowledge. Our explanation only relied on certain independently plausible tenets of epistemic decision theory (Valuable Accuracy, Alethic Proximity, or EV-Aim). So proponents of the impurist reconciliation are put in the uncomfortable position of being forced to reject one of these tenets.

4.2 Aims vs. rational requirements

A second possible reconciliation strategy is to look closer into the relationship between aims and rational requirements. Even if Mia and Tess have attained the aim of inquiry, it might still be rational for them to inquire further. After all, we should distinguish between attaining some aim α and rationally believing that you have attained α . Suppose an archer hits the bullseye on a distant target, but a spectator tells the archer that they missed. It would be rational for the archer to try again.⁸

However, there is an important difference between the archer and our two protagonists, Mia and Tess. The archer has reason to think they did not attain their aim. Not so with Mia and Tess: they are merely told they have the option of gathering more evidence on the topic at hand. Impurist maneuvers aside, this information does not give them any reason to think that their current beliefs on this topic do not amount to knowledge.

Here is one way of putting this point more generally. Presumably there is *some* connection between the aim of inquiry and the rationality of inquiry. I have already suggested one bridge principle connecting the two (Aim-Rationality

⁸A similar distinction has loomed large in recent defenses of knowledge norms of belief and assertion. Proponents of such norms often emphasize the need to distinguish between complying with a norm N and rationally believing that one has complied with N . For relevant discussion, see DeRose 2002; Sutton 2007; Kelp and Simion 2017; Williamson forthcoming, among others.

Bridge). As we saw, this bridge principle, when combined with EV-Aim and Fallibilism, allows us to secure the result that it is rational for Mia and Tess to inquire further. But if we use K-Aim instead of EV-Aim, the derivation no longer goes through.

We can strengthen this point. Aim-Rationality Bridge stated a sufficient condition for when it is rational to continue pursuing your aims. But we can also formulate sufficient conditions for when you are not rationally required to do so:

CALL IT QUILTS If you are pursuing some aim α , and it's rational for you to believe that you have already attained α , then you are not rationally required to continue pursuing α .

On a natural way of filling out our cases, Mia and Tess rationally believe (though they are not certain) that they already know the answers to the question at hand. So Call it Quits—when combined with K-Aim and Fallibilism—predicts that our protagonists are not rationally required to engage in further inquiry. This conflicts with the intuition that they rationally ought to do so.

The same point holds even if we fall back to weaker sufficient conditions, e.g.:

WEAK CALL IT QUILTS If you are pursuing some aim α , and you *know* that you have already attained α , then you are not rationally required to continue pursuing α .

Presumably, if one can know p without being rationally certain of p , then one can also know that one knows p without being rationally certain of p .⁹ And so we should be able to cook up versions of our cases where Mia and Tess not only rationally believe that they have attained knowledge on the question at hand, but they know that they have done so. Still, if they are not certain of the answers, it may well be rational for them to continue inquiring.

So I doubt that we can reconcile K-Aim and Fallibilism simply by distinguishing between the aim of inquiry and the rational norms governing inquiry. Here too, it is worth noting that this reconciliation strategy will be unappealing to anyone sympathetic to §3's diagnosis of the source of the tension. According to the diagnosis offered there, the tension arises from the fact that K-Aim, when combined with Fallibilism, is inconsistent with EV-Aim. This suggests that the fundamental problem concerns the aim of inquiry itself, rather than the connection between aims and rational norms.

⁹Most theories that make room for fallible knowledge will also make room for fallible higher-order knowledge along these lines. Consider the view that knowledge is just safe belief, and that belief does not require credence 1. Presumably, one can safely believe that one safely believes p without assigning credence 1 to p .

4.3 Aim pluralism

This leads us to a final reconciliation strategy, which seeks to defuse the tension between K-Aim and EV-Aim. Why not hold that there are multiple aims of inquiry? On this view, when someone comes to know the answer to a question, they will thereby achieve one of the aims of inquiry. But they will not have thereby achieved every aim of inquiry, since they will not have arrived at an epistemically optimal credal state.

While this is a simple path to harmony, it has some unappealing features. First, the pluralist approach does not provide helpful guidance for inquirers. Go back to Poirot. We wanted our theory to tell us when he is entitled to stop investigating and move on to other cases. What does the pluralist approach say? Well, it says that he will have achieved one of the aims of inquiry when he comes to know that the butler did it. And he will have achieved another aim only when he comes to achieve certainty on the matter. How is this helpful for answering the question of what Poirot rationally ought to do?

Perhaps if we embrace pluralism about aims, we should also be pluralists about rational “ought”s. Relative to the K-Aim, Poirot ought to conclude his case when he comes to know the butler did it. Relative to the EV-Aim, he ought to continue investigating. But this only worsens the problem. We’d like some way of adjudicating between these conflicting ought claims; we’d like our theory to tell us what Poirot ought do *full stop*. The pluralist approach provides no help.

A further strike against the pluralist approach is that it has trouble making sense of our intuitions about the cases from §2. In **Ancient History**, there is a clear-cut intuition that Tess rationally ought to look at the textbook. But if we are pluralists about the aims and oughts of inquiry, we shouldn’t expect our intuitions to be so clear-cut. Rather, we should find ourselves pulled in different directions. We should be inclined to say, “Well, there’s one sense in which it would be perfectly rational for her to refuse to glance at the textbook. And there’s another sense in which she is rationally obligated to take a look.” The fact that we do not react in this way is strong evidence against the pluralist response.

4.4 Looking forward

We tried out three strategies for reconciling K-Aim and Fallibilism. All proved unsuccessful. Since the two principles cannot reside happily together, we must make a choice.

This conclusion is, I think, both surprising and significant. After all, K-Aim and Fallibilism are both popular doctrines; both are defended by prominent epistemologists; both enjoy a good deal of intuitive and theoretical appeal. So if I’ve

convinced you that two doctrines cannot live side-by-side, my main work in this paper is done.

That said, I do want to develop the argument a step further. I now turn to evaluate how we should choose between Fallibilism and K-Aim: which should we keep, and which should get the boot?

5 Rejecting Fallibilism

Let's start by considering the prospects of rejecting Fallibilism.¹⁰ Once we give up Fallibilism, we can no longer derive any inconsistency between K-Aim and EV-Aim. After all, suppose S knows p . By infallibilism, S is rationally required to have credence 1 in p . And since p is true (by the factivity of knowledge), S's credal state *vis-à-vis* p is guaranteed to be the most accurate possible.

The combination of K-Aim with infallibilism is certainly a coherent package. But is it plausible? I want to register a couple doubts on this score.

First, consider how this position handles our cases from §2. According to infallibilism, the sort of cases presented there cannot arise. Take **Ancient History**. According to the infallibilist, as long as Tess rationally assigns some positive credence to the prospect that the Western Roman Empire fell in some year other than 476, then she does not know that it fell in 476. This strikes me a difficult consequence to stomach. As we noted in §4.1, we can imagine Tess remarking to herself: "I'm pretty sure I know when the Roman Empire fell." Similarly, we can imagine that as Tess takes the test, her teacher is hovering over her shoulder, and mutters to herself: "Looks like Tess knew the answer to that one." But presumably neither Tess nor her teacher are under the impression that Tess has credence 1 in her answer. This suggests that our ordinary concept of knowledge makes room for knowledge without absolute certainty.¹¹

¹⁰While Fallibilism was taken for granted throughout much of the 80s and 90s, the last two decades have seen an infallibilist resurgence. For example, Williamson defends the view that an agent's evidence is just their knowledge (2000: chp. 9). Combined with the idea that agents should assign their evidence credence 1, this leads to the view that we should have credence 1 in everything we know. Williamson embraces this consequence, though he frames it in terms of "degrees of outright belief" rather than "credences", since he denies that degrees of outright belief can be reduced to betting dispositions. (For the purposes of this paper, I will not assume any particular connection between credences and betting behavior, so the difference between degrees of outright belief and credences will not affect any of my arguments.)

¹¹Here my intuitions align with those of many philosophers who have discussed Radford's unconfident examinee. While many have balked at Radford's claim that the examinee has knowledge while lacking belief, many agree with the weaker claim that the examinee has knowledge while lacking certainty (e.g., Armstrong 1969; Stanley 2008; McGlynn 2014; Beddor 2020a). For related evidence for Fallibilism, see the discussion of "I think I know..." constructions in Worsnip 2015.

A second challenge for the infallibilist solution comes from the observation that we seem to stand in a stronger epistemic position towards some things we know than towards others. Compare:

b : It rained in Britain sometime in the last six months.

$b \vee \neg b$: It rained in Britain sometime in the last six months or it did not.

Even if I have not been to Britain in the last six months—or checked the weather reports there—it seems I can know b . Nonetheless, I stand in a stronger epistemic position towards the tautology, $b \vee \neg b$. Intuitively, this difference in epistemic position should show up in my credences: my credence in b should be at least slightly lower than my credence in the tautology. But this is inconsistent with infallibilism, which says that I am rationally required to have credence 1 in both propositions.

While this is hardly the final word on the matter, these costs provide some impetus for exploring the other response to our tension: giving up K-Aim.

6 Rejecting K-Aim

6.1 If not knowledge, then what?

If K-Aim is false, then what is the aim of inquiry? §3 already sketched the beginning of an answer: inquiry aims at improving the epistemic value of our credal states (EV-Aim). Formulated thus, EV-Aim does not take a stand on what makes one credal state more epistemically valuable than another. EV-Aim thus offers a schema that can be fleshed out in different ways, depending on how one conceives of epistemic value.

One possibility is that epistemic value reduces to accuracy. This would yield:

ACCURACY AIM The aim of inquiring into a question Q is to make one's credence in the answer to Q as accurate as possible.

On this view, Poirot attains the aim of inquiry when he comes to have credence 1 in the correct answer to the question, *Who committed the murder?*

While it is uncontroversial that accuracy is an important component of epistemic value (§3), it is contentious whether it is the only component. An alternative view is that credal states are epistemically valuable insofar as they have some property that entails accuracy, but which is itself distinct from accuracy. What might such a property be? Here I'll mention three possibilities.

In the traditional epistemology literature, a popular view holds that knowledge is subject to a modal condition, such as safety or sensitivity.¹² For example, safety theorists hold that in order for a belief B to amount to knowledge, B must be true at all sufficiently nearby worlds where it is held on a sufficiently similar basis. We might explore an analogous hypothesis about credal value: the epistemic value of a credal state depends not just on its accuracy at the actual world, but on its accuracy at nearby worlds. This would yield:

SAFETY AIM The aim of inquiring into Q is to attain a safe credence in the answer to Q —i.e., a credal state that is maximally accurate at all nearby worlds where it is held on the same basis.¹³

Rather than looking to modal conditions for inspiration, we might instead consult the virtue epistemological tradition. According to virtue epistemologists, knowledge is a special type of cognitive achievement. Perhaps the most well-developed version of this idea comes from the work of Ernest Sosa (2007; 2015) who argues that a belief amounts to knowledge only if it is *apt*, where a belief is apt just in case it is true in virtue of the exercise of a cognitive ability. This idea can also be extended to provide an account of credal value, yielding:

APTNESS AIM The aim of inquiring into Q is to attain an apt credence in the answer to Q —i.e., a credal state that is maximally accurate in virtue of the exercise of a cognitive ability.¹⁴

For a final option, we might turn to the idea that there is an epistemic status that is more demanding than knowledge: *epistemic certainty*. This idea played an important role in the medieval and early modern traditions. Aquinas, Scotus, and Descartes all took epistemic certainty to be the most exalted epistemic status, it was “perfect cognition.”¹⁵ Over the last century, the notion of epistemic certainty has fallen out of favor. However, in other work I have tried to argue that it is an important notion in its own right, and that it has important work to do in epistemological theorizing.¹⁶ This is not the place to rehearse those arguments. But suppose this idea is on the right track. Then we might identify the epistemic

¹²For the canonical defense of a sensitivity condition, see Nozick 1981. For defenses of safety, see Sosa 1999; Williamson 2000; Pritchard 2005, 2012; Beddor and Pavese 2020.

¹³For related discussion of how to apply a safety condition to credences, see Moss 2013; Beddor and Goldstein forthcoming.

¹⁴Cf. Konek 2016 for a sophisticated discussion of how to make sense of an aptness condition on credences.

¹⁵Most philosophers in this tradition focused on *scientia*, which was thought to require epistemic certainty. See Pasnau 2017 for an illuminating historical discussion.

¹⁶See Beddor 2020a,b. See also Stanley 2008 for an earlier defense of a certainty norm of assertion.

value of a credal state is the degree to which it is epistemically certain. Combined with EV-Aim, this would yield a new version of Descartes' idea that inquiry aims at epistemic certainty:

CARTESIAN AIM The aim of inquiring into a question Q is to attain epistemic certainty about the answer to Q .

Exactly what this involves will depend on one's theory of epistemic certainty. But here is a plausible necessary condition: in order for a credal state to be epistemically certain, it must be perfectly accurate. Given this assumption, anyone who attains the Cartesian Aim will thereby attain the Accuracy Aim.

We thus have a few natural candidates for what the aim of inquiry might be, if it is not knowledge.¹⁷ All of these proposals agree that inquiry aims, *inter alia*, at a perfectly accurate credal state, which in turn requires having credence 1 in the proposition at issue. So all of these proposals are inconsistent with a fallibilist version of K-Aim. At the same time, I think these proposals accommodate many of the intuitions and theoretical impulses that motivate K-Aim.

6.2 Knowing the answer and settling questions

Are there any other costs to replacing K-Aim with EV-Aim? In recent work, [Kelp 2011, 2014, 2018](#) has argued that the K-Aim can be motivated by considering the close connection between *knowing the answer* and *settling a question*. Here's a bare-bones version of Kelp's argument:

- P1) S attains the aim of inquiring into Q iff S settles Q .
- P2) S settles Q iff S comes to know the answer to Q .
- Ccl) S attains the aim of inquiring into Q iff S comes to know the answer to Q .

In response, I think we should grant P1). However, P2) deserves closer scrutiny. While it certainly has an air of plausibility, on reflection the very cases that cause trouble for K-Aim also cause trouble for P2). Consider Mia in **Murine Research**. She knows her hypothesis h . But has she settled the question of whether h is true? Intuitively, the answer is "No." After all, by her lights there is a genuine possibility that h is false. So the question of whether h is true remains open for her. But if it remains open for her, then she has not settled it.

¹⁷There are interesting questions about the entailment relations between these epistemic statuses. For example, some might think that epistemic certainty entails either safety or aptness (cf. [Beddor 2020a](#)). And on some views aptness itself entails safety; see e.g., [Carter 2016](#). One could also explore views on which epistemic value is some hybrid of these different statuses; see e.g., the hybrid modal-virtue epistemological views defended by [Pritchard 2012](#); [Kelp 2013](#).

What, then, does it take to settle a question? EV-Aim suggests a natural answer. One settles Q when one arrives at the epistemically optimal credal state *vis-à-vis* the answer to Q . Again, what this involves will depend on the details of one's theory of epistemic value. But on all the views that we canvassed above (§6.1), this will entail having credence 1 in the correct answer to Q . This suffices to explain the intuition that Mia has not settled the question of whether h is true, at least prior to reading the results of the new study.

6.3 Unattainable Aims?

A second concern is that by replacing K-Aim with EV-Aim, we've made the aim of inquiry unattainable. According to a familiar line of thought, we shouldn't be absolutely certain of anything. Well, maybe there are some exceptions: logical truths, the cogito, maybe some propositions about our current phenomenal states. But for the vast majority of questions regarding the external world—say, who committed the murder, or when the Roman Empire fell—it would never be rational to maximally confident of the answers.

In response, three points are worth noting. First, this objection cuts equally against the infallibilist option. After all, infallibilists hold that knowledge requires absolute certainty. If absolute certainty is seldom attainable, this would show that knowledge is seldom attainable.¹⁸ So attainability considerations will not help us decide between Fallibilism and K-Aim.

Second, we should question the assumption that it is never rational to have credence 1 in any external world proposition. In many contexts, we would be happy to assert that we are certain of many humdrum, contingent external world facts, e.g., "I'm certain Mia is in the office—I just saw her." Such assertions do not reek of irrationality. Moreover, none of our arguments against infallibilism (§5) show we can never be rationally certain of any external world claim. Perhaps, then, while rational certainty is harder to attain than knowledge, it is still frequently attainable.

But suppose we grant, at least for the sake of argument, that EV-Aim is seldom attainable. Would this show that the EV-Aim is false, or even implausible?

Some may think this answer is "Yes", on the grounds that inquiry would become a Sisyphean task: we will be rationally condemned to continue inquiring into questions that we have no hope of settling. Poor Poirot would never be entitled to announce that the butler did it, and move on to another case. But proponents of EV-Aim can—and should—deny their view has this consequence. Whether it is rational to continue pursuing some aim depends on the probability of coming closer to achieving that aim, together with the other possible aims

¹⁸See Unger 1975 for an infallibilist view that explicitly embraces this consequence.

available to you. When it comes to inquiry, these other possible aims will frequently take the form of other lines of inquiry one can pursue. In many cases, the expected epistemic value of inquiring into other questions will be greater than the expected epistemic value of persevering in one's current line of research.

For example, suppose Poirot is .98 confident that the butler did it. And suppose that it is extremely unlikely that he will encounter any further evidence that will significantly change this degree of confidence (he has interrogated all the witnesses, pursued all promising leads, etc.). Then it might well maximize overall expected epistemic value to move on to the next case.

Now, if Poirot has not fully attained the aim of inquiry into the original case, our view does entail his original investigation has not been fully settled. So if many years later he is given the opportunity to review evidence bearing on this case, it would be rational for him to accept. But this seems exactly right. It is often rational to re-open lines of inquiry that one has set aside when new, unexpected evidence comes to light. For example, in **Murine Research** it would still be rational for Mia to read the results of the new study, even if she had stopped researching h years ago.

7 Consequences

I've argued that we need to choose between Fallibilism and K-Aim. I've also given reason to think that rejecting Fallibilism is the costlier option. Consequently we should abandon K-Aim, replacing it with the idea that inquiry aims at achieving maximally valuable credal states.

My argument has important consequences for a number of related topics in epistemology. By way of conclusion, I'll briefly investigate some of these implications.

7.1 Interrogative attitudes

In an important series of papers, [Friedman 2013, 2017, 2019](#) draws attention to what she calls, "interrogative attitudes." These include attitudes such as *wondering whether p* , *being curious as to whether p* , and *deliberating whether p* . As Friedman notes, these attitudes seem to be intimately tied to inquiry. Someone who is inquiring into p is naturally described as wondering whether p ; normally, they will also be curious as to whether p .

Friedman also argues that these interrogative attitudes aim at knowledge: they are all "relieved" when the agent comes to know the answer to the relevant question ([Friedman 2013](#): 145). This idea has considerable pre-theoretic plausibil-

ity. After all, it's natural to describe someone who wonders whether p as "wanting to know" whether p is true.

However, I think our trouble cases for K-Aim show that this idea cannot be right. Take **Murine Research**. When Mia gets the email, it seems natural to describe her as wondering whether h is true. It likewise seems natural to describe her as curious as to whether h is true. But (assuming we retain Fallibilism) she already knows h .

In a way this upshot is unsurprising, given our earlier arguments. After all, suppose we accept Friedman's claim that one can rationally maintain an interrogative attitude towards Q as long as one is rationally inquiring into Q . And suppose we also accept my argument that it can be rational to inquire into Q even when one knows the answer to Q . It follows that in such cases it is rational for the agent to maintain the relevant interrogative attitudes to Q , despite knowing the answer. We can also arrive at this conclusion another way, by considering the connection between interrogative attitudes and "settling" a question. After all, it seems one can rationally be curious about p , and rationally wonder whether p , as long as one has not settled the question of whether p . As we noted in §6.2, our protagonists have not settled the questions under investigation, despite knowing the answers.

What is the aim of interrogative attitudes, if not knowledge? Our earlier discussion suggests a natural answer. Just like inquiry, these attitudes aim at the attainment of epistemically optimal credal states.

7.2 Inquiry and belief

Some authors have also suggested that there is a close normative connection between inquiry and belief. Here again, Friedman offers the most sustained defense of this connection (Friedman 2019). Friedman defends the following norm:

DON'T BELIEVE AND INQUIRE (DBI) One ought not to inquire into the question of whether p if one believes p (or one believes $\neg p$).

Our earlier examples also cause trouble for DBI. Presumably, knowledge entails belief. So Mia believes h , and Tess believes r , but it is still rational for them to inquire further into the truth of these propositions.

Here too, the upshot is not entirely negative. Rather than positing a close normative connection between inquiry and belief, we could instead posit a connection between inquiry and psychological certainty. Indeed, we could accept something much in the spirit of DBI, simply swapping out talk of "belief" with "certainty." This would deliver the correct results in our cases. If Tess were absolutely certain of r (and rationally so), it would be very strange for her to consult

the textbook with the aim of figuring out whether r is true. Moreover, the framework offered here explains why: doing so would do nothing to further the goal of maximizing expected epistemic value.¹⁹

7.3 The Norm of Practical Reasoning

The arguments developed in this paper also have implications for debates over the norm of practical reasoning. A number of epistemologists have advanced knowledge-action norms along the following lines:

KNOWLEDGE-ACTION NORM (KN) If S knows p , then S is permitted to take p for granted in practical reasoning.²⁰

Assuming Fallibilism is true, our cases from §2 provide counterexamples to KN. Take **Ancient History**. If Tess knows r , then by KN she is permitted to take r for granted in practical reasoning. And so she should be permitted to ignore any possibilities in which r is false. But if she is permitted to ignore all $\neg r$ possibilities, then there is no point checking the textbook. Thus the challenge developed in this paper raises more general doubts about whether there are any important connections between knowledge and practical rationality.²¹

Of course, I am not the first to propose counterexamples to KN. Tellingly, many of the counterexamples proposed to date involve agents whose knowledge

¹⁹Some philosophers might question whether this is a genuine alternative to DBI. Recently, a number of epistemologists have advanced the surprising view that belief entails psychological certainty (e.g., Clarke 2013; Greco 2015; Dodd 2017; Moss 2018, 2019). But the same considerations that cut against infallibilism about knowledge also cast doubt on this infallibilist conception of belief. Just as it would be natural for an uncertain test-taker to remark, “I think I know the answer to this one...”, it also would be natural for them to remark, “I believe the answer is...”. Similarly, if Poirot is asked about the progress of his investigation, we can well imagine him saying, “I believe that the butler is the culprit. But I’m not sure, so we’ll have to investigate further.” This suggests the ordinary conception of belief—the sort conveyed in our everyday “belief” talk—is compatible with the absence of certainty. So insofar as DBI is intended to invoke this ordinary conception of belief rather than some stipulative sense (a point that Friedman explicitly avows) our criticism still applies. For complementary evidence that ordinary “belief” talk picks out a weaker state than certainty, see Hawthorne et al. 2016; Beddor and Goldstein 2018; Holguín 2020.

²⁰For sympathetic discussion of a norm along these lines, see Hawthorne 2004; Hawthorne and Stanley 2008; Fantl and McGrath 2002, 2009; Weatherson 2012; Weisberg 2013; Ross and Schroeder 2014; Moss 2018. Note that while some of these authors take knowledge to be both necessary and sufficient for practical reliance, I am focusing on just the sufficiency claim.

²¹Here I assume that if S is permitted to take p for granted in practical reasoning, S is permitted to perform whatever action maximizes expected value conditional on p . Something like this principle is often implicitly assumed in the literature on the norms of practical reasoning, and it is explicitly endorsed by some leading defenders of KN; see Fantl and McGrath 2002, 2009; Weatherson 2012; Ross and Schroeder 2014.

falls shy of certainty. For example, [Brown 2008a,b](#) and [Reed 2010](#) both offer cases with the following structure: an agent knows p , but not with absolute certainty. They are given the opportunity to perform an action—e.g., providing an affirmative answer to the question, “Is p true?”—that will have a modest payoff if p is true, and disastrous consequences otherwise. According to both Brown and Reed, the agent is not permitted to perform the relevant action, contrary to what KN predicts. These other cases helpfully illustrate—and reinforce—the tension between Fallibilism and KN that I am drawing out.

At the same time, the present paper helps advance the critical discussion around knowledge-action norms in two ways. First, it blunts the main response to extant criticisms of KN, which is to insist that knowledge depends on pragmatic factors. According to this response, the agents in Reed and Brown’s cases lose their knowledge of p as soon as they confront a high stakes decision that hinges on whether p is true (e.g., [Fantl and McGrath 2009](#): 62-63). However, we have already seen that an impurist strategy along these lines will not resolve the tension between K-Aim and Fallibilism (§4.1). So even if impurism helps with Reed and Brown’s cases, it will not provide a sufficiently general defense of KN.²²

Second, and more importantly, our earlier discussion suggests a positive picture that avoids the difficulties facing KN. As before, the key idea would be to replace talk of knowledge with talk of epistemically optimal credal states:

OPTIMAL CREDENCE-ACTION NORM If S ’s credal state *vis-à-vis* p is maximally epistemically valuable, then S is permitted to take p for granted in practical reasoning.²³

A norm along these lines avoids the problems facing the KN. Since Tess is not completely certain of r , we cannot use Optimal Credence-Action Norm to derive the conclusion that she is permitted to take r for granted in practical reasoning. A similar solution also applies to Brown and Reed’s counterexamples to KN. In their cases, the agents are not certain of the propositions that they know, hence their credences *vis-à-vis* these propositions are not maximally epistemically valuable.

7.4 The value of knowledge (or absence thereof)

I’ll conclude by considering a final, big picture implication of the arguments developed here. As we noted in §1, some philosophers have been attracted to K-Aim

²²For other worries about this impurist defense of KN, see [Roeber 2018](#); [Beddor 2020b](#).

²³If we take the optimal credal states to be those which are epistemically certain (§6.1), then this is equivalent to a version of the epistemic certainty norm of practical reasoning defended in [Beddor 2020b](#).

because it offers to shed light on the value of knowledge. If we reject K-Aim, what should we say about the importance of knowledge?

An initial point: while I've been arguing that knowledge is not the aim of inquiry, I've been careful to leave open the possibility that the aim of inquiry entails knowledge. Whether this is so depends on one's theory of epistemic value. If one thinks epistemic value is entirely a matter of accuracy, then presumably one will reject even the necessity of knowledge for successful inquiry—after all, one can have credence 1 in a true proposition without knowing that proposition. But we also canvassed a number of other conceptions of epistemic value—for example, views on which maximal epistemic value is a matter of safe credences, or apt credences, or epistemic certainty. On these views, attaining maximal epistemic value may well entail knowledge.

That said, this point will probably come as scant consolation. After all, the key explanatory notion in EV-Aim is the notion of an epistemically optimal credal state. Even if this state entails knowledge, it is not at all clear that knowledge is part of what makes this state valuable.

For this reason, I think we should abandon the project of explaining the value of knowledge in terms of its connection with inquiry. This leaves us with two options. One is to try to find some other work for knowledge to perform. Perhaps even if knowledge does not play an important role in inquiry, it serves an important function in the explanation of action or assertion or what have you.²⁴ Another option is more radical: perhaps we should reject the idea that knowledge has any important explanatory work to do at all. Perhaps when we look into the matter, we will find that much of the explanatory work traditionally allocated to knowledge is better served by the notion of an epistemically valuable credal state.

References

- Charity Anderson and John Hawthorne. Knowledge, Practical Adequacy, and Stakes. In Gendler and Hawthorne, editors, *Oxford Studies in Epistemology*, pages 234–257. Oxford University Press, Oxford, 2019.
- D.M. Armstrong. Does Knowledge Entail Belief? *Proceedings of the Aristotelian Society*, 70:21–36, 1969.
- Bob Beddor. The Toxin and the Dogmatist. *Australasian Journal of Philosophy*, 97(4): 727–740, 2019.
- Bob Beddor. New Work for Certainty. *Philosophers' Imprint*, 20(8):1–25, 2020a.

²⁴Though note that my earlier arguments undercut some of these alternative applications of knowledge. For example, in §7.3 I argued that we should reject knowledge-action norms, such as KN. For those who hold that norm of action is also the norm of assertion, these arguments can also be used to cast doubt on a knowledge norm of assertion. For independent arguments that the norm of assertion is more demanding than knowledge, see [Stanley 2008](#); [Sosa 2009](#); [Beddor 2020a,b](#).

- Bob Beddor. Certainty in Action. *The Philosophical Quarterly*, 70(281):711–737, 2020b.
- Bob Beddor and Simon Goldstein. Believing Epistemic Contradictions. *Review of Symbolic Logic*, 11:87–114, 2018.
- Bob Beddor and Simon Goldstein. Mighty Knowledge. *The Journal of Philosophy*, forthcoming.
- Bob Beddor and Carlotta Pavese. Modal Virtue Epistemology. *Philosophy and Phenomenological Research*, 101(1):61–79, 2020.
- Jessica Brown. Subject-Sensitive Invariantism and the Knowledge Norm for Practical Reasoning. *Noûs*, 42(2):167–189, 2008a.
- Jessica Brown. Knowledge and Practical Reason. *Philosophy Compass*, 3(6):1135–1152, 2008b.
- Jessica Brown. *Fallibilism: Evidence and Knowledge*. Oxford University Press, Oxford, 2018.
- Catrin Campbell-Moore and Bernhard Salow. Avoiding Risk and Avoiding Evidence. *Australasian Journal of Philosophy*, Advance Publication Online, 2020. doi:10.1080/00048402.2019.1697305.
- J. Adam Carter. Robust virtue epistemology as anti-luck epistemology. *Pacific Philosophical Quarterly*, 97(1):140–155, 2016.
- Roger Clarke. Belief is Credence One (In Context). *Philosophers' Imprint*, 13(11):1–18, 2013.
- Stewart Cohen. How to be a Fallibilist. *Philosophical Perspectives*, 2:91–123, 1988.
- Keith DeRose. Assertion, Knowledge, and Context. *Philosophical Review*, 111(2):167–203, 2002.
- Dylan Dodd. Belief and Certainty. *Synthese*, 194(11):4597–4621, 2017.
- Jeremy Fantl and Matthew McGrath. Evidence, Pragmatics, and Justification. *Philosophical Review*, 111(1):67–94, 2002.
- Jeremy Fantl and Matthew McGrath. *Knowledge in an Uncertain World*. Oxford University Press, Oxford, 2009.
- Jane Friedman. Question-Directed Attitudes. *Philosophical Perspectives*, 27(1):145–174, 2013.
- Jane Friedman. Why Suspend Judging? *Noûs*, 51(2):302–326, 2017.
- Jane Friedman. Inquiry and Belief. *Noûs*, 53(2):296–315, 2019.
- Hilary Greaves and David Wallace. Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility. *Mind*, 115(459):607–632, 2006.
- Daniel Greco. How I Learned to Stop Worrying and Love Probability 1. *Philosophical Perspectives*, 29(1):179–201, 2015.
- John Hawthorne. *Knowledge and Lotteries*. Oxford University Press, New York, 2004.
- John Hawthorne and Jason Stanley. Knowledge and Action. *Journal of Philosophy*, 105(10):571–590, 2008.
- John Hawthorne, Daniel Rothschild, and Levi Spectre. Belief is Weak. *Philosophical Studies*, 173:1393–1404, 2016.
- Ben Holguín. Thinking, guessing, and believing. manuscript, 2020.
- James Joyce. A nonpragmatic vindication of probabilism. *Philosophy of Science*, 65(4):575–603, 1998.

- Klemens Kappel. On saying that someone knows: Themes from craig. In Haddock, Millar, and Pritchard, editors, *Social Epistemology*, pages 69–88. Oxford University Press, Oxford, 2010.
- Christoph Kelp. What’s the point of ‘knowledge’ anyway? *Episteme*, 8(1):53–66, 2011.
- Christoph Kelp. Knowledge: The safe-apt view. *Australasian Journal of Philosophy*, 91(2): 265–278, 2013.
- Christoph Kelp. Two for the knowledge goal of inquiry. *American Philosophical Quarterly*, 51:227–232, 2014.
- Christoph Kelp. Inquiry, Knowledge and Understanding. *Synthese*, 2018.
- Christoph Kelp. Theory of Inquiry. *Philosophy and Phenomenological Research*, forthcoming.
- Christoph Kelp and Mona Simion. Criticism and Blame in Action and Assertion. *Journal of Philosophy*, 114(2):76–93, 2017.
- Jason Konek. Probabilistic knowledge and cognitive ability. *Philosophical Review*, 125, 2016.
- Jonathan Kvanvig. The Value of Understanding. In Pritchard, Millar, and Haddock, editors, *Epistemic Value*, pages 95–111. Oxford University Press, Oxford, 2009.
- Aidan McGlynn. *Knowledge First?* Palgrave MacMillan, New York, 2014.
- Sarah Moss. Scoring rules and epistemic compromise. *Mind*, 120(480):1053–1069, 2011.
- Sarah Moss. Epistemology Formalized. *Philosophical Review*, 122(1):1–43, 2013.
- Sarah Moss. *Probabilistic Knowledge*. Oxford University Press, Oxford, 2018.
- Sarah Moss. Full Belief and Loose Speech. *Philosophical Review*, 128(3):255–291, 2019.
- Robert Nozick. *Philosophical Explanations*. Harvard University Press, Cambridge, 1981.
- Graham Oddie. Conditionalization, Cogency, and Cognitive Value. *British Journal for the Philosophy of Science*, 48:533–541, 1997.
- Robert Pasnau. *After Certainty: A History of Our Epistemic Ideals and Illusions*. Oxford University Press, Oxford, 2017.
- Richard Pettigrew. *Accuracy and the Laws of Credence*. Oxford University Press, Oxford, 2016.
- Duncan Pritchard. *Epistemic Luck*. Oxford University Press, New York, 2005.
- Duncan Pritchard. Anti-luck virtue epistemology. *Journal of Philosophy*, 109:247–249, 2012.
- Colin Radford. Knowledge: By Examples. *Analysis*, 27:1–11, 1966.
- Baron Reed. A Defense of Stable Invariantism. *Noûs*, 44(2):224–244, 2010.
- Baron Reed. Fallibilism, Epistemic Possibility, and Epistemic Agency. *Philosophical Issues*, 23:40–69, 2013.
- Blake Roeber. The Pragmatic Encroachment Debate. *Noûs*, 52(1):171–195, 2018.
- Jacob Ross and Mark Schroeder. Belief, Credence, and Pragmatic Encroachment. *Philosophy and Phenomenological Research*, 88(2):259–288, 2014.
- Gillian Russell and John Doris. Knowledge by Indifference. *Australasian Journal of Philosophy*, 86(3):429–437, 2008.
- Patrick Rysiew. Speaking of knowing. *Noûs*, 41(4):627–662, 2007.
- Patrick Rysiew. Epistemic scorekeeping. In Brown and Gerken, editors, *Knowledge Ascriptions*. Oxford University Press, New York, 2012.

- Miriam Schoenfield. Bridging rationality and accuracy. *Journal of Philosophy*, 112(12): 633–657, 2016.
- David Sosa. Dubious Assertions. *Philosophical Studies*, 146:269–272, 2009.
- Ernest Sosa. How to Defeat Opposition to Moore. *Philosophical Perspectives*, 13:137–149, 1999.
- Ernest Sosa. *Apt Belief and Reflective Knowledge*, volume 1. Oxford University Press, Oxford, 2007.
- Ernest Sosa. *Judgment and Agency*. Oxford University Press, Oxford, 2015.
- Jason Stanley. *Knowledge and Practical Interests*. Oxford University Press, New York, 2005a.
- Jason Stanley. Fallibilism and Concessive Knowledge Attributions. *Analysis*, 65(2):126–131, 2005b.
- Jason Stanley. Knowledge and Certainty. *Philosophical Issues*, 18(1):35–57, 2008.
- Jonathan Sutton. *Without Justification*. MIT Press, Cambridge, MA, 2007.
- Peter Unger. *Ignorance*. Clarendon Press, Oxford, 1975.
- Brian Weatherson. Knowledge, Bets, and Interests. In Brown and Gerken, editors, *Knowledge Ascriptions*, pages 75–103. Oxford University Press, New York, 2012.
- Jonathan Weisberg. Knowledge in Action. *Philosophers' Imprint*, 13(22):1–23, 2013.
- Timothy Williamson. *Knowledge and its Limits*. Oxford University Press, Oxford, 2000.
- Timothy Williamson. Justifications, Excuses, and Sceptical Scenarios. In Dorsch and Dutant, editors, *The New Evil Demon Problem*. Oxford University Press, Oxford, forthcoming.
- Alex Worsnip. Possibly false knowledge. *Journal of Philosophy*, 112(5):225–246, 2015.