INQUIRY FOR FALLIBILISTS

Abstract

Why engage in inquiry? Many philosophers answer this question in terms of knowledge: the goal of inquiring into some question is to come to know its answer. While this view holds considerable appeal, this paper argues that it stands in tension with another highly attractive thesis: fallibilism. Forced to choose between these two theses, I argue that the balance of considerations supports preserving fallibilism and rejecting the idea that inquiry aims at knowledge. I go on to articulate an alternative view, according to which inquiry aims at attaining epistemically optimal credal states. This alternative picture is fully compatible with fallibilism, and it coheres nicely with a rich body of work in epistemic decision theory. I proceed to highlight the implications of this replacement for several important topics in epistemology, including the dogmatism paradox, the nature of interrogative attitudes, the norm of practical reasoning, and the value of knowledge (or lack thereof).

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. How long must he continue? When is he entitled to conclude his inquiry and announce that the butler did it?

Until recently, this sort of question received scant attention from epistemologists. Mainstream epistemology has tended to focus on what sort of doxastic attitudes you should adopt, given the evidence you have; it has been less interested in the conditions under which you should seek out more evidence. But no longer: recent epistemology has taken a 'zetetic turn', giving central place to questions about the aims and norms governing inquiry.[1] In the recent flurry of work

[1] I borrow the label 'zetetic' from Friedman 2020, forthcoming, who has been one of the most prominent champions of the zetetic turn. See also Thorstad 2021.

on this topic, one can discern a consensus beginning to emerge: *inquiry aims at knowledge*. That is:

K-Aim: The aim of inquiring into a question Q is to come to know the answer to Q.[2]

K-Aim holds considerable appeal. It delivers a simple answer to our question about Poirot: he is entitled to conclude his inquiry when—and only when—he comes to know who committed the murder. K-Aim also offers to shed light on the value of knowledge. On the resulting picture, knowledge is important because it functions as an *inquiry-stopper*. Finally, K-Aim is bolstered by the idea that "interrogative attitudes" such as *curiosity* and *wondering* aim at knowledge (Friedman 2013, 2017; Sapir and van Elswyk 2021). After all, we typically engage in inquiry in order to satisfy our curiosity.

Despite its appeal, this paper argues that K-Aim stands in tension with another highly attractive principle: fallibilism. While different epistemologists formulate this principle in slightly different ways, here I'll focus one 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.[3]

Here's a quick preview of the argument. 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 acquire this evidence. But this conflicts with K-Aim, which says that our agent has already attained the aim of inquiry.

After developing the argument in more detail (§§2-3), I consider some strategies for reconciling K-Aim and Fallibilism and find them wanting (§4). So we are

[2]While different authors formulate this idea in slightly different ways, versions of K-Aim are defended in Sartwell 1992; Schaffer 2005; Kvanvig 2009; Kappel 2010; Kelp 2011, 2014, 2021a,b, forthcoming; Rysiew 2012; Whitcomb 2017; Sapir and van Elswyk 2021. Other philosophers posit a similar connection between inquiry and knowledge without presupposing that inquiry has any distinctive *telos*. For example, Friedman 2017 proposes a normative requirement, according to which one ought not to inquire into some question if one knows the answer to this question. My arguments in this paper generalize straightforwardly to this normative variant of K-Aim.

[3]Fallibilism—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 Rysiew 2007; Reed 2013; Worsnip 2015; Brown 2018, among many others.

forced to choose between the two principles. I argue that the balance of considerations tips in favor of Fallibilism and against K-Aim (§§5-6). 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 credences. This alternative picture is fully compatible with Fallibilism, and it fits naturally with a rich body of work in epistemic decision theory. I conclude by exploring the implications of this replacement for a range of related topics, including the the dogmatism paradox, the nature of interrogative attitudes, the norms of practical reasoning and belief, and the value of knowledge (§7).

2 Bringing out the tension: two cases

Start with a case of fallible knowledge:

> **Murine Research** Mia is a scientist who forms the hypothesis *m*: *A particular drug (Accuphine) causes hyperactivity in mice.* Mia conducts a number of experiments that support *m*. Eventually, she conducts enough experiments to know that *m* is true. But she still is not completely certain of *m*, 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.

If Fallibilism is true, cases of this sort are presumably 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. Their email announces 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 ought to read the results of the latest study. After all, she is not absolutely certain of *m*, and the results of the latest study might change her mind one way or the other. Now, there are ways of filling in the details of the case that might erode this intuition—for example, if reading the study would take valuable time away from Mia's other research projects. For our purposes, we can stipulate that this is not the case: Mia is deciding between spending the next hour reading the study or playing Tetris. Once this stipulation is in place, it is hard to deny that Mia is under some rational pressure to read the results.

However, this intuition stands in tension with K-Aim. After all, reading the results constitutes further inquiry into the question: *Is m 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 various 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 case, or cases of inductive knowledge more generally. It is 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*.[4] 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." 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.

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.

[4] This case bears some resemblance to Radford's 'unconfident examinee' (Radford 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.

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 credences.[5]

What makes a particular credence 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's credence in p is not maximally accurate, than A's credence in p is not maximally epistemically valuable.

This idea should be appealing to anyone who subscribes to the "veritist" idea that the ultimate epistemic goods are the attainment truth and avoidance of 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 does it mean for a credence to be accurate? Intuitively, the accuracy of your credences is a matter of their "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 A's credence in a true proposition p is higher than B's, then A's credence in p is more accurate than B's.

There are various formal measures of accuracy that respect Alethic Proximity. 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 bother? 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.

[5] For 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.

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, her credence in *r* is not maximally accurate. By Valuable Accuracy, her credence in this proposition is not maximally valuable. By EV-Aim, she has not attained the aim of inquiring into the question, *When did the Roman Empire fall?*. (Similar points can be made, *mutatis mutandis*, using **Murine Research**.)

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

Aim-Rationality Bridge If you rationally ought to pursue aim α, and you rationally ought to believe that you have not yet attained α, then you rationally ought to continue pursuing α.

Mia and Tess are aware that they are not certain of the answers to the relevant questions. So they rationally ought to believe that they have not yet attained maximally valuable credences with regards to the question at hand. By EV-Aim and Aim-Rationality Bridge, they rationally ought to continue their inquiries.

We could also arrive at the same result using a different route. In an important contribution to 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.[6] 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.[7] To illustrate, consider again **Ancient History**. Since Tess is not certain of *r*, there's a 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

[6] Oddie's theorem builds on Good 1967, which showed that, given certain assumptions, it always maximizes expected utility to gather and conditionalize on cost-free evidence. Whereas Good focuses on practical value, Oddie extends this result to epistemic value.

[7] One important assumption in Oddie's argument is that our measure of epistemic value is strictly proper: the expected epistemic value of a credence function c, when calculated using c itself, must exceed the expected epistemic value of any other credence function. For discussion, see Lewis 1971; Oddie 1997; Greaves and Wallace 2006. While strict propriety suffices for Oddie's theorem to hold, it is not necessary. See, for example, the measure of epistemic value developed in Horwich 1982.

case she will raise her credence in *r*. 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. 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.

Now, suppose we appeal to a plausible principle linking up expected epistemic value with one's rational requirements:

Rationality Requires Maximizing Expected Value An agent is epistemically required to do whatever maximizes expected epistemic value.

Combining this principle with Oddie's theorem gives us another route to the conclusion that Tess is rationally required to consult her textbook, and that Mia is required to look at the results of the new study.[8]

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 presuppose the decision theoretic machinery introduced in this section. So 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.[9]

[8]The assumption that it is always rational to maximize expected value (epistemic or otherwise), has been questioned, on the grounds that it can be rational to be risk averse (Buchak 2010; Campbell-Moore and Salow 2020). But for our purposes we need not take a stand on this issue. It's enough that rationality requires maximizing expected value for risk neutral agents, and we can stipulate that Mia and Tess are risk neutral.

[9]My defense of EV-Aim bears some affinities to a view independently developed by Falbo forthcoming, according to which inquiry aims at 'epistemic improvement'. While in many respects Falbo's view is congenial to the picture put forward here, there are important differences. One major difference is that I aim to develop a systematic framework for thinking about epistemic value—a framework that ties this notion to credal accuracy specifically, and to epistemic decision theory more generally. By contrast, Falbo expresses skepticism about whether there is anything systematic to say about the relevant types of improvement guiding inquiry. A more subtle difference concerns cases where an agent inquires into a question even though they have no chance of gaining either knowledge or certainty. On Falbo's view, even if the inquiry only results in a slight epistemic improvement, they could well have attained the aim of inquiry. But, intuitively, their inquiry is importantly complete. On the view developed here, this agent has not attained the aim of inquiry, since their resulting credence is epistemically suboptimal. Still, it could be rational for them to inquire, since doing so will maximize expected epistemic value.

4 Reconciliation strategies

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

4.1 Practical vs. epistemic requirements

An initially tempting thought is that we need to distinguish between practical and epistemic oughts. According to this response, it can be practically rational pursue some inquiry, even when it is not epistemically rational to do so. For example, suppose an eccentric billionaire offers you a million dollars to inquire into some proposition that you already know. If you need the million, then you ought to inquire. But this case is hardly a refutation of K-Aim; after all, K-Aim is supposed to articulate a distinctly epistemic aim for inquirers. Perhaps, some might suggest, we should say something similar about Mia and Tess. Granted, they ought to inquire further into the matters at hand, but their reasons for inquiring are practical, not epistemic.

But do our protagonists lack any epistemic reason for inquiring further? This is doubtful. Take Mia. She may have practical reasons for reading the results of the researcher's study (perhaps doing so will help her get her next paper published). But she also has a further reason: she wants to find out whether m is true! This seems like an epistemic reason *par excellence*.

The epistemic decision theory introduced in the previous section offers one way of developing this intuition. We started with the premise that accuracy is epistemically valuable (Valuable Accuracy), together with the premise epistemic value of a credence is its proximity to the truth (Alethic Proximity). As we saw, this was enough to show that Mia and Tess have not yet attained maximally epistemically valuable credences on the questions at hand. Consequently, insofar as they want to promote the epistemic value of their credences, they have reason to inquire further. Surely this qualifies as an epistemic goal—one that grounds an epistemic, not merely practical, obligation.[10]

[10]These points also serve to defuse a related response to our cases. According to this response, Mia has attained the aim of inquiry on the question of whether m is true, but she has not attained the aim of inquiry on the question, *What were the results of the researcher's study?* This response does not do justice to the intuitive thought that Mia's main reason for investigating the latter question is that she has not completely settled the former. This intuitive thought is codified by our epistemic decision theoretic framework: she ought to read the results of the researcher's study in order to promote the accuracy–and hence the epistemic value—of her credence in m.

4.2 Impurism

A second reconciliation strategy also appeals to practical considerations, but in a different way. According to 'impurism', knowledge depends on practical factors.[11] While there are different ways of using impurism to try to reconcile our principles, the basic strategy 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 considerations rob our protagonists of their knowledge. So they have not achieved the aim of inquiry 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 stakes-based impurism is ill-equipped to resolve our tension; after all, we can stipulate that the 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 ought to inquire remains.

Could a different form of impurism fare better? 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 A knows p only if p is practically adequate in A's situation.

This condition entails that when our agents are given the opportunity to receive cost-free evidence, they lose their knowledge. Take **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 not practically adequate for Tess. By the Practical Adequacy Condition, she ceases to know 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.

[11] AKA 'pragmatic encroachment.' For defenses, see e.g., Fantl and McGrath 2002, 2009; Stanley 2005a; Weatherson 2012; Ross and Schroeder 2014.

A second difficulty for the Practical Adequacy Condition comes from Anderson and Hawthorne 2019, who note that it generates an unwelcome epistemic instability. Imagine that as Tess starts to flip to the relevant page in her textbook, 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, Tess 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 have epistemically suboptimal credences. Nothing in this explanation presupposed a purist theory of knowledge. Our explanation only relied on independently plausible tenets of epistemic decision theory (Valuable Accuracy, Alethic Proximity, EV-Aim). Proponents of the impurist reconcilation are put in the uncomfortable position of being forced to reject one of these tenets.

4.3 Derivative norms

In recent years, one major strategy for defending knowledge norms of assertion, action, and belief from ounterexamples is to appeal to various derivative norms. Could this strategy reconcile K-Aim and Fallibilism?

Let's examine how one might go about developing this strategy. Start with the idea that the primary norm for some activity governed by aim α is to pursue that activity until you have achieved α. Combined with K-Aim, this yields:

Primary Knowledge Norm If you ought to inquire into a question Q, then you ought to continue inquiring until you come to know the answer to Q.[12]

Mia and Tess satisfy the Primary Knowledge Norm. But, according to the response under consideration, there are derivative norms of inquiry that they do not satisfy. These derivative norms explain why they ought to gather more evidence.

What are these derivative norms? One option is to go dispositional (Williamson forthcoming; cf. Lasonen-Aarnio 2010, forthcoming). For example, Williamson

[12] Another approach would be to formulate the primary norm as the claim that if you know the answer to Q, then you ought not inquire into Q (Friedman 2017; Sapir and van Elswyk 2021). My arguments here can be adapted to target this alternative formulation.

forthcoming proposes that for any norm N, there is a derivative norm: *Have a general disposition to satisfy N* (the 'Dispositional Norm'). And this in turn yields a further derivative norm: *Act as someone who complies with the Dispositional Norm would do in this situation* (the 'Occurrent Dispositional Norm'). To illustrate with Williamson's example, suppose A promise to give B money, while secretly intending to give the cash to C instead. However, B disguises themselves as C, thereby tricking A into giving B the money. A fulfills the primary norm of promising (A kept their promise), but their behavior is nonetheless subject to reproach. According to Williamson, this is because A violated the Dispositional Norm and the Occurrent Dispositional Norm: A did not act as a promise-keeper would typically act in this situation.

Combining Williamson's proposal with the Primary Knowledge Norm yields:

Occurrent Dispositional Knowledge Norm You ought to act in a given situation as someone who is disposed to comply with the Primary Knowledge Norm would act in this situation.

However, this derivative norm does not explain why Mia and Tess ought to continue their inquiries. Meet Minny: Minny is reliably disposed to do the bare minimum, from the the point of view of the Primary Knowledge Norm. Whenever Minny inquires into some question, she continues inquiring until she knows its answer. But, having attained such knowledge, she stops investigating. Indeed, if you offer to give her further evidence bearing on the question, she will decline. By stipulation, Minny is disposed comply with the Primary Knowledge Norm. Now, if Mia ignored the researcher's email, or Tess refused to consult the textbook, they would be acting just as Minny would act in their situation.

Would some other derivative knowledge norm fare better? In Williamson's promising example, even though A complies with the primary norm on promising, A justifiably believes they are violating this norm. Perhaps this is why A's action is blameworthy. A diagnosis along these lines is developed by DeRose (2002: 180, 2009: 93), who suggests that for any primary norm N, there is a secondary norm that forbids one from performing actions that one reasonably thinks violate N. DeRose's suggestion seems very plausible; indeed, we relied on a variant of this idea in §3 (Aim-Rationality Bridge). Combining DeRose's suggestion with the Primary Knowledge Norm yields:

Reasonable Belief Knowledge Norm If you ought to inquire into Q, then you ought not cease your inquiry if you have reason to believe that you do not know the answer to Q.

But this norm also fails to explain our cases. Mia and Tess do not have any reason to think that they do not know the answers to the questions under investigation. As we noted above, it would be natural for Tess to think to herself, 'I believe I know the answer to this one. But I'm not absolutely certain, so I might as well check.' So the Reasonable Belief Knowledge Norm does not explain why they ought to inquire further.

Similar points apply if we strengthen the derivative norm. In the promising example, perhaps the fault lies in the fact that A does not *know* that they are fulfilling their promise. More generally, we might hypothesize that for any norm N, there is a derivative norm forbidding you from performing actions if you do not know that those actions comply with N (cf. Benton 2013). Combined with the Primary Knowledge Norm, this yields:

Second Order Knowledge Norm If you ought to inquire into Q, then you ought to continue inquiring until you know that you know the answer to Q.

If one can know p without being rationally certain of p, presumably one can also know that one knows p without being rationally certain of p. Indeed, most theories that make room for fallible knowledge also make room for fallible second order knowledge.[13] For a toy example, consider the view knowledge is just safe belief, and that belief only requires a credence above .9. Now imagine that your credence in p is .95, and your credence that you safely believe p is also .95. If the latter belief is safe, then you know that you know p, even though you are less than certain of p. So we should be able to cook up versions of our cases where Mia and Tess know that that they have attained knowledge on the question at hand. In such cases, they will satisfy the Second Order Knowledge Norm. Still, if they are not certain of the answers, it may be rational for them to continue inquiring.

This point generalizes to further strengthenings of the derivative norm. Above, we floated the idea that for any activity governed by norm N, there is a derivative norm according to which one ought to know whether one is complying with N. Now, we might apply this idea to the Second Order Knowledge Norm, yielding a Third Order Knowledge Norm, and so on (cf. Goldstein 2022). Continued applications would yield an omega knowledge norm:

[13] According to the KK thesis, whenever one knows p, one knows that one knows p. If we accept the KK thesis, then it is trivially true that fallible knowledge entails fallible second order knowledge. For present purposes, I will set this view aside, since it will clearly be unwelcome to those who seek to defend K-Aim by appealing to a higher order knowledge norm. That said, one advantage of replacing K-Aim with EV-Aim is that it allows us to remain neutral on the viability of the KK thesis. For criticism of the KK thesis, see Williamson 2000, chp.5; for defenses, see e.g., Greco 2014, 2015b; Goodman and Salow 2018; Dorst 2019.

Omega Knowledge Norm If you ought to inquire into Q, then you ought to continue inquiring until you know that know that you know (*ad infinitum*) the answer to Q.

Even this norm is not enough. If one can know that one knows p without having credence 1 in p, there is no principled reason to think that some magical number of further iterations of knowledge will all of a sudden entail credence 1 in p. This suggests that someone could have omega knowledge of p without being rationally certain of p. In such cases, the Omega Knowledge Norm cannot explain why they are under rational pressure to consult cost-free evidence bearing on p.[14]

Let's take stock. While the strategy of appealing to derivative norms holds initial appeal, the devil is in the details: we found no way of spelling out the derivative norm that handles the full range of cases. Here too, the epistemic decision theoretic framework from §3 offers a useful diagnosis of the source of the trouble. According to the diagnosis offered there, the the reason why Mia and Tess ought to continue their inquiries is that inquiry aims at maximizing epistemic value, and any credence shy of 1 is guaranteed to fall short of the maximal epistemic value. But if knowledge does not require certainty, it is hard to see why satisfying any norms that derive from the K-Aim would require certainty.

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.

[14]To make this a bit more concrete, consider cases where researchers have found it fruitful to postulate omega knowledge. Game theorists, linguists, and computer scientists frequently posit common knowledge, which entails omega knowledge. (A proposition is common knowledge if everyone in the group knows that everyone in the group knows it, and everyone in the group knows that everyone knows that everyone knows it, and so on.) One application of common knowledge is to solve coordination problems (Lewis 1969). For example, if we get lost on the ski slopes, we want to meet at the same location. For any particular location, I will only go there if I know that you will go there, but I also know that you know this, and that you will only go there if you know that I will go there, etc. A simple way of solving such a problem is to make a public announcement: 'I will meet you at the base lodge.' This proposition now becomes common knowledge, and hence is something that you omega know. Nothing in this set-up requires that you have credence 1 in my announcement. You might assign some slight credence to the possibility that I will get confused, or that you misheard me.

Of course, not all philosophers agree that a public announcement suffices for omega knowledge. On some views, omega knowledge is hardly ever attainable, because each order of knowledge requires an additional margin for error. (See Williamson 2000; for defenses of the possibility of omega knowledge, see the defenses of the KK thesis in fn.13, as well as Goldstein 2022, forthcoming.) For our purposes, we can remain neutral on this question. The important dialectical point is that if omega knowledge is attainable, it does not require completely certainty.

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. Once we give up Fallibilism, we can no longer derive any inconsistency between K-Aim and EV-Aim. After all, suppose A knows p. By infallibilism, A is rationally required to have credence 1 in p. Since p is true (by the factivity of knowledge), A's credence in p is guaranteed to be the most accurate possible.

The combination of K-Aim with infallibilism is clearly a coherent package. But is it plausible?

Consider how this position handles our cases from §2. Infallibilism entails our cases are impossible, at least as described. Take **Ancient History**. According to infallibilists, as long as Tess rationally assigns some positive credence—however slight—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 as a difficult consequence to stomach. As we noted in §4, we can imagine Tess remarking to herself: 'I'm pretty sure I know when the Roman Empire fell: 476. But I'm not completely sure that's when it fell, so I should probably check.' Such a thought seems perfectly natural. If infallibilism is true, then it is hard to make sense of Tess' reasoning here: Tess' lack of certainty automatically entails that she does not know when the Roman Empire fell. Why, then, does she (falsely) think that she knows this?

The problem for infallibilism is not due to any idiosyncratic features of this particular case. Conjunctions of the form, 'I believe I know p, but I'm not completely certain of p' are typically felicitous. For example, we could equally well imagine Mia remarking: 'I believe I know what effect Accuphine has on mice: it causes hyperactivity. But I'm not completely certain of this; it's possible that further data will prove me wrong.' If infallibilism is correct, any such conjunction commits the speaker to a false belief: either the speaker's belief that they know p is mistaken, or they are mistaken about their lack of certainty in p.[15]

[15]Here I am building on Worsnip 2015, which makes a closely related argument on behalf of

We can press a related concern from the third person perspective. Imagine that Tess is not given the opportunity to check her textbook. As Tess fills out the test, her teacher is hovering over her shoulder, and remarks to herself: 'Looks like Tess knew the answer to that one.' This remark seems perfectly natural. But it is hard to see how this knowledge ascription is warranted if infallibilism is true; Tess' teacher has no grounds for thinking that Tess has credence 1 in the answer.[16]

Here too, the problem is perfectly general. If infallibilism is true, we should be very hesitant to ascribe knowledge to other people: I should only be willing to say you know something if I am warranted in thinking that you are completely certain of it. (Indeed, if knowledge is the norm of assertion, then I should only be willing to say you know something if I am completely certain that you are completely certain of it.) But as a matter of fact, we do not exhibit any such hesitation. As we often emphasize to undergraduates who are enamored with the skeptic's argument, ordinary discourse is rife with knowledge ascriptions; 'knows' is one of the most commonly used verbs in English (Nagel 2014).

A further strike against infallibilism comes from the fact that we frequently talk about knowing something *with certainty*. This expression is not redundant. If I say that the gardener knows for certain where the butler was on the night of the murder, then I make a stronger claim if I merely say that gardener knows where the butler was on the night of the murder. This is not just an idiomatic feature of English; as Beddor 2020 documents, a wide variety of languages carve out a distinction between merely knowing something and knowing something with certainty. But for infallibilists, it is hard to make sense of this distinction: according to infallibilism, knowledge entails complete certainty, so anytime someone knows something, they know it for certain.

As a further data point, a quick Google search will unearth naturally occurring attributions of fallible knowledge. We often claim that someone knows something with 'near certainty', implying that the knowledge is not held with complete certainty. Some examples:

---

fallibilism. Worsnip focuses on constructions involving epistemic modals, in particular, 'I think I know p, but I might be mistaken.' I have presented my argument directly in terms of certainty denials, so as to avoid controversies about the semantics of epistemic modals. For further discussion, see Marushak 2021.

[16]Here 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 2020). As Goodman and Holguín forthcoming note, this intuition can be supported on theoretical grounds. It is very natural to judge that the unconfident examinee remembers the answer to the question. If remembering entails knowing, as Williamson 2000 argues, we are led to the conclusion that the examinee knows the answer.

> "We now know with near-certainty that Wall Street execs committed felonies."[17]
>
> "In less than half an hour, the doctor knows with near certainty which influenza virus—if any—is present in the patient's respiratory tract."[18]
>
> "[T]he Clippers already locked up home court in the first round of the playoffs and know with near certainty they'll be the No. 4 seed in the Western Conference."[19]
>
> "When business resumes, we know with near-certainty that it will likely be a deluge."[20]

Examples along these lines sound perfectly coherent, and can be readily multiplied. But if infallibilism is true, all of these ascriptions are necessarily false: *near* certainty precludes knowledge.

I have laid out four problems for infallibilism. To recap: infallibilism has trouble making sense of our first person reasoning about knowledge, in particular, our willingness to say things like, 'I think I know p, but I'm not completely certain of p.' It also has trouble making sense of our third person knowledge ascriptions, in particular, our willingness to attribute knowledge to others without being sure whether they credence 1 in the relevant propositions. It is unable to explain the contrast between 'knows' and 'knows for certain', and it predicts that our ascriptions of knowledge with 'near certainty' are necessarily false. Now, I do not pretend that these difficulties will fully persuade committed infallibilists. Infallibilists may plead semantic blindness; perhaps ordinary speakers do not realize that knowledge requires credence 1, hence their (mistaken) tendency to think and talk in fallibilist terms. Still, I think most philosophers will admit that positing widespread semantic blindness is a last resort, or at the very least a theoretical cost. This cost provides impetus for exploring the other response to our tension: rejecting K-Aim.[21]

[17] https://www.businessinsider.com/my-plan-to-finally-make-wall-street-pay-for-its-crimes-2012-3

[18] https://www.nature.com/articles/d41586-019-02754-7

[19] https://www.nba.com/clippers/news/five-things-watch-five-games-left

[20] https://www.nuix.com/blog/forming-good-remote-work-habits-law-firms

[21] Infallibilists may concede that their view faces costs, but insist that these costs need to be weighed against the explanatory benefits. For example, Dodd 2011 argues that infallibilism provides the best explanation of the infelicity of concessive knowledge attributions ('I know p, but it might be that $\neg p$'). While a full discussion of this issue is outside the scope of this paper, a few alternative explanations of the infelicity of concessive knowledge have been proposed in recent years, each of which is both independently motivated and compatible with fallibilism. See Stanley 2005b; Worsnip 2015; Beddor 2020.

6 Rejecting K-Aim

6.1 If not knowledge, then what?

What is the aim of inquiry, if not knowledge? We have already sketched the beginning of an answer: inquiry aims at maximizing the epistemic value of our credences (EV-Aim). This answer can be developed in different ways, depending on how one conceives of epistemic value.

One possibility is that epistemic value reduces to accuracy. This gives us a simple way of fleshing out EV-Aim: the aim of inquiry is to make your credences as accurate as possible.

Another possibility is that credences 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? I'll briefly sketch three options.

In the traditional epistemology literature, one prominent view maintains that knowledge is subject to a modal condition, such as safety or sensitivity.[22] 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 credence depends not just on its accuracy at the actual world, but on its accuracy at nearby worlds. On this view, the ideal credence is not only maximally accurate, it is also maximally safe—that is, maximally accurate at all nearby worlds where it is held on the same basis.[23]

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*—that is, true in virtue of the exercise of a cognitive ability. This idea could be extended to provide an account of credal value. On the resulting view, the ideal credence is maximally apt—that is, maximally accurate in virtue of the exercise of a cognitive ability.[24]

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. For philosophers such as Aquinas, Scotus, and Descartes, epistemic certainty—or *scientia*—was a particularly exalted epistemic status; it was "perfect cognition."[25] The idea that there

[22] For the canonical defense of a sensitivity condition, see Nozick 1981. For defenses of safety, see Sosa 1999; Williamson 2000; Pritchard 2005, 2012, among others.

[23] For discussion of modal conditions on credences, see Moss 2018; Beddor and Goldstein 2021.

[24] See Konek 2016 for a defense of an aptness condition on credences.

[25] See Pasnau 2017 for relevant historical discussion. See Beddor 2020 for a recent attempt to

is an epistemic status with this profile fits quite naturally with our arguments for Fallibilism in §5. If knowing p does not always make it rational to be subjectively certain of p, then presumably there is some epistemic status stronger than knowledge that does warrant subjective certainty. Epistemologists looking to develop an alternative to K-Aim might consider dusting off the notion of epistemic certainty and putting it to work in a theory of epistemic value. On the resulting view, the epistemic value of your credence in p is the degree to which p is epistemically certain for you. Assuming that maximal epistemic certainty entails maximal accuracy, this provides another view on which epistemic value entails accuracy, but is not simply reducible to accuracy.

We thus have a few natural candidates for what the aim of inquiry might be, if it is not knowledge. These candidates are not necessarily rivals. For example, one might maintain that maximal epistemic certainty entails maximal safety or maximal aptness (or both).[26] For our purposes, we need not choose between these proposals. The important point is that 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 true answer to the question under investigation. So all of these proposals are inconsistent with a fallibilist version of K-Aim. At the same time, these proposals accommodate many of the intuitions and theoretical impulses that rendered K-Aim *prima facie* attractive.

6.2 Unattainable aims?

If we replace K-Aim with EV-Aim, have we made the aim of inquiry unattainable? According to many philosophers, 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 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. According to infallibilists, knowledge requires absolute certainty. If absolute certainty is seldom attainable, this would show that knowledge is seldom attainable. So attainability considerations will not help us decide between Fallibilism and K-Aim.

That said, we should question the assumption that it is never rational to have credence 1 in any external world proposition. People are often happy to claim that

rehabilitate the notion of epistemic certainty and assign it a starring role in epistemology.

[26] 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.

they are certain—indeed, absolutely certain—of all sorts of humdrum, contingent facts about the external world. Some examples 'from the wild':

> "Dr. Anthony Fauci said he is 'absolutely certain' the Omicron coronavirus variant will become the dominant variant in the US soon."[27]
>
> "Scientists are absolutely certain that this warming trend is due to human activity."[28]
>
> "Hunter Biden says he is '100 percent certain' he will be cleared of wrongdoing in tax investigation."[29]
>
> "Mayor says scientists 'are absolutely certain that a lot of these planets have good conditions for life."'[30]

Now, some might dispute whether these ordinary ascriptions of absolute certainty entail credence 1. However, note that it sounds contradictory to follow up an ascription of certainty in p with an admission that the relevant agent does not assign credence 1 to p. Just try it: 'I'm absolutely certain that Mia is in her office, but there's some possibility that she's elsewhere'; 'He's 100 percent certain he'll be cleared of wrongdoing, but he has some credence that he won't be'—these sound terrible. Alternatively, some might insist that our everyday ascriptions of complete certainty are just hyperbole: literally false but rhetorically effective. However, most epistemologists would be reluctant to take this stance with regards to our ordinary knowledge ascriptions. The considerations of charity that count against knowledge skepticism also count against certainty skepticism.

Perhaps, then, while rational certainty is harder to attain than knowledge, it is still frequently attainable.[31] But suppose we grant, 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 will 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

[27] https://www.cnn.com/us/live-news/omicron-covid-19-variant-12-16-21/index.html

[28] https://www.brookings.edu/wp-content/uploads/2019/09/20190920-global-response-to-the-climate-crisis.pdf

[29] https://thehill.com/homenews/sunday-talk-shows/546384-hunter-biden-says-he-is-100-certain-he-will-be-cleared-of/

[30] https://www.wtnh.com/news/the-latest-3-scientists-win-nobel-prize-in-physics/amp/

[31] For defenses of the idea that rational certainty is comparatively abundant, see Miller 1978; Klein 1981; Beddor 2020; Goodman and Holguín forthcoming, among others.

is rational to continue pursuing some aim depends on the probability of coming closer to achieving that aim, together with the other possible aims 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 m years ago.

7 Why it matters

I've argued that we should replace K-Aim with EV-Aim. This replacement has important implications for a number of topics in epistemology.

7.1 The dogmatism paradox

First up: the Kripke-Harman dogmatism paradox.[32] Suppose that when Mia receives the email, she reasons as follows: 'I know m is true. If I read the results of the latest study, I might find corroborating evidence that m is true, in which case I will retain my knowledge. But I might encounter evidence that m is false, which may defeat my knowledge. The safest course, then, is to delete the email!' Mia's reasoning here seems absurd. But where lies her mistake?

The framework developed in this paper provides a diagnosis. Mia is correct that deleting the email may protect her knowledge of m. But knowledge is not the epistemic *summum bonum*. As long as Mia is less than certain of m, her credence

[32]The paradox was first formulated by Kripke in a 1972 lecture to the Moral Sciences Club at Cambridge. For discussion, see Harman 1973; Kripke 2011; Lasonen-Aarnio 2014; Fraser forthcoming; [citation removed for review]; among others.

in m is epistemically suboptimal. By deleting the email, she is consigning herself to remain in this suboptimal state. What if she reads? Well, she is not guaranteed to maximize epistemic value (she is right that the results of the study may be misleading). But she is guaranteed to maximize expected epistemic value (recall the Oddie's theorem, §3). Hence the rational course is to read. This diagnosis of the dogmatist's error is a straightforward consequence of the positive proposal advanced here. By contrast, proponents of K-Aim face a real challenge here; for them, it is far less clear where the dogmatist's reasoning goes wrong.

7.2 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.[33] This idea has considerable pre-theoretic plausibility. After all, it's natural to describe someone who wonders whether p as 'wanting to know' whether p is true.

However, our cases from §2 put pressure on this view. Take **Murine Research**. When Mia gets the email, it seems natural to describe her as wondering whether m is true. It likewise seems natural to describe her as curious as to whether m is true. Indeed, it is precisely because she has these interrogative attitudes that it is rational for her to read the results of the latest study.

What is the aim of interrogative attitudes, if not knowledge? Our earlier discussion suggests a natural answer. Just like inquiry, these attitudes aim at attaining epistemic ideal credences.

7.3 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 $\neg p$).

[33] Friedman 2013: 145. Sapir and van Elswyk 2021 also defend a connection between interrogative attitudes and knowledge.

Our examples also cause trouble for DBI. Mia believes m, and Tess believes r, but it is still rational for them to inquire further.

Here too, the upshot is not entirely negative. We could accept something much in the spirit of DBI, provided we replace *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. If Tess were absolutely certain of r, then by her lights consulting the textbook will do nothing to improve the epistemic value of her credence in r.[34]

7.4 The norm of practical reasoning

The arguments developed in this paper also have implications for debates over the norm of practical reasoning. Over the last couple decades, many epistemologists have advanced knowledge-action norms along the following lines:

Knowledge-Action Norm (KN) If A knows p, then A is permitted to take p for granted in practical reasoning.[35]

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.

Of course, I am not the first to propose counterexamples to KN. Tellingly, many of the counterexamples proposed to date involve agents whose knowledge falls shy of certainty. For example, Brown 2008 and Reed 2010 offer cases with

[34]Some philosophers might question whether this proposal is a genuine alternative to DBI. According to some philosophers, belief entails certainty (e.g., Levi 1991; Clarke 2013; Greco 2015a; Dodd 2017; Moss 2019). But the same considerations that cut against infallibilism about knowledge also cast doubt on this infallibilist conception of belief. The fact that it would be natural for Tess to remark, 'I believe I know the answer to this one…' suggests that the ordinary conception of belief does not require 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; Rothschild 2020; Goodman and Holguín forthcoming. For complementary criticisms of DBI, see Falbo 2021; Sapir and van Elswyk 2021.

[35]For 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.

the following structure: an agent knows p, but not with absolute certainty. This agent is 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 go impurist. 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 (Fantl and McGrath 2009: 62-63). However, we have already seen that impurism fails to satisfactorily resolve the tension between K-Aim and Fallibilism (§4.2). So even if impurism helps with Reed and Brown's cases, it will not provide a sufficiently general defense of KN.

More importantly, our discussion suggests a positive picture that avoids the difficulties facing KN. As before, the key is to shift from knowledge to epistemically ideal credences:

Ideal Credence-Action Norm If A's credence in p is maximally epistemically valuable, then A is permitted to take p for granted in practical reasoning.

A norm along these lines avoids the problems facing KN. Since Tess is not completely certain of *r*, we cannot use the Ideal 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 are not maximally epistemically valuable.

7.5 The value of knowledge (or absence thereof)

I'll conclude by considering a final, big picture implication of the arguments developed here. Many philosophers have been attracted to K-Aim, in part, because it promises 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 argued 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. The key explanatory notion in EV-Aim is the notion of an epistemically valuable credence. Even if (maximal) epistemic value entails knowledge, it is questionable whether knowledge is doing any of the explanatory heavy-lifting.

For this reason, I think we should abandon the project of explaining the value of knowledge in terms of its connection with inquiry. We are left 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.[36] 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 epistemically valuable credences.

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. New Work for Certainty. *Philosophers' Imprint*, 20(8):1–25, 2020.

Bob Beddor and Simon Goldstein. Mighty knowledge. *Journal of Philosophy*, 118(5):229–269, 2021.

Matthew Benton. Dubious objections from iterated conjunctions. *Philosophical Studies*, 162(2):355–358, 2013.

Jessica Brown. Subject-Sensitive Invariantism and the Knowledge Norm for Practical Reasoning. *Noûs*, 42(2):167–189, 2008.

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

[36] However, some of my earlier arguments undercut some of these alternative applications of knowledge. For example, 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 leveraged to cast doubt on a knowledge norm of assertion.

Lara Buchak. Instrumental Rationality, Epistemic Rationality, and Evidence-Gathering. *Philosophical Perspectives*, 24(1):85–120, 2010.

Catrin Campbell-Moore and Bernhard Salow. Avoiding Risk and Avoiding Evidence. *Australasian Journal of Philosophy*, 98(3):495–515, 2020.

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.

Keith DeRose. *The Case for Contextualism: Knowledge, Skepticism, and Context*. Clarendon Press, Oxford, 2009.

Dylan Dodd. Against fallibilism. *Australasian Journal of Philosophy*, 89(4):665–685, 2011.

Dylan Dodd. Belief and Certainty. *Synthese*, 194(11):4597–4621, 2017.

Kevin Dorst. Abominable KK Failures. *Mind*, 128(512):1227–1259, 2019.

Arianna Falbo. Inquiry and confirmation. *Analysis*, 81(4):622–631, 2021.

Arianna Falbo. Inquiring minds want to improve. *Australasian Journal of Philosophy*, forthcoming.

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.

Rachel Fraser. The will in belief. In Szabo Gendler and Hawthorne, editors, *Oxford Studies in Epistemology*. Oxford University Press, Oxford, forthcoming.

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.

Jane Friedman. The Epistemic and the Zetetic. *The Philosophical Review*, 129(4):501–536, 2020.

Jane Friedman. Zetetic epistemology. In *Towards an Expansive Epistemology: Norms, Action, and the Social Sphere*. Routledge, forthcoming.

Simon Goldstein. Omega knowledge. ms., 2022.

Simon Goldstein. A preface paradox for intentions. *Philosophers Imprint*, forthcoming.

I.J. Good. On the Principle of Total Evidence. *British Journal for the Philosophy of Science*, 17(4):319–321, 1967.

Jeremy Goodman and Ben Holguín. Thinking and being sure. *Philosophy and Phenomenological Research*, forthcoming.

Jeremy Goodman and Bernhard Salow. Taking a Chance on KK. *Philosophical Studies*, 175:183–196, 2018.

Hilary Greaves and David Wallace. Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility. *Mind*, 115(459):607–632, 2006.

Daniel Greco. Could KK be OK? *Journal of Philosophy*, 111(4):169–197, 2014.

Daniel Greco. How I Learned to Stop Worrying and Love Probability 1. *Philosophical Perspectives*, 29(1):179–201, 2015a.

Daniel Greco. Iteration principles in epistemology: Arguments for. *Philosophy Compass*, 10(11):754–764, 2015b.

Gilbert Harman. *Thought.* Princeton University Press, Princeton, 1973.

John Hawthorne. *Knowlege and Lotteries.* Oxford University Press, New York, 2004.

John Hawthorne and Jason Stanley. Knowlege 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.

Paul Horwich. *Probability and Evidence.* Cambridge University Press, Cambridge, 1982.

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*, 198:1583–1593, 2021a.

Christoph Kelp. *Inquiry, Knowledge, and Understanding.* Oxford University Press, Oxford, 2021b.

Christoph Kelp. Theory of Inquiry. *Philosophy and Phenomenological Research*, forthcoming.

Peter Klein. *Certainty: A Refutation of Scepticism.* University of Minnesota Press, Minneapolis, 1981.

Jason Konek. Probabilistic knowledge and cognitive ability. *Philosophical Review*, 125, 2016.

Saul Kripke. Two Paradoxes of Knowledge. In *Philosophical Troubles.* Oxford University Press, Oxford, 2011.

Jonathan Kvanvig. The Value of Understanding. In Pritchard, Millar, and Haddock, editors, *Epistemic Value*, pages 95–111. Oxford University Press, Oxford, 2009.

Maria Lasonen-Aarnio. Unreasonable Knowledge. *Philosophical Perspectives*, 24:1–21, 2010.

Maria Lasonen-Aarnio. The Dogmatism Puzzle. *Australasian Journal of Philosophy*, 92(3): 417–432, 2014.

Maria Lasonen-Aarnio. Perspectives and good dispositions. *Philosophy and Phenomenological Research*, forthcoming.

Isaac Levi. *The Fixation of Belief and its Undoing: Changing Beliefs Through Inquiry.* Cambridge University Press, Cambridge, 1991.

David Lewis. *Convention.* Harvard University Press, Cambridge, 1969.

David Lewis. Immodest Inductive Methods. *Philosophy of Science*, 38(1):54–63, 1971.

Adam Marushak. Fallibilism and consequence. *Journal of Philosophy*, 118(4):214–226, 2021.
Aidan McGlynn. *Knowledge First?* Palgrave McMillian, New York, 2014.
Richard Miller. Absolute Certainty. *Mind*, 87:46–65, 1978.
Sarah Moss. Scoring rules and epistemic compromise. *Mind*, 120(480):1053–1069, 2011.
Sarah Moss. *Probabilistic Knowledge*. Oxford University Press, Oxford, 2018.
Sarah Moss. Full Belief and Loose Speech. *Philosophical Review*, 128(3):255–291, 2019.
Jennifer Nagel. *Knowledge: A Very Short Introduction*. Oxford University Press, Oxford, 2014.
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.
Jacob Ross and Mark Schroeder. Belief, Credence, and Pragmatic Encroachment. *Philosophy and Phenomenological Research*, 88(2):259–288, 2014.
Daniel Rothschild. What it takes to believe. *Philosophical Studies*, 177:1345–1362, 2020.
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.
Yasha Sapir and Peter van Elswyk. Hedging and the Ignorance Norm on Inquiry. *Synthese*, 199(3-4):5837–5859, 2021.
Crispin Sartwell. Why knowledge is merely true belief. *Journal of Philosophy*, 89(4): 167–180, 1992.
Jonathan Schaffer. Contrastive knowledge. In *Oxford Studies in Epistemology*, volume 1, pages 235–271. Oxford University Press, Oxford, 2005.
Miriam Schoenfield. Bridging rationality and accuracy. *Journal of Philosophy*, 112(12): 633–657, 2016.
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. *Knowlege 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.
David Thorstad. Inquiry and the Epistemic. *Philosophical Studies*, 178:2913–2928, 2021.
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.
Dennis Whitcomb. One Kind of Asking. *The Philosophical Quarterly*, 166/67:148–168, 2017.
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.