

# 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?

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.<sup>1</sup> In the recent flurry of work 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  $Q$ .<sup>2</sup>

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 forthcoming). After all, we typically engage in inquiry in order to satisfy our curiosity.

---

<sup>1</sup>I borrow the label ‘zetetic’ from Friedman 2020, forthcoming, who has been one of the most prominent champions of the zetetic turn.

<sup>2</sup>While different authors formulate this idea in slightly different ways, versions of K-Aim are defended in Kvanvig 2009; Kappel 2010; Kelp 2011, 2014, 2021a,b, forthcoming; Rysiew 2012; Friedman 2013, 2017; Whitcomb 2017; Sapir and van Elswyk forthcoming, among others.

Despite its appeal, this paper argues that K-Aim stands in tension with another highly attractive epistemological doctrine: fallibilism. While 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  $p$ .<sup>3</sup>

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 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 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.

---

<sup>3</sup>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.

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.<sup>4</sup> 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  $r$ .<sup>5</sup> 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

---

<sup>4</sup>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.

<sup>5</sup>This 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.

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.

### 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.<sup>6</sup>

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.

---

<sup>6</sup>For a sampling of important contributions to epistemic decision theory, see [Joyce 1998](#); [Greaves and Wallace 2006](#); [Moss 2011](#); [Pettigrew 2016](#); [Schoenfeld 2016](#); among many others.

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.

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. By EV-Aim, it follows that 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  $\alpha$ , and you rationally ought to believe that you have not yet attained  $\alpha$ , then you rationally ought to continue pursuing  $\alpha$ .

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.<sup>7</sup>

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. 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.

## 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

One way of trying to reconcile K-Aim and Fallibilism is to opt for some form of 'impurism': the view that knowledge depends on practical factors.<sup>8</sup> While there are different ways of developing 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 considerations rob our protagonists of their knowledge. So they have not achieved the aim of inquiry after all.

---

<sup>7</sup>We 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. 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 is epistemically rational to do whatever maximizes expected epistemic value.

<sup>8</sup>AKA 'pragmatic encroachment.' For defenses, see e.g., Fantl and McGrath 2002, 2009; Stanley 2005; Weatherson 2012; Ross and Schroeder 2014.

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.

Perhaps a different form of impurism will 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.

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 reconciliation are put in the uncomfortable position of being forced to reject one of these tenets.

## 4.2 Aims vs. rational requirements

A second reconciliation strategy is to distinguish between aims and rational requirements. According to this response, 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  $\alpha$  and rationally believing that you have attained  $\alpha$ . 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.<sup>9</sup>

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.

We can put this point more generally. Presumably, there is some point at which you are rationally permitted to cease pursuing an aim. The following seems like a plausible sufficient condition:

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

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 they are not rationally required to engage in further inquiry.

---

<sup>9</sup>A 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; Kelp and Simion 2017; Greco forthcoming; Williamson forthcoming, among others.

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  $\alpha$ , and you *know* that you have already attained  $\alpha$ , then you are not rationally required to continue pursuing  $\alpha$ .

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  $p$ .<sup>10</sup> And 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. Still, if they are not certain of the answers, it may 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.

### 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 have achieved one of the aims of inquiry. But they will not have achieved every aim of inquiry, since they will not have maximized epistemic value.

While this is a simple path to harmony, it incurs significant costs. 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 achieves complete certainty on the matter. How is this helpful for answering the question of what Poirot rationally ought to do?

Perhaps if we are pluralists about aims, we should also be pluralists about rational "ought"s. Relative to the K-Aim, Poirot ought to conclude his case when

---

<sup>10</sup>Most theories that make room for fallible knowledge will also make room for fallible higher-order knowledge. 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  $p$ .

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 want some way of adjudicating between these conflicting ought claims; we want our theory to say what Poirot ought to do *full stop*.

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. 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. 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 as 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.<sup>11</sup>

We can reinforce this point by adducing ascriptions of fallible knowledge “from the wild”. As [Beddor 2020](#) notes, a quick Google search unearths many naturally occurring attributions of knowledge with “near certainty”, implying that the knowledge is not held with complete certainty. Some examples:

“We now know with near-certainty that Wall Street execs committed felonies.”<sup>12</sup>

“[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.”<sup>13</sup>

“When business resumes, we know with near-certainty that it will likely be a deluge.”<sup>14</sup>

Taken together, these considerations provide reason to think that our ordinary notion of knowledge is fallibilist. They thereby provide impetus for exploring the other response to our tension: rejecting K-Aim.

## 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 improving the epistemic value of our

<sup>11</sup>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](#)). For related evidence for Fallibilism, see the discussion of “I think I know...” constructions in [Worsnip 2015](#).

<sup>12</sup><https://www.businessinsider.com/my-plan-to-finally-make-wall-street-pay-for-its-crimes-2012-3>

<sup>13</sup><https://www.nba.com/clippers/news/five-things-watch-five-games-left>

<sup>14</sup><https://www.nuix.com/blog/forming-good-remote-work-habits-law-firms>

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.<sup>15</sup> 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.<sup>16</sup>

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.<sup>17</sup>

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.”<sup>18</sup> 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

<sup>15</sup>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.

<sup>16</sup>For related discussion of how to apply modal conditions to credences, see Moss 2018; Beddor and Goldstein 2021.

<sup>17</sup>See Konek 2016 for a defense of an aptness condition on credences.

<sup>18</sup>See Pasnau 2017 for relevant historical discussion. See Beddor 2020 for a recent attempt to rehabilitate the notion of epistemic certainty and assign it a starring role in epistemology.

$p$  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.<sup>19</sup> 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?

Some might worry 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. 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, for the sake of argument, that EV-Aim is seldom attainable. Would this show that the EV-Aim is false, or even implausible?

---

<sup>19</sup>These candidates are not necessarily rivals. For example, one might maintain that maximal epistemic certainty entails safety or aptness (or both). 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](#).

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 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 Consequences

I’ve argued that we should replace K-Aim with EV-Aim. This replacement has important implications for a number of related topics in epistemology. By way of conclusion, I’ll briefly investigate some of these consequences.

### 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 observes, these attitudes seem to be intimately tied to inquiry. When Poirot inquires into who committed the murder, it is natural to describe him as wondering who did the deed.

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; see also Whitcomb 2017; Sapir and van Elswyk forthcoming). 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, and 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 an answer. The functional role of interrogative attitudes is to motivate us to engage in further inquiry with the goal of maximizing the epistemic value of our credences. And this goal is only fully satisfied when we have achieved certainty on the matter at hand.

## 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 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  $r$ .<sup>20</sup>

---

<sup>20</sup>Some philosophers might question whether this proposal is a genuine alternative to DBI. Recently, a number of epistemologists have advanced the surprising view that belief entails certainty (e.g., Clarke 2013; Greco 2015; Dodd 2017; Moss 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 uncertain Tess to remark, “I think I know the answer to this one...”, it also would be natural for her to remark, “I believe the answer is...” This suggests the ordinary con-

### 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 A knows  $p$ , then A is permitted to take  $p$  for granted in practical reasoning.<sup>21</sup>

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 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 the impurist, the protagonists of Brown and Reed’s cases do not know  $p$  after all, since the high stakes of the decision deprives them of knowledge ([Fantl and McGrath 2009](#): 62-63). But we have already seen that impurism fails to satisfactorily resolve the tension between K-Aim and Fallibilism (§4.1). So even if impurism helps with Brown and Reed’s cases, it will not provide a sufficiently general defense of KN.

---

ception 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](#).

<sup>21</sup>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](#).

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 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.

#### 7.4 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 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 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 other conceptions of epistemic value—for example, views on which epistemic value involves safety, or aptness, 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.<sup>22</sup> Another

---

<sup>22</sup>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

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. *The Journal of Philosophy*, 118(5):229–269, 2021.
- 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.
- 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.
- Jane Friedman. The Epistemic and the Zetetic. *Noûs*, 129(4):501–536, 2020.
- Jane Friedman. Zetetic epistemology. In *Towards an Expansive Epistemology: Norms, Action, and the Social Sphere*. Routledge, forthcoming.
- Hilary Greaves and David Wallace. Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility. *Mind*, 115(459):607–632, 2006.

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.

- Daniel Greco. How I Learned to Stop Worrying and Love Probability 1. *Philosophical Perspectives*, 29(1):179–201, 2015.
- Daniel Greco. Justification and Excuses in Epistemology. *Noûs*, forthcoming.
- 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.
- 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.
- 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. *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.
- Jacob Ross and Mark Schroeder. Belief, Credence, and Pragmatic Encroachment. *Philosophy and Phenomenological Research*, 88(2):259–288, 2014.
- 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*, forthcoming.
- 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. *Knowledge and Practical Interests*. Oxford University Press, New York, 2005.
- Jason Stanley. Knowledge and Certainty. *Philosophical Issues*, 18(1):35–57, 2008.
- 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.