## INQUIRY BEYOND KNOWLEDGE

#### **Bob Beddor**

Philosophy and Phenomenological Research, forthcoming

#### **Abstract**

Why engage in inquiry? According to many philosophers, 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: knowledge does not require absolute certainty. Forced to choose between these two theses, I argue that we should reject the idea that inquiry aims at knowledge. I go on to develop an alternative view, according to which inquiry aims at maximizing the epistemic value of our credences. This alternative view makes room for knowledge that falls shy of certainty, and it coheres nicely with a rich body of work in epistemic decision theory. I proceed to highlight the implications of this replacement for some important topics in epistemology, including the dogmatism paradox, the nature of interrogative attitudes, and the norm of practical reasoning.

# 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 investigating. 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. In the recent flurry of work on this topic, one can discern a consensus beginning to emerge: *inquiry aims at knowledge*. That is:

<sup>&</sup>lt;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. For discussion, see also Thorstad 2021; Falbo forthcoming.

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 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).

Despite its appeal, this paper argues that K-Aim stands in tension with another highly attractive principle:

Knowledge Without Certainty: It's possible for a rational agent to know p without being absolutely certain that p.<sup>3</sup>

I present two arguments that K-Aim stands in tension with Knowledge Without Certainty. The first is case-based ( $\S 2$ ). To preview: take 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. The second argument is principle-based ( $\S 3$ ). I show that the conjunction of K-Aim and Knowledge Without Certainty conflicts with some plausible principles of epistemic decision theory.

After developing these arguments in more detail (§§2-3), I consider some strategies for reconciling K-Aim and Knowledge Without Certainty and find them wanting

<sup>&</sup>lt;sup>2</sup>The idea that there is an important connection between inquiry and knowledge traces back to Plato's *Meno*. In the contemporary literature, 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; Carter and Hawthorne forthcoming. Other philosophers posit a similar connection between inquiry and knowledge without presupposing that inquiry has any distinctive *telos*. For example, Friedman 2023 expresses skepticism about the idea that inquiry has an aim. However, she also defends a normative requirement according to which one ought not to inquire into some question if one knows the answer to this question (Friedman 2017). My arguments in this paper generalize straightforwardly to this normative variant of K-Aim.

 $<sup>^3</sup>$ Knowledge Without Certainty bears a close relation to fallibilism, the view that someone can know p even though their evidence does not eliminate the possibility that p is false (Cohen 1988; Reed 2013; Worsnip 2015; Brown 2018). It's a familiar thought that rational agents proportion their degrees of belief to the evidence. According to a natural way of unpacking this evidentialist principle, if your evidence does not eliminate the possibility that p is false, you should not be completely certain that p. Combined with fallibilism, it follows that an agent can know p without being completely certain of p. While fallibilism—combined with this evidentialist constraint—entails Knowledge Without Certainty, Knowledge Without Certainty is logically weaker than fallibilism. Consequently, infallibilists are also welcome to take it on board. (Thanks to a referee for helpful points here.) I discuss the advantages of Knowlege Without Certainty at greater length in §5.

(§4). So we are forced to choose between the two principles. I argue that the balance of considerations tips in favor of Knowledge Without Certainty and against K-Aim (§§5-6). These considerations 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 Knowledge Without Certainty, and it plays nicely with 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, and the norm of practical reasoning (§7).

# 2 Bringing out the tension: two cases

In order to bring out the tension between K-Aim and Knowledge Without Certainty, start with a case of inductive 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 Knowledge Without Certainty 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.

On at least some ways of developing this scenario, it seems Mia has a *pro tanto* reason to look at the results of the study. In support of this judgment, imagine Mia reasons as follows:

Mia's Reasoning I'm pretty sure that m is true. But I am not completely certain of m, and I'd like to settle the matter. Reading the results of this study might help me figure out for certain whether m is true, or at least come closer to this goal.

This seems like perfectly good means-ends reasoning. First, the desire to be certain (or closer to certain) whether *m* is true seems like a perfectly rational desire. Many of us reasonably desire to be sure whether we are right about some topic, particularly if it is a topic that is important to us.<sup>4</sup> Second, Mia rationally believes that reading the study is a means to achieving this goal. Given that rational means-ends reasoning typically generates a *pro tanto* reason to perform the believed means to one's ends, we should regard **Mia's Reasoning** as generating a *pro tanto* reason for her to read the study.<sup>5</sup>

However, this judgment 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.

Let me briefly consider two ways defenders of K-Aim might try to respond to this point. First, proponents of K-Aim may reply that while Mia has attained the aim of inquiry into the question, *Is m true?*, there are various other questions whose answers Mia does not know. For example, she might be interested in the question, *What were the results of the researcher's study?* She does not know the answer to this question, hence it is rational for to read the results.

However, I don't think this response gets to the heart of the problem. We can imagine that Mia is only derivatively interested in discovering the results of the researcher's latest study. She does not care about the results for their own sake, but only insofar as they will help her answer the question of whether m is true. In this version of the case, she still has a reason to read the results. To support this verdict, consider the internal monologue imagined above (Mia's Reasoning). Mia's reason for reading the study is that it will help her settle the question of whether m is true, not that it will help her answer some other question.

Alternatively, some might grant Mia has a reason to read the study, but deny that this constitutes genuine inquiry. Rather, she is simply double-checking that m is true, which differs from inquiring into whether m is true.

In response, I am happy to grant that Mia is double-checking whether m is true. But I think we should push back on the assumption that double-checking is incompatible with inquiry. Suppose that Mia goes through the reasoning outlined above, and

<sup>&</sup>lt;sup>4</sup>There is a substantial body of work in psychology on people's "need for closure", understood as a desire to arrive at a settled answer to some question. See e.g., Kruglanski and Webster 1996.

<sup>&</sup>lt;sup>5</sup>Two quick clarifications. First, the claim that Mia has a *pro tanto* reason to read the study is weaker than the claim that she *ought* to read it. After all, she might also have a *pro tanto* reason to spend her time on other activities (a point to which I'll return in §3). Second, we can remain neutral for the moment on whether the reasons in question are practical or epistemic. in §3, I'll give some reason thinking the answer is *both*: she has a practical reason to inquire, insofar as doing so will help satisfy her desire to attain certainty about *m*. She also has an epistemic reason to inquire, insofar as doing so will is expected to increase the epistemic value of her credences.

<sup>&</sup>lt;sup>6</sup>Thanks to a referee for suggesting this response.

decides to read the study. Her decision—and her resulting behavior—bears many of the hallmarks of inquiry. First, she is trying to answer a question under investigation: *Is m true?* Second, she is trying to answer this question in a distinctive way: she is intentionally gathering and examining evidence that she deems relevant to settling this question. Finally, she revises her attitudes towards the question under investigation in light of this evidence. In all of these respects, her behavior resembles paradigmatic inquiries. In paradigmatic inquiries, someone seeks to answer some question Q. In order to do so, they intentionally gather evidence relevant to answering Q. They then revise their attitudes towards Q in light of the evidence. These resemblances make a strong presumptive case that Mia is engaged in genuine inquiry.<sup>7,8</sup>

There are other ways one might try to resolve the tension between K-Aim and Knowledge Without Certainty; I will consider some of these 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. From her studying, she knows *r*: *The Western Roman Empire fell in 476 CE.* However, it has been some time since she reviewed. So she rationally assigns some credence to the possibility that she got the dates wrong.

If Knowledge Without Certainty is true, then presumably cases of this sort are possible. 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

<sup>&</sup>lt;sup>7</sup>This is not to say that every case of double-checking is a form of inquiry. When someone obsessively checks again and again that the stove is turned off, it is at least debatable whether they are really inquiring. Crucially, however, this compulsive stove-checker's behavior lacks the paradigmatic features of inquiry present in **Murine Research**. In particular, it is doubtful whether the obsessive re-checker is motivated by a desire to discover the truth, rather than simply indulging in a reflexive habit. It is also not clear whether the obsessive re-checker is getting new evidence after, say, the fourth check, and whether they are revising their degrees of belief in light of that evidence. Of course, we could try to flesh out the case so that these features are present, but then it becomes more tempting to describe the stove checker's actions as a genuine (if perhaps ill-advised) form of inquiry. For further defense of the idea that double-checking can constitute genuine inquiry, see Friedman 2019b; Woodard forthcoming.

 $<sup>^8</sup>$ One possible response would be to hold that genuine inquiry is constitutively tied to interrogative attitudes, such as wondering and curiosity, whereas mere checking is not. I will defer consideration of this response until §7.2, where I argue that Mia's desire to figure out for certain whether m is true qualifies as an interrogative attitude.

 $<sup>^9</sup>$ 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.

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.

On at least some ways of developing this scenario, it seems Tess has a *pro tanto* reason to consult her textbook. We can support this verdict using a variant of **Mia's Reasoning**. Imagine Tess wants to be sure (or closer to sure) whether *r* is true. She also rationally believes that by checking the textbook, she can achieve this goal. Given that rational means-ends reasoning generates a *pro tanto* reason to perform the believed means to one's ends, she has a *pro tanto* reason to check.

But this verdict conflicts with K-Aim. Since Tess knows r, K-Aim says that she has attained the aim of inquiry on the question of whether r is true. So there should be no reason—at least no reason stemming from the aims of inquiry—for her to inquire further.

Taking stock: I have presented two cases that draw out a tension between K-Aim and Knowledge Without Certainty. In both cases, the agent knows the answer to some question under investigation. But, I have argued, they still have reason to inquire further. Now, there remains a question as to whether my judgments about these cases are widely shared. This is an empirical question, and not one that I can hope to settle here. That said, let me briefly report a preliminary study I conducted on this front. I recruited 100 participants on *Prolific* (a site for online surveys) and presented them with a slightly simplified version of **Ancient History**. Participants were asked, 'Should Tess check the date?' They were given two options:

- (a) Yes, she should check the date to make sure.
- (b) No, there's no point in checking since she already knows the answer.

90% of the participants selected option (a). The details of the study and the anonymized raw results can be found here.  $^{10}$ 

Of course, there are complicated and controversial questions about what role folk intuitions should play in philosophy. For my purposes, the important point is that these results provide some evidence that my judgments about this sort of case are not idiosyncratic; there is reason to think they are widely (though not universally) shared. <sup>11</sup>

<sup>&</sup>lt;sup>10</sup>I phrased the question in terms of what Tess should do rather than what she has *pro tanto* reason to do, since the former locution is less technical.

<sup>&</sup>lt;sup>11</sup>For complementary arguments that agents sometimes have reason to inquire beyond knowledge, see Falbo 2023; Woodard forthcoming.

# 3 The argument from epistemic decision theory

Our argument need not rest on cases alone. This section develops a complementary argument that K-Aim stands in tension with Knowledge Without Certainty—an argument grounded in epistemic decision theory.

## 3.1 Three principles

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>12</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.

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.

A wide variety of formal measures of accuracy 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.

<sup>&</sup>lt;sup>12</sup>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.

## 3.2 An inconsistent quintet

The principles I've laid out so far are fairly minimal. But they cannot be combined with both K-Aim and Knowledge Without Certainty. That is, Valuable Accuracy, Alethic Proximity, EV-Aim, Knowledge Without Certainty and K-Aim are jointly inconsistent.

To see this, take a case of knowledge without certainty, such as **Ancient History**. (This case is just for the purpose of illustration; any other case of knowledge without certainty will do just as well.) We stipulated 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: if her credence in r were .99 or 1, it would be more accurate still. By Valuable Accuracy, her credence in r is not maximally valuable. By EV-Aim, she has not attained the aim of inquiring into the question, *When did the Roman Empire fall*? But this contradicts K-Aim, which says that she has attained the aim of inquiry.

Consequently, those who wish to keep both K-Aim and Knowledge Without Certainty will be forced to give up one of our three principles. However, each of these principles is highly plausible in its own right.

Start with Valuable Accuracy. This thesis encodes the minimal "veritist" idea that attaining truths and avoiding errors are epistemic goods. From an epistemic perspective, we care about getting things *right*. Importantly, Valuable Accuracy is considerably weaker than the strong veritist claim that accuracy is the sole source of epistemic value; it merely claims that this is one important contributing factor. (More on this in §6.)

Next, consider Alethic Proximity. As noted above, Alethic Proximity is a straightforward generalization of the intuition that a .9 credence in a truth is more accurate than a .2 credence in a truth, and a .99 credence in a truth is more accurate still. Moreover, denying Alethic Proximity means that someone who assigns positive credences to falsehoods can have a *maximally* accurate credences, which seems absurd.

Perhaps, then, we should abandon EV-Aim. Don't we already have a perfectly good alternative, namely, K-Aim? However, this would be too quick. Epistemic value is, by definition, something we care about from the epistemic point of view: we want to attain epistemically valuable states. Given this, it's extremely natural to think that the epistemic value of inquiry is closely tied to the epistemic value of the states to which it leads. This idea is already implicit in K-Aim. Part of the reason why so many epistemologists have been attracted to K-Aim is that knowledge is thought to be an epistemically valuable. But this suggests that K-Aim really derives its appeal from EV-Aim.

#### 3.3 Back to reasons

I have shown that three plausible principles are inconsistent with the conjunction of K-Aim and Knowledge Without Certainty. This argument holds even if one does not share the intuitions about my cases in §2. That said, the framework developed here has some interesting interactions with the cases from §2. In particular, we can use this framework to provide further justification for the conclusion that Mia and Tess have a reason to continue inquiry.

One way to do so would be to appeal to a plausible bridge principle linking up one's beliefs about whether one has fulfilled the aims governing some activity with one's reasons for performing that activity, for example:

AIM-RATIONALITY BRIDGE If you have a reason to pursue aim  $\alpha$ , and you rationally ought to believe that you have not yet attained  $\alpha$ , then you have a reason 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 have a reason to continue their inquiries.

We can reach the same conclusion by a different route. In an important contribution to epistemic decision theory, Oddie 1997 proves that, given certain assumptions, 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.<sup>13</sup> 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 change her credence in *r*. In particular, there's a chance that she will read that the Western Roman Empire did fall in 476, in which case she will 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.

 $<sup>^{13}</sup>$  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. 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.

How does this bear on what Tess should do? According to many decision theorists, practical rationality requires maximizing expected practical value. Many epistemic decision theorists have embraced the epistemic counterpart of this idea:

MAXIMIZE EXPECTED EPISTEMIC VALUE Epistemic rationality requires maximizing expected epistemic value.<sup>14</sup>

Now, from this it does not necessarily follow that Tess is epistemically required to check her textbook. There may be other inquiries available to Tess: perhaps she could spend the five minutes before the test reading about some other chapter of Roman history. Oddie's theorem applies equally well to these alternative investigations. If Tess decides to read about the Second Punic War (for example), and if doing so might affect her credences, then the expected epistemic value of the gathering this evidence and conditionalizing on it exceeds the expected epistemic value of sticking with her current credal state. So nothing in Oddie's theorem tells us which of these two inquiries is epistemically superior.

While we cannot use Oddie's theorem to establish that Tess is epistemically required to check whether r is true, we can use it to establish the weaker claim that she has a *pro tanto* (epistemic) reason to check. This leaves open the possibility that she also has a *pro tanto* (epistemic) reason to devote her time to other inquiries. Which inquiry does she have most epistemic reason to pursue? The answer will depend on further factors—for example, the probability of making epistemic progress on the question at hand vs. other questions. If the expected epistemic benefit of checking r exceeds the expected epistemic benefit of pursuing these other inquiries, then the framework developed here predicts that Tess epistemically ought to check r. If the expected epistemic benefit of pursuing some other inquiry is greater, she epistemically ought to pursue that inquiry instead.

Let me pause to address a concern and a question. First, the concern. I have construed the epistemic requirement to maximize expected epistemic value in a fairly liberal fashion. On my construal, this requirement extends to inquiries. This liberal construal is admittedly controversial. After all, inquiry is an *activity*. According to a more conservative picture, epistemic reasons and requirements only apply to doxastic *states*. A full defense of the more liberal construal falls outside the scope of this paper. However, I think it enjoys intuitive support. As Singer and Aronowitz 2022 observe, it seems we have epistemic reasons to perform inferences (say, inferring the conclusions supported by our evidence) and to consider hypotheses (say, when a scientist

 $<sup>^{14}</sup>$ See Greaves and Wallace 2006; Easwaran 2013; Schoenfield 2018; Dorst 2019b, among others. Following these authors, I have formulated this requirement in terms of maximizing *overall* epistemic value. One could also opt for a question-relative version of this requirement, according to which agent who is inquiring (or ought to be inquiring) into questions  $Q_1 \dots Q_n$  is epistemically required to maximize expected epistemic value with respect to  $Q_1 \dots Q_n$ . My argument in this section would go through equally well using this question-relative variant. (Thanks to a referee for helpful questions here.)

is presented with a hypothesis that gives a compelling explanation of their data). But *inferring* and *considering* are activities, not doxastic states. So anyone who is persuaded that we have epistemic reasons to infer and consider is already on the path towards a more liberal construal of the epistemic realm. <sup>15</sup> Furthermore, the liberal construal seems to be shared by many defenders of K-Aim. According to them, K-Aim is supposed to articulate a distinctly epistemic aim of inquiry, which is supposed to generate distinctly epistemic reasons.

That said, those who prefer the more conservative picture can still take on board many of my claims. On this picture, we practically ought to inquire (even beyond knowledge) insofar as we practically ought to maximize the expected epistemic value of our credences (cf. Maguire and Woods 2020). This might be because we desire to find out for certain whether some proposition is true, as in Mia's Reasoning. Or it might be because improving the epistemic value of our credences will help fulfill further desires—say, Tess' desire to pass the exam—whose satisfaction hinges on the accuracy of these credences.<sup>16</sup>

On to the question. I suggested that inquiry aims at maximizing epistemic value (EV-Aim). I also defended the idea that epistemic rationality requires maximizing *expected* epistemic value. One might wonder: well, which is it?<sup>17</sup> In response, consider the analogy with the practical domain. Should agents aim at maximizing utility or at maximizing expected utility? While this is a matter for debate, one popular answer is *both*. The injunction to maximize utility governs what one objectively ought to do, whereas the injunction to maximize expected utility governs what one subjectively ought to do.

Extending this diagnosis to the epistemic case yields the following picture. EV-Aim provides an objective aim of inquiry; it tells us how we ought to conduct our inquiries from the point of view of the objective epistemic *ought*. Maximize Expected Epistemic Value tells us how we ought to conduct our inquiries from the point of view of the subjective epistemic *ought*. However, for our purposes we need not commit to this particular way of understanding the relation between the two norms. Whatever we say about the practical case should carry over to the epistemic case.<sup>18</sup>

<sup>&</sup>lt;sup>15</sup>For further arguments in this vein, see McWilliams 2023; Flores and Woodard forthcoming.

<sup>&</sup>lt;sup>16</sup>Of course, there will be some cases where the practical and epistemic incentives pull in different directions. Suppose the question on the Fall of Rome won't be graded, whereas the question on the Punic War will. And suppose it matters a great deal how Tess performs on the exam. Then Tess (practically and all-things-considered) ought inquire into the latter question, even if the expected epistemic benefit of inquiring into the former is greater. Both the liberal and the conservative construals can agree with this assessment. However, the liberal construal will also hold that there is a distinctively epistemic sense in which she should pursue the former inquiry. It's just that practical considerations conflict with—and outweigh—her epistemic reasons. Thanks an editor for helpful points here.

 $<sup>^{\</sup>rm 17}{\rm Thanks}$  to a referee for raising this question.

<sup>&</sup>lt;sup>18</sup>Of course, there will be cases where the objective ought and the subjective ought pull in different directions, as reflected in the Miners Puzzle. This puzzle has spawned a large literature (see Kolodny

Let's take stock. §2 offered cases to draw out the tension between K-Aim and Knowledge Without Certainty. In this section, I bolstered this case-based argument with a principle-based argument. I laid out three plausible principles from epistemic decision theory and showed that they are inconsistent with the conjunction of K-Aim and Knowledge Without Certainty. This principle-based argument also provides a deeper theoretical understanding of the cases in §2. According to the diagnosis that emerges, 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 Knowledge Without Certainty, this conflicts with K-Aim.<sup>19</sup>

# 4 Reconciliation strategies

I now turn to consider some ways of trying to reconcile K-Aim and Knowledge Without Certainty. I argue that each of these reconcilation strategies faces problems. 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 Knowledge Without Certainty is to appeal to "impurism", the view that knowledge depends on practical factors. 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. But knowledge is sensitive to

and MacFarlane 2010; Jackson 1985; Carr 2015; Műnoz and Spencer 2021, among many others). Here too, we need not take a stand on how to resolve this puzzle. Whatever solution one opts for in the practical domain will also carry over here.

It is also worth noting that the main argument of this paper can get off the ground using *either* the objective aim of inquiry or its subjective cousin. We saw in §3.2 that EV-Aim—when combined with Alethic Proximity and Valuable Accuracy—is inconsistent with the conjunction of K-Aim and Knowledge Without Certainty. The same inconsistency arises using the subjective version of EV-Aim. Assuming Tess is aware of Alethic Proximity and Valuable Accuracy, then she can expect to increase the value of her credences by inquiring, and so—by the view that we ought to maximize expected epistemic value—it follows that she has a reason to inquire.

<sup>19</sup>My defense of EV-Aim bears some affinities to a view independently developed by Falbo 2023, according to which inquiry aims at 'epistemic improvement'. 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. This point also distinguishes my approach from that of Woodard forthcoming, who also expresses sympathy for pluralism about the goals of inquiry. These differences aside, our projects are largely congenial: I view my approach as meshing nicely with the central claims put forward by Falbo and Woodard.

<sup>20</sup>AKA 'pragmatic encroachment.' For defenses, see e.g., Fantl and McGrath 2002, 2009; Stanley 2005; Weatherson 2012; Ross and Schroeder 2014.

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 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. 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 further 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, "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.

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. It seems implausible that knowledge can be gained and lost so easily.

A further problem is more general. No version of the impurist reconciliation strategy will address the argument from epistemic decision theory (§3). As we saw in §3.2,

K-Aim and Knowledge Without Certainty are directly inconsistent with three plausible principles from epistemic decision theory. Nothing in this proof presupposed a purist theory of knowledge. Thus proponents of the impurist reconciliation strategy are put in the uncomfortable position of being forced to reject one of these principles.

#### 4.2 Derivative norms

A prominent way of defending knowledge norms of assertion, action, and the like is to appeal to derivative norms. Could this strategy reconcile K-Aim and Knowledge Without Certainty?

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  $\alpha$  is to pursue that activity until you have achieved  $\alpha$ . 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.<sup>21</sup>

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.

What are these derivative norms? One option is to go dispositional. For example, Williamson 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'). 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. Minny is disposed to 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.

 $<sup>^{21}</sup>$ 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.

Would some other derivative knowledge norm fare better? DeRose (2002: 180, 2009: 93) 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 plausible; we relied on a variant of this idea in §3 (Aim-Rationality Bridge). Combining this 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 noted above, we can imagine Tess thinking 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. Some might suggest 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*. Suppose that Tess believes—but isn't certain—that she knows when the Western Roman Empire fell. By stipulation, this higher order belief (that is, her belief that she knows when it fell) is true. This belief is also justified: it is based on her memory of having read that it fell in 476, and her memory is reliable. Finally, this belief is free from Gettierization. So there seems to be no reason why it couldn't amount to knowledge.

To bolster this point, note that most theories that make room for fallible knowledge also make room for fallible second order knowledge. 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 .96, and your credence that you safely believe p is .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 expect there to be versions of our cases where the agent knows that they have attained knowledge on the question at

 $<sup>^{22}</sup>$ 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 criticism of the KK thesis, see Williamson 2000, chp.5; for defenses, see e.g., Greco 2014, 2015; Dorst 2019a.

hand, without being certain of the answer. In such cases, they will satisfy the Second Order Knowledge Norm. Still, it may be rational for them to continue inquiring.<sup>23</sup>

It thus proves difficult to develop the derivative norm strategy in a way that handles the full range of cases. Moreover, the argument from epistemic decision theory provides grounds for pessimism about any solution along these lines. As we saw in §3.2, there is an inconsistency between K-Aim, Knowledge Without Certainty, and EV-Aim (at least when the latter is coupled with some plausible further assumptions about the epistemic values of credences). Appealing to derivative norms will not resolve this underlying inconsistency.

### 4.3 Epistemic satisficing

Let's put one final reconciliation strategy on the table. According to this strategy, inquiry beyond knowledge is often epistemically valuable, but it is always supererogatory. Translated into decision theoretic terms, this proposal amounts to a form of epistemic satisficing. Once Mia and Tess know their respective propositions, they have done well enough, from the epistemic point of view. At this point, they are permitted to cease inquiry. Granted, they could do better: if they were to become rationally certain of their propositions, they would be in a superior epistemic state. Consequently, this approach accommodates the intuition that they have a *pro tanto* reason to inquire further.<sup>24</sup>

In my view, this is the most promising strategy for reconciling K-Aim and Knowledge Without Certainty. Still, not all is smooth sailing. First, note that this satisficing strategy only preserves a weakened version of K-Aim. On the resulting picture, knowledge is only the aim of inquiry for those who aspire to do just well enough, epistemically speaking. For inquirers who are maximizers, it remains true that knowledge is not the goal of inquiry. Consequently, the resulting view will not be strong

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.

If one can know that one knows p without having credence 1 in p, there is no principled reason to think that any number (even infinite) of further iterations of knowledge will entail credence 1 in p. This suggests that an agent could have omega knowledge of p without being rationally certain of p. In such cases, the Omega Knowledge Norm cannot explain why they will have  $pro\ tanto$  reason to consult decisive evidence bearing on p.

 $<sup>^{23}</sup>$  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:

<sup>&</sup>lt;sup>24</sup>Thanks to a referee and Duncan Purves for suggesting this response. Cf. Woodard forthcoming, who defends the view that we sometimes have reason to inquire beyond knowledge, but leaves open the possibility that it is rationally permissible to stop inquiry on attaining knowledge (fn.5).

enough to underwrite some of the claims that proponents of a knowledge-centered conception of inquiry have championed. For example, Friedman 2017 and Sapir and van Elswyk 2021 hold that if you know p, then you ought not continue inquiring into whether p is true. The satisficing strategy will not be able to underwrite this principle.

Even if we are willing to make peace with a weakening of K-Aim, further trouble is in store. Arguably, there are ways of developing our scenarios that elicit the judgment that our protagonists ought to continue inquiry. Imagine that Tess cares deeply about getting r right. And imagine that the expected benefit (epistemic and otherwise) of checking r greatly exceeds the expected benefit of any other use of her time. As noted in §3, there is at least some temptation to say that Tess ought to check whether r is true, at least given these stipulations. Indeed, we might even regard her as criticizable if she failed to check. The satisficing solution is unable to account for this judgment.

A more general concern lurks in the background. Satisficing views in ethics are usually motivated by a desire to avoid the demandingness problem for maximizing versions of consequentialism (e.g., Slote 1985). The underlying intuition is that morality does not require us to constantly make heroic sacrifices—for example, by donating all of our disposable income to charity. The epistemic satisficing solution is very different: it makes no reference to demandingness. According to the present proposal, one is never required to inquire beyond knowledge, even when doing so would yield great epistemic rewards at minimal cost, or even no additional cost. To make this point vivid, consider:

A Choice of Tomes Eva is conducting research at her local library. Her librarian, who is conveniently omniscient, offers her the choice of two books. If she reads from the Book of Knowledge, she will come to know the answers to many of the questions she is investigating. If she reads from the Book of Certainty, she will come to know the very same answers with complete certainty. (Being omniscient, the librarian knows with certainty how reading these tomes would affect Eva's epistemic state.) Moreover, the librarian informs Eva, both books are equally quick and pleasurable reads.

According to the satisficing strategy, Eva ought to read one of the books, but it would be permissible to read the Book of Knowledge rather than the Book of Certainty. This is true even though reading the latter is guaranteed to deliver strictly better epistemic results at no extra cost. This strikes me as rather counterintuitive.

I will leave it as an open question whether some more satisfactory version of the satisficing solution is forthcoming. For our purposes, it's enough that the satisficing approach faces significant hurdles. We have yet to find a problem-free way of reconciling K-Aim and Knowledge Without Certainty.

### 4.4 Looking forward

We tried out three strategies for reconciling K-Aim and Knowledge Without Certainty. All ran into serious hurdles. Since the two principles cannot reside happily together, we must make a choice.

This conclusion is, I think, both surprising and significant. K-Aim and Knowledge Without Certainty 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 these doctrines: which should we keep, and which should get the boot?

# 5 Rejecting Knowledge Without Certainty

Let's start by considering the prospects of rejecting Knowledge Without Certainty. Once we give up Knowledge Without Certainty, we can no longer derive any inconsistency between K-Aim and EV-Aim. Suppose A knows p. If knowledge rationally requires certainty, then 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.

But is this position plausible? Consider how it handles our cases from §2. According to the position under consideration, these cases are impossible, at least as described. Take **Ancient History**. If knowledge requires certainty, then 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, 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 knowledge requires certainty, 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 generalizes beyond this particular case. Conjunctions of the form, "I believe I know p, but I'm not completely certain of p" are typically felicitous. If knowledge requires certainty, then 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. <sup>25</sup>

<sup>&</sup>lt;sup>25</sup>Here I build on Worsnip 2015, which makes a related argument on behalf of fallibilism. Worsnip

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 remark is warranted if knowledge requires certainty; Tess' teacher has no grounds for thinking that Tess has credence 1 in the answer.<sup>26</sup>

Here too, the problem is perfectly general. If knowledge requires certainty, 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. But in practice 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).

Another argument for Knowledge Without Certainty 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 than 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 2020a documents, a wide variety of languages hailing from different language families similarly distinguish between knowledge and knowledge with certainty. If knowledge requires complete certainty, it is hard to make sense of this distinction.

As a further data point, we often claim that someone knows something with "near certainty", implying that the knowledge is not held with complete certainty:

"We now know with near-certainty that Wall Street execs committed felonies." 27

"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." <sup>28</sup>

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.

<sup>&</sup>lt;sup>26</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; Beddor 2020a,b). 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.

 $<sup>^{27}</sup> https://www.businessinsider.com/my-plan-to-finally-make-wall-street-pay-for-its-crimes-2012-thtps://www.businessinsider.com/my-plan-to-finally-make-wall-street-pay-for-its-crimes-2012-thtps://www.businessinsider.com/my-plan-to-finally-make-wall-street-pay-for-its-crimes-2012-thtps://www.businessinsider.com/my-plan-to-finally-make-wall-street-pay-for-its-crimes-2012-thtps://www.businessinsider.com/my-plan-to-finally-make-wall-street-pay-for-its-crimes-2012-thtps://www.businessinsider.com/my-plan-to-finally-make-wall-street-pay-for-its-crimes-2012-thtps://www.businessinsider.com/my-plan-to-finally-make-wall-street-pay-for-its-crimes-2012-thtps://www.businessinsider.com/my-plan-to-finally-make-wall-street-pay-for-its-crimes-2012-thtps://www.businessinsider.com/my-plan-to-finally-make-wall-street-pay-for-its-crimes-2012-thtps://www.businessinsider.com/my-plan-to-finally-make-wall-street-pay-for-its-crimes-2012-thtps://www.businessinsider.com/my-plan-to-finally-make-wall-street-pay-for-its-crimes-2012-thtps://www.businessinsider.com/my-plan-to-finally-make-wall-street-pay-for-its-crimes-2012-thtps://www.businessinsider.com/my-plan-to-finally-make-wall-street-pay-for-its-crimes-2012-thtps://www.businessinsider.com/my-plan-to-finally-make-wall-street-pay-for-its-crimes-2012-thtps://www.businessinsider.com/wall-street-pay-for-its-crimes-2012-thtps://www.businessinsider.com/wall-street-pay-for-its-crimes-2012-thtps://www.businessinsider.com/wall-street-pay-for-its-crimes-2012-thtps://www.businessinsider.com/wall-street-pay-for-its-crimes-2012-thtps://www.businessinsider.com/wall-street-pay-for-its-crimes-2012-thtps://www.businessinsider.com/wall-street-pay-for-its-crimes-2012-thtps://www.businessins-pay-for-its-crimes-2012-thtps://www.businessins-pay-for-its-crimes-2012-thtps://www.businessins-pay-for-its-crimes-2012-thtps://www.businessins-pay-for-its-crimes-2012-thtps://www.businessins-pay-for-its-crimes-2012-thtps://www.businessins-pay-for-its-crimes-2012-thtps://www.businessins-pay-for-its-crimes-2012-thtps:$ 

<sup>28</sup> https://www.nature.com/articles/d41586-019-02754-7

"When business resumes, we know with near-certainty that it will likely be a deluge." <sup>29</sup>

If knowledge requires complete certainty, all of these ascriptions are necessarily false: *near* certainty precludes knowledge.

I have laid out four problems with maintaining that knowledge rationally requires certainty. These problems provide impetus for exploring the other response to our tension: rejecting K-Aim.  $^{30}$ 

# 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.

One prominent view in epistemology maintains that knowledge is subject to a modal condition, such as safety or sensitivity. For example, safety theorists hold that in order for a belief B to amount to knowledge, B must be true at all nearby worlds where it is held on the same evidential basis. We might explore an analogous hypothesis about credal value: the epistemic value of a credence depends not just

 $<sup>^{29} \</sup>verb|https://www.nuix.com/blog/forming-good-remote-work-habits-law-firms|$ 

 $<sup>^{30}</sup>$ Infallibilists may argue that even if their view incurs costs, it is supported on independent grounds. For example, one argument for infallibilism comes from E=K: the view that an agent's evidence at a time is whatever they know at that time (Williamson 2000). Given the assumption that one rationally ought to conditionalize on one's evidence, it follows that everyone rationally ought to assign credence 1 to everything they know. While a full discussion of this issue falls outside the scope of this paper, I think the problems I've raised for infallibilism also cut against E=K. For example, the fact that we distinguish between knowing and knowing with certainty suggests that there is a stronger epistemic position than mere knowledge. A natural way to gloss this distinction is in terms of evidence: if A knows p with certainty, whereas B (merely) knows p, then A's belief that p is better supported by their evidence than B's. But proponents of E=K must reject this diagnosis; according to them, both parties' evidence entails p. For further discussion, see Beddor 2020a, which uses related arguments to motivate the view that an agent's evidence consists in those propositions that are epistemically certain for them.

<sup>&</sup>lt;sup>31</sup>For the canonical defense of a sensitivity condition, see Nozick 1981. For defenses of safety, see e.g., Sosa 1999; Williamson 2000; Manley 2007; Pritchard 2005, 2012; Beddor and Pavese 2020.

on its accuracy at the actual world, but on its accuracy at nearby worlds. On this view, epistemically optimal credences are not only maximally accurate, they are also maximally safe—that is, maximally accurate at all nearby worlds where they are held on the same evidential basis.<sup>32</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, epistemically optimal credences are maximally apt—that is, maximally accurate in virtue of the exercise of a cognitive ability.<sup>33</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." The idea that there is an epistemic status with this profile fits naturally with our arguments for Knowledge Without Certainty 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 candidates for what the aim of inquiry might be, if not knowledge. These candidates are not necessarily rivals. One might maintain that maximal epistemic certainty entails maximal safety or maximal aptness (or both).<sup>35</sup> For

<sup>&</sup>lt;sup>32</sup>For discussion of modal conditions on credences, see Moss 2018; Beddor and Goldstein 2021.

 $<sup>^{33}</sup>$ Sosa also proposes that there are higher grades of knowledge, which require not only that one aptly believes p, but that one aptly believes that one aptly believes p (Sosa 2007, 2015, 2021). Insofar as one follows Sosa in holding that these higher grades of knowledge represent a higher epistemic achievement, one might propose a corresponding view about epistemic value: a maximally valuable credence in p requires having a maximally accurate credence in the aptness of one's credence in p.

<sup>&</sup>lt;sup>34</sup>See Pasnau 2017 for relevant historical discussion. See Beddor 2020a,b for a recent attempt to rehabilitate the notion of epistemic certainty and assign it a starring role in epistemology.

<sup>&</sup>lt;sup>35</sup>On some views aptness itself entails safety; see e.g., Carter 2016. See also Beddor and Pavese 2020, who unpack a virtue epistemological condition on knowledge (maximal adroitness) in terms of safety. One could also explore views on which epistemic value is some hybrid of these statuses; see e.g., the hybrid modal-virtue epistemological views defended by Pritchard 2012; Kelp 2013.

our purposes, we need not choose between these proposals. The important point is that all of these proposals agree that inquiry aims at epistemically optimal credences, which requires having credence 1 in the true answer to the question under investigation. At the same time, these proposals accommodate many of the intuitions and theoretical impulses that made K-Aim attractive.

## 6.2 The Threat of Revenge

Some might wonder whether our earlier arguments against K-Aim can be revamped to undermine EV-Aim. Imagine that Mia is certain of m and Tess is certain of r. As before, suppose that they are given the opportunity to acquire decisive evidence bearing on these propositions. Don't they still have an epistemic reason to inquire? If so, don't we have an equally compelling argument against EV-Aim?<sup>36</sup>

An initial observation: once we stipulate that Mia and Tess are certain of the relevant propositions, some of my arguments that they have reason to inquire no longer apply. In §2 I asked the reader to imagine that Mia rehearses the following line of reasoning:

Mia's Reasoning, Reprised I'm pretty sure that m is true. But I am not completely certain of m, and I'd like to settle the matter. Reading the results of this study might help me figure out for certain whether m is true, or at least come closer to this goal.

If Mia is absolutely certain of m, this reasoning no longer applies. From her perspective, the question,  $Is\ m\ true?$ , is no longer an open question. After all, she does not assign any credence to the possibility that m is false. In assigning credence 1 to m, she treats the matter as completely settled.

A similar point applies to the argument from epistemic decision theory. In §3.3 I used Oddie's theorem to argue that our agents have a *pro tanto* reason to continue inquiring. This reasoning relied on two assumptions: (i) our agents update by conditionalization, (ii) further inquiry might change their credences in the relevant propositions. But if our agents have credence 1 in p, then if (i) is satisfied (ii) is not. If Pr(p)=1, then for any evidence e,  $Pr(p\mid e)=1$ ; certainties can never be lost by conditionalizing on further evidence. Thus if Mia's current credence in m is 1, her credence in m will remain 1 regardless of what she reads. Consequently, unless we make further assumptions about epistemic value (more on this below), we have no reason to hold that the expected epistemic value of her current credence in m is less than the expected epistemic value of her credence in m after reading.

So I think that both the intuitive and the theoretical arguments for further inquiry are less compelling once we stipulate that our protagonists are certain of the relevant

<sup>&</sup>lt;sup>36</sup>Thanks to a referee for raising this issue.

propositions. But for those who share the intuition that our protagonists have reason to inquire further in such cases, all is not lost: there is room to accommodate this intuition in terms of EV-Aim. We saw in §6.1 that EV-Aim can be fleshed out in different ways, depending on how one understands epistemic value. All of the views that we canvassed agreed that credence 1 in a true proposition is a necessary condition on achieving maximal epistemic value. But not all agreed that this is sufficient. Take the view that the epistemic value of a credence depends not just on its accuracy at the actual world, but on its accuracy across a range of nearby worlds where it is held on the same evidential basis. A view like this allows that we may have reason to inquire further into propositions we believe with certainty.




Figure 1: Unsafe Certainties

A toy example: suppose that at world  $w_1$  the butler committed the murder (p). Suppose also that Poirot is certain of p on the basis of the gardener's testimony (t). But Poirot could inquire further by matching the butler's DNA to the blood on the murder weapon (b). Imagine there is a nearby world  $w_2$  where Poirot possesses just t but p is false (a world where the gardener committed the murder and then falsely claimed the butler did it), but there is no nearby world where Poirot possesses both t and b but p is false, because there is no nearby world where DNA evidence can be faked. (See Fig. 1.) The view under consideration predicts that even though Poirot has credence 1 in a true proposition at  $w_1$ , his credence in this proposition is not yet maximally epistemically valuable. It is not maximally valuable because it is not maximally safe—there is a nearby world  $(w_2)$  where he has the same credence on his present basis, but p is false. In such a case, EV-Aim predicts he may have reason to continue inquiry at  $w_1$ . True, further inquiry will not change his credence in p (assuming he updates by conditionalization, this will remain 1). But further inquiry will augment its basis, thereby improving its epistemic value.

The point generalizes. For any view on which epistemic value involves more than credal accuracy, situations will arise where someone has credence 1 in a truth but falls short of maximal epistemic value. In such cases, EV-Aim allows that they may well

have epistemic reason to inquire further.<sup>37</sup>

#### 6.3 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 be maximally confident of the answers.

In response, three points are worth noting. First, this objection cuts equally against those who hold that knowledge requires absolute certainty. If knowledge requires absolute certainty and absolute certainty is seldom attainable, this would show that knowledge is seldom attainable. So attainability considerations will not help us decide between Knowledge Without Certainty and K-Aim.

Second, 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 they 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." 38

"Scientists are absolutely certain that this warming trend is due to human activity." <sup>39</sup>

"Mayor says scientists 'are absolutely certain that a lot of these planets have good conditions for life'." <sup>40</sup>

Now, some might dispute whether these ordinary ascriptions of absolute certainty entail credence 1. However, 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: "She's absolutely certain that this warming trend is due to human activity, but she has some credence that it isn't"—this sounds terrible. Alternatively, some

<sup>&</sup>lt;sup>37</sup>What about cases where someone has credence 1 in a true proposition p, and their credence is maximally safe (and apt, etc.), but they are not certain that they have credence 1 in p? Don't they have reason to inquire further into p? I am happy to grant that this agent may have reason to gather more evidence about p. But it's doubtful that their reason for doing so is that they are trying to figure out the truth of p. (By their lights, there is no possibility that p is false, otherwise their credence in p would be less than 1.) More plausibly, their reason for gathering more evidence about p is to settle whether they are certain of p, not to settle the truth of p.

 $<sup>^{38}</sup> https://www.cnn.com/us/live-news/omicron-covid-19-variant-12-16-21/index.html \\$ 

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

<sup>40</sup> https://www.wtnh.com/news/the-latest-3-scientists-win-nobel-prize-in-physics/amp/

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.<sup>41</sup> But suppose we grant, for the sake of argument, that EV-Aim is seldom attainable.<sup>42</sup> 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 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. Then it might well maximize overall expected epistemic value to move on to the next case.<sup>43</sup>

Now, even if epistemic rationality requires Poirot to move on to the next case, our view does entail that he has not attained the aim of inquiry with respect to his original investigation. Some might balk at this. After all, talk of "aims" seems to set a standard for success or failure. On the face of it, it seems overly harsh to say that Poirot

<sup>&</sup>lt;sup>41</sup>For defenses of the idea that rational certainty is comparatively abundant, see Miller 1978; Klein 1981; Beddor 2020a,b; Goodman and Holguín forthcoming, among others.

<sup>&</sup>lt;sup>42</sup>Note that if we make this concession, then the objection raised in §6.2 loses much of its bite. If we can never rationally have credence 1 in propositions about the external world, then there will never be cases where Mia and Tess (rationally) have credence 1 in the relevant propositions, but still have the opportunity to inquire further.

 $<sup>^{43}</sup>$ As observed in fn.13, one could also endorse a question-relative norm, according to which an epistemically rational agent who is inquiring (or ought to be inquiring) into  $Q_1 \dots Q_n$  does whatever maximizes expected epistemic value with respect to  $Q_1 \dots Q_n$ . If we adopt this question-relative norm, the argument that Poirot ought to move onto the next case is less straightforward. Still, it seems plausible that in any realistic case, agents are (or ought to be) investigating a variety of questions. (There is something defective about single-mindedly pursuing a single question at the expense of all others.) So even a question-relative version of the norm is potentially able to undergird the judgment that Poirot is rationally required to move onto the next case.

has *failed* in his original inquiry.<sup>44</sup> However, a couple of points may assuage this concern. First, I think we should hesitate to say that Poirot has completely *succeeded* in his original inquiry. Poirot has not fully settled the question of who committed the murder. After all, if new evidence bearing on this question comes to light, Poirot would have a reason to consider this evidence. But it is hard to see why this is so if he has fully settled the matter. So insofar as one holds that (completely) successful inquiry into *Q* requires settling *Q* (Kelp forthcoming, 2021b), we have reason to deny that Poirot's inquiry has been a complete success. Second, we could view talk of the "aim" of inquiry as shorthand for an epistemic *ideal*, towards which inquirers should aspire, even if they are—through no fault of their own—frequently unable to reach this lofty standard. While Poirot has failed to achieve this epistemic ideal, he has successfully achieved the secondary aim of coming as close to this ideal as is practically possible, given his circumstances.<sup>45</sup>

# 7 Implications

Adopting EV-Aim has implications for a number of important topics in epistemology.

## 7.1 The dogmatism paradox

First up: the Kripke-Harman dogmatism paradox. <sup>46</sup> 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 in m is epistemically suboptimal. By deleting the email, she is consigning herself to remain

 $<sup>^{44}</sup>$ Thanks to an editor for raising this issue.

<sup>&</sup>lt;sup>45</sup>More revisionary responses to this concern are also available. One would be to adopt a satisficing framework. But, as noted in §4.3, this approach will not help if we stipulate that Poirot is a maximizer rather than a satisficer. Another approach would be to reinterpret the qualifier "as possible" in EV-Aim to refer to what is *practically* possible for the agent. On this view, Poirot has attained the EV-Aim after all. One concern for this approach comes from cases where an agent does not know, or even rationally believe, the answer to *Q*, and there is no inquiry that is practically available that would help answer *Q*. According to the view under consideration, our agent has trivially attained the aim of inquiry into *Q*. But, intuitively, their inquiry is in an important sense incomplete. That said, most of the arguments of this paper are fully compatible with this alternative approach.

<sup>&</sup>lt;sup>46</sup>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; Beddor 2019; Fraser forthcoming; among others.

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 Oddie's theorem, §3). So it is rational to read. This diagnosis of the dogmatist's error follows from the positive proposal advanced here. By contrast, proponents of K-Aim face a real challenge; 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, 2019a draws attention to what she calls, 'interrogative attitudes.' These include attitudes such as *wondering whether* p and *being curious as to whether* p. Friedman 2019a posits a close connection between these attitudes and inquiry, arguing that whenever someone inquires into some question Q, they have an interrogative attitude towards Q.

Friedman also argues that these attitudes aim at knowledge: they are "relieved" when the agent comes to know the answer to the relevant question.<sup>47</sup> This idea has considerable pre-theoretic plausibility. After all, it's natural to describe someone who wonders whether Q as "wanting to know" the answer to Q.

But are all interrogative attitudes relieved by knowledge? Our cases from §2 call this into question. Recall **Mia's Reasoning**. Even though Mia knows *m*, she still has a further desire that is not yet satisfied. She wants to figure out *for certain* whether *m* is true. This desire bears all the hallmarks of an interrogative attitude. It is directed towards a question (*Is m true*?). More importantly, it is an attitude that necessarily involves treating this question as open—as not yet fully settled. Finally, this attitude has the motivational profile of interrogative attitudes: it motivates her to try to settle this question by consulting further evidence and revising her credences in light of this evidence. This suggests that there are some interrogative attitudes that sustain inquiry but are *not* relieved by knowledge—attitudes such as *desiring to find out for sure whether*, or *hoping to conclusive settle whether*. When it comes to these attitudes, nothing shy of certainty will suffice.

Now, it may well be that some interrogative attitudes do not impose such exacting demands. As Friedman 2017 and Sapir and van Elswyk 2021 note, it sounds odd to say, "Mia knows m is true, but she wonders [/is curious] whether it is true." Indeed, this suggests that the positive picture developed here can make room for knowledge to play a role in inquiry after all. Perhaps knowledge is the aim of some interrogative

<sup>&</sup>lt;sup>47</sup>Friedman 2013: 145. Sapir and van Elswyk 2021 also defend a connection between interrogative attitudes and knowledge.

<sup>&</sup>lt;sup>48</sup>Friedman 2019a proposes that a key feature of interrogative attitudes is that anyone who has an interrogative attitude is trying figure something out. Mia's desire to figure out for certain whether m is true satisfies this criterion.

<sup>&</sup>lt;sup>49</sup>Though see Woodard 2022 for discussion of whether this oddity can be explained on pragmatic grounds.

attitudes (in particular, wondering and being curious whether), even if it is not the aim of inquiry more generally.

### 7.3 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.<sup>50</sup>

Assuming Knowledge Without Certainty 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 Knowledge Without Certainty and KN that I am drawing out.

At the same time, this paper advances the critical discussion around knowledge-action norms in two ways. First, it blunts the main defense 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). We have already seen that impurism fails to satisfactorily resolve the tension between K-Aim and Knowledge Without Certainty (§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 more promising alternative to KN. The key is to shift from knowledge to epistemically ideal credences:

<sup>&</sup>lt;sup>50</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.

IDEAL CREDENCE-ACTION NORM If A's credence in a true proposition p is maximally epistemically valuable, then A is permitted to take p for granted in practical reasoning.<sup>51</sup>

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.  $^{52}$ 

## References

Charity Anderson and John Hawthorne. Knowledge, Practical Adequacy, and Stakes. In Gendler and Hawthorne, editors, *Oxford Studies in Epistemology*, pages 234–257. Oxford University Press, Oxford, 2019.

D.M. Armstrong. Does Knowledge Entail Belief? *Proceedings of the Aristotelian Society*, 70: 21–36, 1969.

Bob Beddor. The Toxin and the Dogmatist. *Australasian Journal of Philosophy*, 97(4):727–740, 2019.

Bob Beddor. New Work for Certainty. Philosophers' Imprint, 20(8):1-25, 2020a.

Bob Beddor. Certainty in Action. *The Philosophical Quarterly*, 70(281):711–737, 2020b.

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

Bob Beddor and Carlotta Pavese. Modal Virtue Epistemology. *Philosophy and Phenomenological Research*, 101(1):61–79, 2020.

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.  $No\hat{u}s$ , 42(2):167–189, 2008.

Jessica Brown. *Fallibilism: Evidence and Knowledge*. Oxford University Press, Oxford, 2018. Jennifer Carr. Subjective *Ought. Ergo*, 2(27):678–710, 2015.

J. Adam Carter. Robust virtue epistemology as anti-luck epistemology. *Pacific Philosophical Quarterly*, 97(1):140–155, 2016.

Sam Carter and John Hawthorne. Dogmatism and inquiry. Mind, forthcoming.

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.

<sup>&</sup>lt;sup>51</sup>If we take the optimal credences to be those which are epistemically certain (§6.1), then this is equivalent to a version of the epistemic certainty norm of practical reasoning defended in Beddor 2020b.

<sup>&</sup>lt;sup>52</sup>For valuable feedback, I thank Zach Barnett, Kyle Blumberg, Sam Carter, Peter van Elswyk, Arianna Falbo, Nico Kirk-Giannini, Simon Goldstein, Ben Holguín, Mattias Skipper, audiences at the University of Florida, Goethe University Frankfurt, the London School of Economics, and the Inquiry Network, as well as an anonymous referee and editor at *PPR*.

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

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

Kevin Dorst. Lockeans maximize expected accuracy. Mind, 128(509):175-211, 2019b.

Kenny Easwaran. Expected accuracy supports conditionalization – and conglomerability and reflection'. *Philosophy of Science*, 80(1):119–142, 2013.

Arianna Falbo. Inquiring minds want to improve. *Australasian Journal of Philosophy*, 101(2), 2023.

Arianna Falbo. Should epistemology take the zetetic turn? *Philosophical Studies*, 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.

Carolina Flores and Elise Woodard. Epistemic norms on evidence-gathering. *Philosophical Studies*, forthcoming.

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, 2019a.

Jane Friedman. Checking Again. Philosophical Perspectives, 29(1):84-96, 2019b.

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

Jane Friedman. The Aim of Inquiry. Philosophy and Phenomenological Research, 2023.

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

Simon Goldstein. Omega Knowledge. Oxford University Press, Oxford, 2022.

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.

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

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.

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

Frank Jackson. Internal conflicts in desires and morals. *American Philosophical Quarterly*, 2: 105–114, 1985.

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.

Niko Kolodny and John MacFarlane. Ifs and oughts. Journal of Philosophy, 107(3):115–143, 2010

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

Arie Kruglanski and Donna Webster. Motivated closure of the mind: 'seizing' and 'freezing'. *Psychological Review*, 103:263–83, 1996.

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

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

Barry Maguire and Jack Woods. The game of belief. *Philosophical Review*, 129(2):211–249, 2020.

David Manley. Safety, content, a priority, self-knowledge. *Journal of Philosophy*, 104(8):403–423, 2007.

Adam Marushak. Fallibilism and consequence. Journal of Philosophy, 118(4):214-226, 2021.

Aidan McGlynn. Knowledge First? Palgrave McMillian, New York, 2014.

Emily C McWilliams. Evidentialism and epistemic duties to inquire. *The Philosophical Quarterly*, 73(4):965–982, 2023.

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.

Daniel Műnoz and Jack Spencer. Knowledge of objective 'oughts': Monotonicity and the new miners puzzle. *Philosophy and Phenomenological Research*, 103:77–91, 2021.

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.

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.

Miriam Schoenfield. An accuracy-based approach to higher order evidence. *Philosophy and Phenomenological Research*, 96(3):690–715, 2018.

Daniel Singer and Sara Aronowitz. What epistemic reasons are for: Against the belief-sandwich distinction. In *Meaning, Decision, and Norms: Themes from the Work of Allan Gibbard*, pages 74–94. Maize Books, University of Michigan Publishing, 2022.

Michael Slote. Common-sense Morality and Consequentialism. Routledge and Kegan, Boston, 1985.

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.

Ernest Sosa. Epistemic Explanations. Oxford University Press, Oxford, 2021.

Jason Stanley. Knowlege and Practical Interests. Oxford University Press, New York, 2005.

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.

Elise Woodard. The ignorance norm and paradoxical assertions. *Philosophical Topics*, 49(2): 321–332, 2022.

Elise Woodard. Why double check? Episteme, forthcoming.

Alex Worsnip. Possibly false knowledge. Journal of Philosophy, 112(5):225-246, 2015.