

ARTICLE

Opportunity Costs and Resource Allocation Problems: Epistemology for Finite Minds

Endre Begby

Simon Fraser University, Burnaby, BC, Canada
Email: endre.begby@gmail.com

Abstract

Overwhelmingly, philosophers tend to work on the assumption that epistemic justification is a normative status that supervenes on the relation between a cognitive subject, some body of evidence, and a particular proposition (or “hypothesis”). This article will explore some motivations for moving in the direction of a rather different view. On this view, we are invited to think of the relevant epistemic norm(s) as applying more widely to the competent exercise of epistemic agency, where it is understood that cognitive subjects are simultaneously engaged in a number of different epistemic pursuits (distinct “lines of inquiry”), each placing irreconcilable demands on our limited cognitive resources. In effect, adopting this view would require shifting our normative epistemic concern away from the question of how a subject stands with respect to the evidence bearing on the hypothesis at stake in any one line of inquiry, and over onto the question of how well they cope with the inherent risks of *epistemic resource management* across several lines of inquiry. While this conclusion brings to light important connections between practical and epistemic rationality, it does not collapse the distinction between them. It does, however, constitute a step in the direction of a more systematically developed account of “non-ideal epistemology.”

Keywords: epistemology; non-ideal epistemology; epistemology of inquiry; evidence; cognitive finitude

1. Evidence, Justification, and “Non-Ideal Epistemology”

Epistemologists broadly agree that beliefs are justified (or not) in light of evidence.¹ But how much evidence? Presumably not *all* of it: there is a *lot* of evidence in the world, and no cognitively finite agent could ever hope to have it all, nor would they be able to process it if they did. But at the same time, it seems rather too permissive to say that beliefs are justified (or not) by the evidence that the agent actually has in her possession since an agent might have neglected to gather evidence which

¹This is not to say, of course, that epistemologists are in broad agreement that “evidentialism,” as espoused, for instance, by Feldman and Conee (1985) is the correct view to take. I will have more to say about evidentialism in section 4 of this essay. But meanwhile, notice that most contemporary epistemologists remain “evidentialists” in some suitably irenic sense, namely in holding that *only* evidence can justify belief. (Cf. Nolfi, 2018.) For instance, although evidentialism is often associated with “internalism,” even externalists (ala Williamson, 2000) agree that only evidence can justify belief: the difference is simply that they advocate a comparatively restrictive account of evidence, according to which your evidence consists of all and only the propositions that you know (i.e., “E = K”). Similar things could be said about a development like “pragmatic encroachment,” which is explicitly billed as “anti-evidentialist” (e.g., Fantl and McGrath, 2002). While they deny that justification supervenes on evidence, they do not deny that only evidence can justify belief. Notice, in particular, how the “practical stakes” of the situation does not itself contribute to the agent’s epistemic justification: instead, they help determine the amount or quality of evidence that would be required for the belief to be justified.

would, for all she is entitled to believe, have defeated her justification.² In light of such reflections, it is now quite common to seek a compromise of sorts: beliefs are justified (or not) in light of the *evidence available to the agent*, whether or not she has actually availed herself of it.³ Call this, “the available evidence thesis.”

In my view, this development constitutes an important and timely concession to “non-ideal epistemology,” the approach to epistemology that emphasizes how human agents discharge their epistemic responsibilities under significant situational and cognitive constraints. Nonetheless, it is important to note that this compromise still proposes to measure justification in terms of a subject’s standing with respect to the evidence bearing on particular, individual hypotheses. However, it is characteristic of the human epistemic situation that we are typically—and as I will argue, maybe even always and necessarily—engaged in multiple concurrent lines of inquiry. Of course, we can *prioritize* one line of inquiry over others, and this is precisely the fact whose epistemological consequences my alternative approach attempts to bring to the fore. This article will argue that, when we are asking questions about epistemic justification, what we are asking about is, in no small part, whether cognitive subjects are making reasonable prioritizations in managing their limited epistemic resources, i.e., whether they are responding well to the *epistemic resource allocation problem*.

The epistemic resource allocation problem arises as a consequence of the fact that any inquiry involves *opportunity costs*:⁴ given our cognitive limitations, the pursuit of evidence along one line of inquiry will necessarily involve siphoning resources away from other (possible) lines of inquiry. As such, any decision to prioritize one line of inquiry over another will leave potentially significant evidence unaccounted for, evidence that nonetheless remains *available* in any reasonable sense of the term. On the standard view, any such (knowing) neglect of available evidence would jeopardize a subject’s justification for believing the provisional outcome of the deprioritized lines of inquiry. I will argue that this is the wrong result: we want our epistemic norm(s) to put agents in a position to make *informed decisions* about where their limited epistemic resources are best deployed at any time, even as this necessarily entails neglecting available evidence. Non-ideal epistemology is in part defined by the task of providing substantive, normative guidelines for good epistemic resource management. I argue that a plausible guideline of this sort could be developed in terms of a notion of the *relative expected informativeness* of different lines of inquiry, leading to rather different—and to my mind, rather more relevant—prescriptions for the exercise of epistemic agency.⁵

²Cf. Harman, 1973 for a classic case, or, more recently, the concept of “normative defeat” as discussed in Lackey, 2008 and Goldberg, 2017, among others. Some high-profile evidentialists resist this move and argue that justification should only ever be measured in terms of the evidence that the agent actually has in her possession (e.g., Feldman & Conee, 1985). For recent criticisms of this sort of view, see McWilliams, 2021, pp. 7170–7171; Flores & Woodard, 2023. I will be adding my own—rather different—argument for why I think this is an unsustainable position in section 4 of this essay.

³Cf. Kelly, 2006/2014: “Given the thesis that evidence is that which justifies belief, one’s intuitions about the evidence that is available to an individual in a hypothetical scenario will shape one’s views about what the individual would be justified in believing in that scenario.” This specific focus on available evidence seems mostly to be taken as a matter of course in contemporary mainstream epistemology but is frequently foregrounded as a matter of concern in discussions of the interrelations between moral and epistemic responsibility, e.g., Zimmerman, 2008, Sher, 2009.

⁴See Rescher (1989) for an early account (drawing on Peircean inspirations), and Dallman (2017) for a more recent perspective.

⁵It might seem that my position has certain affinities with reliabilism (cf. Goldman (1979) for a canonical development). To be sure, reliabilism also denies that justification generally supervenes on particular bodies of evidence, insofar as it holds that evidence may not be necessary for justification at all. Even so, we note that Goldman lists several types of recognizably evidence-based reasoning as *instances* of “reliable processes” that can bestow justification on belief (e.g., perception or inference, as contrasted with guessing or wishful thinking). In these cases, the lead question of this paper arises again, namely *how much evidence is required to justify belief?* As before, the answer cannot be, “all the evidence.” By contrast, an answer like “just as much as is required to justify belief” is not only unhelpful, but, as I will argue, also deeply misleading: simply, we should not seek to answer this question in terms of the subject’s standing with respect to the evidence bearing on individual propositions. Instead, we should try to answer it in terms of how epistemic agents distribute their epistemic resources over several lines of inquiry. As I

Here is an overview of the argument to follow: in Section 2, I will put pressure on the available evidence thesis by showcasing examples in which a subject's neglect of available evidence does not appear to undermine their epistemic justification. In Section 3, we will up the ante by looking at cases in which agents *knowingly* neglect readily available evidence but appear to be well within their epistemic rights to do so. When an agent knowingly neglects available evidence for one hypothesis, it may be because she makes the epistemically responsible decision to prioritize inquiry along a different line. When such decisions are well-made, agents are due a measure of epistemic praise, not blame. In Section 4, I argue that the observation that epistemic agents are typically engaged in multiple, concurrent lines of inquiry does not simply express some aspirational vision of human epistemic flourishing, i.e., that we are better off as epistemic agents if we cultivate broad interests rather than focusing narrowly on a single inquiry. Rather, it is internal to the logic of inquiry itself and arguably constitutes a necessary feature of any kind of epistemic agency. This sets the stage for comparing our approach with recent discussions about the "epistemology of inquiry" (e.g., Whitcomb, 2010; Friedman, 2017) and also for mounting an effective response to traditional evidentialists (e.g., Feldman and Conee, 1985) who resist the idea that justification is ever sensitive to evidence not currently in our possession. Additionally, I argue that this is not just another instance of "pragmatic encroachment" on epistemology: it is certainly true that practical stakes sometimes help determine the choice of which inquiry to prioritize; but in other contexts, we are required to make these choices on more clearly epistemic grounds. (In other words, while my argument's emphasis on *agency* does help bring to light important connections between practical rationality and epistemic rationality, it does not collapse the distinction between the two.) In Section 5, I explain how these dynamics cannot be straightforwardly subsumed under familiar epistemic norms of maximizing truth (e.g., Kornblith, 1983) or knowledge (e.g., Williamson, 2000). Instead, I argue that they should be thought of in terms of strategies for maximizing the *overall expected informativeness* of one's inquiry: "informativeness" is now understood to distribute over multiple lines of inquiry taken as a whole. There is no plausible way to disaggregate this notion so as to produce epistemic norms uniquely dictating how one should conduct oneself in the several individual lines of inquiry that one is engaged in. In Section 6, I offer my final conclusions and tie them together with current reflections about the importance of pursuing a program of "non-ideal epistemology."

2. The Available Evidence Thesis

From within a certain philosophical frame of mind, the available evidence thesis has an air of a truism about it: it is something that seems like it just *has to be true* somehow. But on the other hand, it turns out to be extremely tricky to say exactly what is entailed by it, i.e., to say what it means for a piece of evidence to be "available" in the relevant sense.⁶ Maybe this explains why, even among philosophers who seem disposed to endorse the claim, the thesis is rarely put forward with any real sense of enthusiasm. In short, the available evidence thesis seems *dubiously helpful* even if it is true. The following line from Michael Zimmerman may be seen to capture the mood perfectly: "We must distinguish between evidence that is *available* to someone and evidence of which that person in fact *avails* himself. Available evidence is evidence of which someone *can*, in some sense, and *ought*, in some sense, to avail himself. I confess that the exact sense of this 'can' and 'ought' still elude me" (Zimmerman, 2008, pp. 35–36).

Certain cases may seem straightforward indeed. Drawing on Holly Smith's classic discussion of culpable ignorance (Smith, 1983), consider the pediatric doctor who fails to read up on the latest

argue in Section 5, there is no way to uniquely quantify such an answer in terms of the number of true or justified beliefs that an agent should expect to sit back with at the end of the day.

⁶For an early catalog of difficulties, see Lycan, 1977, responding to Harman, 1973.

medical research and therefore continues to believe that a certain method of treatment—now discredited—is the right choice for a particular condition. (We can imagine that the journals in which this research is conclusively presented sit proudly on display in his office, spines uncracked.) It is one thing to note that he is thereby at risk of prescribing a harmful treatment to his patients, and so, that there is also a *moral* dimension to our evaluation of his ignorance. But even before we get around to considering possible harm to other people, it seems like we would also say that he is a poor *epistemic* agent, and that his belief in the appropriateness of this line of treatment is epistemically unjustified simply because he is neglecting available contrary evidence. (10 years previously, let us say, he might have been justified in his belief even if the treatment was no less harmful, just because the relevant evidence was not yet available.)

But “available” might turn out to be a vague predicate, and once we start digging deeper we will soon enough run into cases that are less clear. For nearly two centuries, the scientific consensus was that all earthly life is supported by photosynthesis. This view was not entirely unopposed: chemosynthesis—energy production by the oxidization of inorganic compounds, such as hydrogen sulfide—was first proposed as a scientific possibility in the late 19th century.⁷ However, clear evidence to support the hypothesis was not gathered until the 1970s, when someone built a submersible capable of venturing down to deep-sea hydrothermal vents at the Galapagos rift. So, in a relevant sense, the evidence for chemosynthetic life forms on the ocean floor was presumably always available. It was, literally, always there, in precisely the place where we eventually found it.⁸

Similarly, it was not so long ago that it was considered an open question in the scientific community whether the universe is constant and unchanging or whether it is currently expanding following an initial Big Bang. In 1965, a project at Bell Labs involved the construction of a super-sensitive radio antenna. Massive efforts were made to eliminate all sources of background interference. Despite their best efforts, they failed to eliminate a persistent and ubiquitous source of noise. Who knows how long they might have gone on trying to eliminate local sources of interference—including cleaning up the constantly replenished supply of bird doo—were it not for a near-simultaneous theoretical development, namely the prediction that if the Big Bang theory were true, there should be a ubiquitous cosmic background radiation detectable in the microwave spectrum. In other words, a noise long known to radio astronomers everywhere—and presumed to be just that, mere noise—turned out to be the evidence that precipitated a radical and sweeping scientific discovery, overturning the competing theory of the steady-state universe.⁹

These are examples in which false beliefs persisted despite the availability of the sort of evidence that would eventually overturn them. However, it does not seem outrageous to suppose that these scientists might nonetheless have been justified in continuing to believe as they did. If this is right, it puts real pressure on the idea that epistemic justification stands or falls with how well one’s belief fits the available evidence.

It is possible, though, that sophisticated articulations of the available evidence thesis can accommodate these kinds of examples. Let’s look again at Zimmerman: “Available evidence is evidence of which someone *can*, in some sense, and *ought*, in some sense, to avail himself.” Maybe our first example fails the “can” test: the evidence was not really “available” (in the relevant, technical sense) until someone built a submersible capable of traveling down to these deep-sea hydrothermal vents. It is easy to feel a twinge of sympathy here. But on the other hand, there would have been no decisive barrier to building the technology earlier. It cannot be, in general, that we should not classify evidence as available until someone can physically gather it, if no one is doing anything to put themselves in a position to gather it.

⁷Cf. Ackert, 2006.

⁸But not *only* there: as Dubilier et al. observe (2008, p. 725), “[i]t is surprising in retrospect that it took the discovery of these symbioses in the deep sea for scientists to realize that chemosynthetic symbioses occur worldwide in a wide range of habitats, including easily accessible habitats, such as shallow-water coastal sediments.”

⁹For an account, see Hacking, 1983, pp. 159–160.

The second example might illustrate a different concern, more closely related to the “ought”-test that Zimmerman proposes. The reason these scientists’ justification for believing in the steady-state universe was not in jeopardy, despite the ready availability of the evidence that would eventually overturn the theory, is that they were not in a position to recognize it *as* evidence or as evidence bearing on this particular hypothesis: this recognition required a theoretical development which was yet forthcoming. In this sense, the example could be taken to illustrate a point which is articulated in Harman (1973) and also more recently in Flores and Woodard (2023): the fact that some agent fails to avail themselves of some piece of available evidence is epistemically consequential only if they had reason to believe that the evidence was there (or that this evidence would have the epistemic significance that it did). This is precisely what they did not have, which is why their justification was not (yet) in jeopardy.

Again, I like the general shape of this answer. But it is not without its problems: in general, the narrowness of someone’s intellectual horizon, their lack of ingenuity, and their failure to consider certain epistemic possibilities are not usually good inputs to a narrative about why they are justified in believing as they do. At best, one would want to say that these observations could contribute to a story about why they might be *excused* for failing to avail themselves of this evidence, not why they would be justified in continuing to believe as they did.¹⁰

Now, perhaps everyone will agree that the available evidence thesis, despite its intuitive appeal and truistic appearance, has significant problems. Despite its recognizable vagueness, it could still be “largely correct” or at least “on the right track” or “pointing in the right direction” even if we struggle to articulate precisely the conditions under which it holds.

We can afford to leave this question unresolved here because I think there also exist more telling cases in which a subject will *knowingly* neglect available evidence *which they have every reason to think* could potentially overturn their beliefs, but whose justification for continuing to believe as they does not seem to be jeopardized as a result.

3. Knowingly Neglecting Available Evidence

To my mind, then, the deeper problem with the available evidence thesis is not that “available” is vague and resists precisification (along the “can”-dimension), or that it is already caught up in normative notions (along the “ought”-dimension) and therefore cannot be used to properly ground an account of epistemic justification. Instead, the deeper problem is that it remains locked into the supposition that epistemic justification is a notion that supervenes the relation between a cognitive subject, some body of evidence, and a particular, individual proposition (or “hypothesis”).

To see how we might motivate moving beyond this focus, consider another example. In 1919, Eddington and Dyson famously staged a spectacular experiment to conclusively demonstrate the truth of Einstein’s General Theory of Relativity. Einstein’s theory predicted that the angular deflection of light rays by the sun would be 1.74 arc seconds, or about twice what was predicted by Newtonian gravitation. Dyson determined that a solar eclipse that would occur on May 29, 1919, during which the sun would pass across the Hyades star cluster, could serve as an excellent occasion to test this prediction. In order to carry out the experiments, equipment was shipped to two remote

¹⁰It is worth emphasizing the “in general” here. Maybe it is true, for instance, that Kuhnian “normal scientists” working within a firmly entrenched paradigm are sometimes justified in retaining their beliefs by dismissing new evidence as anomalous (cf. Kuhn, 1962). But even here, the picture is complicated: on the one hand, they might have broadly Bayesian grounds for dismissing as “improbable” evidence that someone with a different background theory might well take to be probative. On the other hand, even as Kuhn might say that there was no *unique point* at which, for instance, Priestley’s continued adherence to phlogiston theory became irrational, he might also say, looking back over the larger trajectory of Priestley’s career, that he should eventually have recognized that what he had discovered was not “dephlogisticated air,” but rather oxygen.

locations, the West African island of Principe and the Brazilian town of Sobral. The results were widely publicized and broadly hailed as the conclusive demonstration of Einstein's theory.¹¹

We can certainly agree that the results broadly favored Einstein's prediction. Even so, the observations were apparently less conclusive than they might have hoped. Several cameras were out of focus, and some delivered results closer to the Newtonian range. The point to note, though, is that this apparently did not move Eddington and Dyson to suggest that they should take the next opportunity to perform the experiments again, hopefully with better luck next time. Instead, they were already strongly convinced (on theoretical and empirical grounds) that the theory was right. Although they were presumably open to the possibility that further experiments *could* yield results inconsistent with General Relativity, they determined that the 1919 results were "good enough" for the purpose, even though they could no doubt be improved. Although, in a real sense, further evidence was there to be had—"available," even if they had to wait until the next convenient opportunity to avail themselves of it—they decided to forego that opportunity. Presumably, they decided to forego it because they understood that pursuing this opportunity would be a waste of resources that would be better spent developing and testing other aspects of the theory.

To my mind, this seems like a *good decision*.¹² Even though they knowingly decided not to collect further available evidence—evidence which *could* have overturned their belief—they were justified in their decision, and, I think, justified in believing that GTR was correct.¹³

Nor do we need the distance that history provides to see this dynamic in action. Contemporary archeologists face a similar problem. No one seriously doubts that Lidar technology might provide gold-standard evidence. But its deployment is also prohibitively resource-intensive.¹⁴ Ideally, one might assume, archeologists would like to Lidar-map pretty much all of the earth's surface. But that would be ridiculous. Given the costs of implementation, one has to be selective.¹⁵ Let's say one is interested in determining the precise location of a particular Roman encampment somewhere in contemporary Greece. Historical sources suggest somewhere in the region of Thessaly. If you knew the exact location to turn the Lidar loose, it would be a no-brainer: it would certainly be much cheaper (and less environmentally impactful) than using shovels. But at the same time, the area might be so large that it is too costly to map all of it, especially if one has access to decent satellite

¹¹As New York Times reported in its headlines on November 10, 1919: "LIGHTS ALL ASKEW IN THE HEAVENS [...] Einstein theory triumphs."

¹²A fuller story would of course acknowledge that other research teams did in fact subsequently perform similar experiments, and that the results were happily folded into the evidence base for GTR. Even so, the point to note here is that few people in the Einstein-camp would have held their breath for these later results to obtain. (In fact, one version of the story has it that Eddington was already sufficiently convinced of GTR prior to 1919 that he debated whether even the initial experiments were worth undertaking, given the technical and logistical challenges involved and the chances of obtaining a "null result" (Eddington, 1919, p. 122, cited in Chandrasekhar, 1976, p. 250). (See Kennefick, 2011 for discussion.))

¹³Moreover, I hope that this example will help put to rest a lingering concern that some readers might have, namely that my account retains an unhealthy focus on *individual subjects* (and on individual epistemic performance) which otherwise seems at odds with the broader push toward non-ideal epistemology. In general, I have a fair degree of sympathy with this type of concern. But I want to say two things in response: first, I do think that individuals (and norms designed for assessing individual performance) will continue to play an important role even in a fully developed (and appropriately "socialized") non-ideal epistemology; second, and relatedly, I think it is a mistake to assume that the "epistemic subjects" to which these norms apply are always going to be human individuals. For instance, when a scientific working group is suitably integrated, it is really no stretch to consider it as an epistemic subject in its own right, and to assess its performance with respect to epistemic norms that are not fundamentally different from those that we apply to human individuals. In particular, although the group will presumably have greater cognitive resources to draw on than do the individual members that compose it, it nonetheless remains a recognizably "finite cognitive subject," and as such, subject to the epistemic resource allocation problem. In this vein, note how the present example is clearly not about Eddington and Dyson *qua* individuals but rather about "Eddington and Dyson (plus associates)" *qua* integrated working group. We can ask whether *this group* made good on the epistemic resource allocation problem for the purposes of vindicating and deepening our understanding of GTR.

¹⁴Cf. Keller, 2013.

¹⁵As Keller, 2013 points out, it would also produce more data than anyone would be able to process.

imagery. Satellite imagery, under favorable conditions, can provide good evidence, though obviously less good than Lidar (or shovels) might provide. While you would love to have the evidence of Lidar, and there is no serious question that the better evidence is “available,” the opportunity costs involved in pursuing such evidence might be prohibitive. At this point, you might well say, historical sources indicate that the encampment is somewhere in this area; satellite imagery is favorable but inconclusive. We could verify this with Lidar technology, but we won’t: our current evidence is “good enough” as is.

In these sorts of cases, scientists may acknowledge that they are setting aside lines of inquiry which *could* produce contrary evidence. In this sense, it is hard to deny that they would be *knowingly neglecting available evidence*. Yet this knowing neglect of available evidence does not, in any obvious sense, seem to undermine their justification for believing in the theory.

If this is right, we need to rethink how we approach the relationship between cognitive subjects, their beliefs, and their evidence. If these examples point in the right direction, it cannot be that epistemic justification for believing in some proposition *p* is straightforwardly measured in terms of the subject’s having availed themselves of (i.e., gathered and processed) all the available evidence. Moreover, their failure to avail themselves of this evidence is not a matter of their inability to recognize that the evidence was there to be gathered or their failure to understand that what was there to be gathered was in fact evidence. Rather, it was a neglect of evidence which they knew was available; moreover, evidence which, they would freely acknowledge, *could* have told against their favored theory.

How should we accommodate this fact? Here’s a rough outline of the story that I propose to tell. The traditional account portrays epistemic justification as a concept that supervenes on the relation between a cognitive subject, some evidence base, and a particular, individual proposition *p*. However, epistemic agents are characteristically engaged in multiple concurrent lines of inquiry. Each of these lines of inquiry imposes irreconcilable demands on our limited epistemic resources. Any kind of epistemic activity imposes *opportunity costs*: in particular, resources spent unearthing (and processing) evidence relating to the hypothesis at stake in one line of inquiry will necessarily entail siphoning resources from some other line of inquiry, as a result of which *some* available evidence will go uncollected, no matter what one does.

We could stick to our guns and argue that they must thereby stand to lose their justification for believing the provisional outcome(s) of the deprioritized line(s) of inquiry. I think this would be the wrong view to take. It would entail, quite generally, that one could never hope to improve one’s epistemic standing with respect to the hypothesis at stake in one line of inquiry without simultaneously jeopardizing one’s epistemic standing with respect to the hypotheses at stake in several others. On the contrary, we should praise agents who make good decisions about how to allocate their limited epistemic resources, even as they do so in the knowledge that they thereby forego potentially significant evidence. If this is right, then we should think about epistemic justification *not* as supervening on the relation between an epistemic subject (whether singular or plural), some body of evidence (whether broadly or narrowly construed), and a particular proposition under inquiry. Instead, we must think of it in terms of how well the subject copes with the inherent risks involved in *epistemic resource management* across these multiple lines of inquiry. Clearly, these decisions can be better or worse made from an epistemic point of view.¹⁶ Quite simply, some lines of inquiry are more likely to turn up the kind of evidence which will, in the larger scheme of things, serve to confirm or disconfirm the theory. This is not to say that the evidence one might gather

¹⁶For an example of *bad* epistemic resource management, we can surely criticize researchers who presume to draw significant inferences about human psychology from rat studies. Crucially, such criticism need not claim that (i) the rat studies do not provide evidence, or that (ii) human studies would not come at a cost. Instead, what we are saying is that, given the significance of the conclusions you are proposing, the cost of the human studies is one that we ought to bear, given the comparative quality of the evidence that such studies are (almost) certain to provide. *Not* conducting human studies here is tantamount to culpably neglecting significant evidence. (These remarks are inspired by O’Connor, 2021.)

along some other line of inquiry is not worth having. Rather, it is to say that it is *less worth having*, at least so far as we are currently in a position to tell, than some evidence that one might hope to gather along another line of inquiry. Or one might think that this is indeed the evidence that we would ideally hope for. But the chances of actually obtaining it are sufficiently slim that we will not even try.

Epistemology for finite minds turns essentially on such zero-sum decision-making: a good epistemic agent is attuned to the fact that every line of inquiry involves some opportunity cost, and that they are thereby continuously faced with an *epistemic resource allocation problem*. They understand that there is no epistemic decision they could make which would not leave some potentially significant evidence uncollected and/or unprocessed. The proper measure of epistemic justification, I argue, is how well they face up to this problem, not how they stand with respect to the available evidence in any particular line of inquiry.

These reflections help explain why *sometimes* (maybe even often) agents are in fact epistemically at fault for failing to pursue available evidence. But importantly, they also contribute an explanation for why these agents sometimes (maybe even often) can perfectly well retain their epistemic justification despite neglecting such evidence. We still have work to do, however, in determining what speaks in favor of such a view and exploring its philosophical consequences.

4. Comparison with Current Alternatives

In the previous section, I offered the observation that epistemic subjects are “characteristically” engaged in multiple, concurrent lines of inquiry, each placing irreconcilable demands on our limited cognitive resources. In this section, I will go further and claim that this is not merely an incidental observation about human beings as epistemic agents—that we have broad interests and wide intellectual curiosity, and that, up to a limit at least, we should be praised for indulging this curiosity. Rather, it is a necessary truth about any relevant sort of inquiry involving cognitively limited agents operating in a complex environment.

Quite simply, for agents such as us, no epistemically responsible inquiry could ever be an inquiry into a single hypothesis to the exclusion of everything else. That is, competent, epistemically responsible inquiry into any one hypothesis will also—necessarily—involve a relevant degree of sensitivity to the outcome of other (actual or potential) inquiries as well. The resource allocation problem now arises in the form of the question, to what extent should one concurrently pursue these subsidiary lines of inquiry?

To see where this is going, note how someone might argue that the inquiry-related considerations I’ve marshaled so far really have no bearing on the traditional focal question of normative epistemology, namely what should one believe, given the evidence currently in one’s possession? To this end, notice, for instance, how traditional evidentialists like Feldman and Conee (1985) would resist even the first step of the line of reasoning that I offered in this paper: i.e., they would argue that the justification of belief is *never* sensitive to a larger body of evidence than what the agent has in her possession at the time in question. Any normative notions we might harbor about what agents should do with the evidence that they *don’t* have belong to a different problem field altogether, for which they propose the catch-all term “the ethics of belief.”¹⁷

But this view cannot be sustained. In adopting this view, traditional evidentialists entitle themselves to the assumption that evidence is in an important sense *given*—that we can always know whether what we have is evidence; moreover, that we are always in a position to know the epistemic significance of this evidence. I think this assumption is demonstrably false: it is clear that

¹⁷Cf. Dougherty (2014).

we are sometimes *not* in a position to know whether a particular observation is evidence (as opposed to mere noise), and if it is evidence, what its epistemic significance should be.¹⁸

For an example, we may turn to the early records of research into Covid-19. One observation was that the virus seemed not only to affect young people less than other demographics but also to spread less easily among them. This observation was crucial to policy decisions in many jurisdictions, for instance, in determining whether schools should be kept open. However, in July 2020, Florida's Department of Health released past-week statistics indicating a test positivity rate among subjects under 18 well in excess of 30%, or about twice what they found in adult demographics. This certainly *would* have constituted startling evidence of widespread, hitherto undetected transmission among youth were it not for the fact that the report was subsequently retracted due to a coding error: quite simply, large amounts of negative test results had been left out. The following week's report gave the much more sensible number of 13.4%.¹⁹

One could say, it is shocking that public health officials didn't catch the mistake prior to publication, and even more shocking that it would take them a full week to own up to the mistake. But our task is to determine what ordinary agents should do at the particular point in time that the report was made available to them.

Transporting ourselves back to this time and place, we can ask, how should an epistemic agent respond to this report? According to Feldman and Conee, S is justified in believing *p* at time *t* only if *p* "fits the evidence S has at *t*" (1985, p. 15). Presumably, then, there is a sense in which as soon as we were in possession of the Florida DoH report, our belief in comparatively low rates of transmission among youth would no longer "fit" our evidence. So, we would have lost our justification for believing *p*, and therefore should, epistemically speaking, cease to believe it.

But this strikes me as extraordinarily naïve. One's first response to the DoH report *should* be to question its status *as* evidence: it is so far beyond what we had reason to expect that it is more likely that the observation is false than that our belief is false.²⁰ This doesn't mean that it *has* to be false, of course: so it is not like we can rationally dismiss it without further ado. Instead, it gives us a job to do, namely that of establishing whether it *is* evidence (as opposed to noise). How does one go about this? By launching a separate line of inquiry, specifically an inquiry into the reporting practices for Covid cases among youth at county level in Florida. This is precisely the inquiry which then subsequently established that the reported number was false by a significant margin.

This example shows that in many cases, our epistemic right to take what is presented to us as evidence to be evidence is dependent on the outcome of a separate inquiry. Even when we are considering what looks from a distance like it should be just a single line of inquiry, it will turn out that we must be prepared to pursue sublines of inquiry as part of it. Further, notice that I have so far just considered the process of establishing observations *as* evidence (Y/N). Perhaps it is rare that observations are so decisively recanted as in the Florida case. But in most situations, even accepting something as evidence still leaves us with the job of determining what the precise significance of the evidence is. And this, again, would depend on the outcomes of further inquiries: for instance, we could imagine finding that the inflated number wasn't an artifact of a coding error per se, but an extrapolation from a smaller, demographically specific sample (say, a kids' Bible camp in some remote area). Now, we would have to figure out just how representative is this sample? Again, this takes further inquiry along separate lines. We cannot, in other words, simply assume that evidence is given, or wears its epistemic significance on its sleeve.

If epistemic agents could responsibly just assume that what is given to them as evidence *is* evidence, and that its evidential significance is apparent, the traditional evidentialist picture would

¹⁸Of course, contemporary externalists also argue, though from a very different starting point, that epistemic agents are not always in a position to know what their evidence is. My argument here draws on different philosophical motivations and does not depend in any way on this externalist line of reasoning.

¹⁹Cf. Wamsley, 2020.

²⁰This is, of course, a general lesson about inductive reasoning that Hume taught us long ago (Hume, 1748).

be in the clear. But in reality, it is not so. Epistemic agents such as us must reckon with their fallibility and cognitive finitude. One manifestation of this finitude is that it takes further inquiry to know what we should do, epistemically speaking, with the evidence in our possession at a given time. In short, I do not accept, and I do not think anyone should accept, a fine line dividing the epistemology of inquiry and the epistemology of belief. In most situations, any epistemologically responsible belief involves sensitivity to multiple open lines of inquiry, whether these lines of inquiry are currently engaged in or not. (Again, this is a matter of epistemic resource allocation.)

To be sure, we can imagine cases in which the evidentialist norm would seem more readily applicable. Suppose you want to know how much money is in your savings jar. To determine this, you conduct an inquiry: find an appropriate surface, pour the contents of the jar out, and proceed to counting. In this case, you might reasonably believe that your evidence is complete (supposing you took measures to ensure that no coin rolled off the counter); moreover, each bit of evidence is discrete and unambiguous (supposing you took measures to exclude the possibility that some of the coins were, say, US quarters rather than Canadian quarters). On these assumptions, and provided you've counted correctly, there should be no further question as to which belief best "fits" the evidence that you have. I won't deny that such cases exist. But at the same time, they are not particularly representative of belief formation "in the wild," where typically, our evidence is incomplete, ambiguous, and dependent on other bits of evidence.

We cannot, then, in general divide the epistemic process into an inquiry stage (where we gather and process the evidence) and a belief-formation stage (where we decide what to believe on the basis of this evidence).²¹ What you *should* believe on the basis of your (current) evidence is, in typical cases, provisionally dependent on the outcome of further evidence-seeking inquiries, the results of which may yet be indeterminate or not forthcoming on a time scale that you can live with.

I'd be happy to think of these arguments as pushing us in the direction of a normative epistemology centered on "inquiry" rather than "belief." But notice how much of the recent resurgence of interest in this area (e.g. Friedman, 2017, 2019) is still very much structured around considering lines of inquiry on a one-by-one basis. The lead question in this literature is simply whether taking up an "inquiring attitude" toward some *p* is consistent with concurrently "believing that *p*." Friedman's answer is "no": taking up a genuinely inquiring attitude with regard to *p* requires "suspending belief" in *p*. It would belong to a different paper to assess whether this is the correct view.²² For now, note simply that Friedman's reasoning still applies only to single, individuated lines of inquiry. By contrast, I have advocated a perspective on which the epistemological focus should be on how subjects distribute their cognitive resources across many lines of inquiry.

Another concern that might arise at this point is that the sort of agent's perspective that I have emphasized is really just another instance of "pragmatic encroachment" on epistemic norms, as described for instance in Fantl and McGrath (2002). Certainly, there is some overlap: for instance, though the point is not typically phrased in these terms, it is a consequence of pragmatic encroachment that someone might be justified in believing that the train will stop in Foxboro even though they have not gathered and processed all the available evidence. According to pragmatic encroachment, the reason they are justified is, broadly speaking, that the evidence currently in their possession is "good enough" for the practical decision they are faced with, and that collecting the further evidence would either be (i) pointless, given that they already *know* (to the extent that their situation requires) that the train will stop in Foxboro, or (ii) that collecting the further evidence will jeopardize their ability to board the train in time.

²¹See also McWilliams, 2021, pp. 7191–7192 for similar reflections on the role of active epistemic agency in "constituting evidence."

²²Preview: I don't think it is. It is certainly fair to demand that inquiring agents should be "open" to any outcome of the inquiry. But this openness does not require suspension of belief.

I would be happy to fold these kinds of cases into the general argument for the position outlined in this paper. But not all inquiry decisions are made under pragmatic pressures of this sort. It is of course true that we can imagine a variety of pressures weighing in on such decisions, and that it is by no means clear that the “all things considered right choice” will always be the epistemically optimal choice. So, for instance, we can imagine a decision to prioritize a line of inquiry being made because of its urgency (say, related to the policy decisions required under the Covid-19 pandemic). Or we could imagine a scientist prioritizing a particular line of inquiry simply because it is more likely to receive funding. But equally, we can abstract away from these sorts of concerns and imagine the choice being made primarily on epistemic grounds: we do right by the relevant epistemic norms when we allocate our epistemic resources in a way that maximizes our expected return on investment.²³ Consider, for instance, the controversy surrounding the Human Genome Project (1990–2003). No one doubts the potential epistemic value of a full mapping of the human genome. No one doubts that the Human Genome Project generated an enormous amount of funding, thereby potentially bringing about a significant amount of empirical insight we might not otherwise have had. But we can still wonder what else we might have accomplished with the epistemic resources that were devoted to this project. This question turns crucially on the opportunity costs of inquiry and the problem of epistemic resource allocation. While these issues certainly intersect with practical decision-making (if only because budgeting decisions are involved), we can also construe the problem specifically on its epistemic merits: the problem with the Human Genome Project was not that it cost a lot of money; rather, the problem was that it distracted us from other, potentially more valuable epistemic pursuits.²⁴ In this sense, it is true that my argument has served to bring considerations of practical rationality and epistemic rationality into closer register than is common in traditional approaches to epistemology. But even so, it does not collapse the two, since in many cases, we should have no problem telling apart the considerations that count *epistemically* in favor of a particular resource allocation decision and those that count in its favor *merely prudentially*.²⁵

If this is correct, our next question must be, what is this epistemic value that we should seek to maximize in epistemic resource allocation? How does it relate to more traditional conceptions of the epistemic value of truth or knowledge? Is it possible that the kind of norm that I have been speaking of could be subsumed under these more traditional norms?

5. Truth, Knowledge, Informativeness

The previous section ended with the suggestion that justification should be construed in terms of decisions which maximize some expected epistemic value. We have not yet said anything to determine what that value might be. This leaves open the possibility that my line of argument might after all be subsumed under more traditional epistemic norms, such as the maximization of truth or knowledge. So, it might be useful to take a moment to see why that strategy will not work.

²³It would be tempting to say “maximizes expected epistemic utility,” but for the fact that “epistemic utility theory” is already taken by a program which, once again, looks primarily at single lines of inquiry (cf. Pettigrew, 2011/2019, 2016).

²⁴Cf. Angier, 1990, Torrey, 2019.

²⁵I concede that this contrast can easily be overstated. All my argument requires is that, sometimes (perhaps even often), recognizably epistemic grounds bear on this sort of decision-making. It does not, I think, require that these grounds always be fully and finally distinguishable from various prudential considerations that *also* bear on epistemic agency, or that we can always tell whether a responsible all-things-considered decision was made primarily on the basis of epistemic or more broadly prudential reasoning. For the purposes of distinguishing my position from standard articulations of pragmatic encroachment, it should be sufficient if prudential and epistemic considerations, in this sense, form a continuum. It is the dialectical burden of pragmatic encroachment, as I understand it, to argue that *strictly non-evidential considerations* (such as the “practical stakes” involved in a decision) can help determine what one is justified in believing, even as we hold the *strictly epistemic considerations* fixed. But at the end of the day, I don’t think anything much depends on this. As noted above, I would be perfectly happy if it were agreed that pragmatic encroachment captures one dimension of a kind of epistemic-cum-practical decision-making, for which my argument provides a larger and better motivated theoretical framework.

A standard line to take these days, often inspired by Williamson, 2000, is to hold that the aim of inquiry is knowledge (e.g., Whitcomb, 2010; Friedman, 2017; Kelp, 2021). Maybe, then, we could retool the argument of the previous pages into service of something like a Knowledge-Norm for Inquiry: we should allocate our epistemic resources in ways that would maximize the expected yield of knowledge.

But it is far from clear that this would work: consider two lines of inquiry in which we have (some) evidence to provisionally support p and q , respectively, but in each of which we currently fall short of knowledge. It is by no means obvious that we would be right to focus our resources on the line of inquiry which is most likely to produce knowledge, for instance, if the knowledge in question would be low-hanging fruit of comparatively little epistemic significance. Instead, we might do better to focus our resources on the other line of inquiry, even if it is overall less likely to produce the sorts of determinate results that knowledge would seem to require. This is precisely what the example of the Human Genome Project teaches. It certainly generated an enormous amount of knowledge. But how valuable, epistemically speaking, is this knowledge?

What about a norm which ties justification to the maximization of truth? Consider, for instance, Hilary Kornblith (1983): “An epistemically responsible agent desires to have true beliefs, and thus desires to have his beliefs produced by processes which lead to true beliefs; his actions are guided by these desires. Sometimes when we ask whether an agent’s belief is *justified* what we mean to ask is whether the belief is the product of an *epistemically responsible action*,” or in other words, “whether [the agent] has done all he should to bring it about that he have true beliefs (ibid, p. 34).”

It is worth noting that Kornblith’s proposal draws on very different philosophical motivations than the often more metaphysically tinged proposals for a Knowledge Norm of Inquiry. In particular, Kornblith ties the notion of epistemic justification directly to “responsible epistemic agency” in a way that would appear to render any categorical distinction between an “epistemology of inquiry” and an “epistemology of belief” problematic at best. In this sense, I acknowledge that Kornblith’s motivations for pursuing a “naturalistic” approach to epistemic normativity are in important ways similar to mine.

Nonetheless, I am skeptical of the idea that we might try to capture this epistemic norm in the form of an injunction to maximize the truth of one’s beliefs. We can notice, first, that Kornblith’s claim is couched in terms of “should” rather than “could.” If it were stated in terms of “could,” the claim would right away be shown false in light of the argument of this paper. Suppose that, in general, the way to increase one’s chances of having a true belief is to gather more evidence. But there is no limit to the amount of evidence one “could” gather. So, justified belief would be out of reach for cognitively limited agents.

Presumably, it is for such reasons that Kornblith prefers to talk in terms of “should” rather than “could.” Now, we have a different problem: since “should” already seems to invoke a normative standard, it is hard to see how it could serve to ground an account of justified belief. For how much evidence “should” one gather?—presumably, just as much as would be required to justify one’s beliefs. Again, this is a notion that may seem to make some sense when applied to lines of inquiry on a one-by-one basis, which is essentially how Kornblith views it. But it does not conveniently scale up to the level where we are considering the allocation of epistemic resources across several lines of inquiry.

Kornblith argues, correctly in my view, that “[t]he manner in which one goes about acquiring evidence bears on the justificatory status of one’s beliefs.” He elaborates the point as follows: “an agent who [...] simply refuses to look at evidence not already in his possession, is acting in an epistemically irresponsible manner.” Accordingly, “it seems unreasonable to say that he is justified in that belief” (Kornblith, 1983, p. 35). Now, if the question of justification never required us to look beyond the agent’s epistemic standing with respect to some particular proposition p , then it would make sense to posit a norm according to which agents should always strive to have *more* evidence regarding p (provided such evidence is “available”). Certainly, if someone were to “simply refuse” to consider such evidence, it could only serve to undermine their justification.

But as I have argued in this paper, this perspective is often misleading. Once we take into account the way in which finite epistemic agents are required to distribute their epistemic resources across a number of different inquiries, we will also come to recognize that someone may indeed be fully epistemically responsible even in “refusing” to gather or process more evidence regarding some particular p .

To argue for this view, we do not need to claim that truth is not an epistemic good, or that responsible epistemic agents are not in some broader sense “aiming at the truth.” Instead, it is sufficient to point out that any inquiry aimed at establishing a truth imposes significant opportunity costs and that epistemically speaking, not all truths are equal: in consideration of our cognitive finitude, a normative injunction to maximize truths known is not generally consistent with good epistemic practice, either in science or in everyday life. In many cases, one might responsibly choose to pursue a line of inquiry even with the understanding that it is less likely to eventuate in a true belief (or a justified true belief) than some other line of inquiry.

What, then? What are responsible epistemic agents trying to achieve in these kinds of cases? Drawing on the previous two proposals, one might consider something like the following: the aim is just to *maximize truths worth knowing*. There is clearly something right about this proposal. Even so, I think it would miss the fact that there is something more basic going on in these cases: what one should be doing, epistemically speaking, is just to try to gain the most consequential information that one can. In what sense “consequential”? Consequential in the sense of having the potential to settle, overturn, or redirect, not just *this* line of inquiry but several others related to it. *Informativeness*, I propose, is the value that epistemic agents should seek to maximize: they do well by the epistemic resource allocation problem if they distribute their resources in ways that can be expected to maximize the informativeness of their epistemic endeavors. In so doing, they can perfectly well retain their justification for believing lots of things, even in the knowledge that there remains available evidence “out there” which could have overturned these beliefs, and even if their epistemic decision-making is not geared toward maximizing the number of truths known.

“Informativeness” might seem vague and impressionistic at this point, and it probably is. It must await a separate paper to flesh it out in more detail, and perhaps to offer a formal framework for its measurement. But I do think that what I’ve said so far suffices at least to get a sense of the broad direction that a more fully developed account would take. The expected informativeness of a line of inquiry, as I understand it, is to be measured primarily along two dimensions: first, the probability that the line of inquiry will in fact turn up the evidence that we seek, and second, the degree of revision this evidence (if it were uncovered) would prompt in our “web of belief.”²⁶ Generally, we should prefer high-impact evidence to low-impact evidence.²⁷ But if the high-impact evidence would be exceptionally difficult to come by, it might introduce prohibitive opportunity costs: accordingly, we might be better off, from an epistemic point of view, to pursue easier-to-come-by evidence, even if, parcel by parcel, this evidence is less informative.

More broadly, why substitute a notion of “informativeness” for a notion of evidence? In some ways, the two notions seem importantly related. For instance, a standard approach within Bayesianism effectively defines evidence in terms of informativeness, via its impact on our credence assignments. That is, E is evidence for H if and only if the probability of H given E is greater than the

²⁶Or perhaps better, our “web of credence-assignments.”

²⁷One might worry here that this runs afoul of advice offered by Kuhn (1962) and Harman (1986), according to which one is entitled to display a degree of “conservatism” in one’s epistemic conduct. But in the end, I do not think there is any real tension here. On the one hand, there is the question of what evidence one should seek to procure; on the other hand, there is the question of what interpretation one should give to this evidence. On my account, these questions intersect: for at the same time, one’s estimation of how consequential some piece of evidence might be is clearly dependent on the question of how one would interpret it, were one to have it. I am open to the possibility that one should sometimes be maximally conservative in one’s interpretation of new evidence. But I don’t think our epistemic norms would ever condone a decision *not to pursue* some evidence just because one’s thinks (or fears) that gathering this evidence would be consequential in the way that I have described.

probability of H alone. An experiment whose outcome would not move our credence assignments in any way is plainly a bad experiment, simply in virtue of being “uninformative”: whatever the outcome, it would not tell us anything about the things that we set out to learn about.²⁸ Likewise, we can capture degrees of informativeness in terms of *how much* E (v. E*) would impact our credence in H.

Naturally, I’m very much inclined toward this way of thinking. But again, I must warn against the supposition that there is some unique H against which we would measure the informativeness of some E. Instead, the value of E would lie in the degree to which it requires us to adjust our credences with respect to a whole range of hypotheses, H1 ... Hn. Correspondingly, of course, it could be that uncovering some E would indeed have a very significant impact on some unique H: but if that H is probabilistically unconnected to most other things that we are currently interested in learning more about, then we could, for excellent epistemic reasons, forego further pursuit of that E. While inquiry in pursuit of E might certainly turn out to be informative, it is also likely to be significantly less informative than some other inquiry that we *could*—and therefore *should*—pursue instead.²⁹

Quite often, though, philosophical discussion conceptualizes evidence in a very different way, more like discrete parcels of uniform value, like barrels of oil. This is the conceptualization that naturally invites the supposition that, all things equal, it is always better to have more rather than less of it. But if evidence—say in the form of discrete “observations”—is of little value in its own right apart from the information that it provides, then it should be clear that we cannot measure the expected informativeness of a line of inquiry simply in terms of the “amount of evidence” (e.g., number of observations) that it promises to yield.³⁰ It is obvious that different bodies of evidence can be differently informative. In some sense, it is certainly true that evidence is something we can collect and store for possible future use, even though we currently have no clue what its epistemic value would be. As an epistemic strategy, though, this quickly founders on the recognition that there is really an infinity of such discrete bits of evidence that we could seek out. It would be seriously misguided to attempt to compile all that we can *now*, just because we could possibly find some epistemic use for it later. Therefore, any epistemic agent is required to make decisions about prioritization. As we have seen, standard philosophical discussions of epistemic normativity appear to have very little to say about what norms should guide these decisions. This paper has aimed to bring this problem to the foreground and suggest a plausible candidate norm: in recognition of the opportunity costs involved in any one inquiry, epistemic agents should allocate their limited cognitive resources in a manner that maximizes the expected informativeness of their inquiries, seen as a whole.

Admittedly, such a norm might still seem to leave us short of providing confident answers to a number of further questions that traditional evidentialist accounts have trained us to think we *need* answers to. These accounts at least hold out the promise of providing clear and determinate answers to questions of the form, “given a body of evidence E, is A justified in believing that *p*?” (for any arbitrary E, A, and *p*). By contrast, my view seems to embroil us in qualitative assessments that might not point to similarly clear verdicts.

But in the end, I don’t think this is a weakness of my view. As I have argued (section 4), traditional evidentialism’s promise to provide such answers is false: there will be many cases in which epistemic agents will not be in any position to determine whether their evidence supports *p* (or to what degree it supports *p*) without conducting further evidence-producing inquiries, specifically inquiries aimed at establishing the significance of the evidence currently in their possession.

²⁸For an argument along these lines, broadly congenial to the argument of this paper, see Zollman, 2021.

²⁹For different developments of similar ideas, see Heesen, 2015; Dallman, 2017.

³⁰Again, consider how critics of the Human Genome Project were not alleging that the evidence it would gather was not, under any circumstances, worth having. Rather, they were claiming that there were better things one might do with the estimated 3 billion dollars that it budgeted.

So far, this argument might seem to push us in the direction of the available evidence thesis. However, as I argued in Sections 2 and 3, there is generally no good answer to the question of how much evidence is “available” to a subject. As such, the available evidence thesis will likewise embroil us in similar qualitative assessment soon enough, by requiring us to ask, “has A done enough to supply themselves with evidence relating to p ?” The question, so far as it goes, is fine: the problem, rather, is that the available evidence thesis will tend to give us wrong or misleading answers to this question, insofar as it brings our exclusive focus to bear on evidence relating to some unique p .

By contrast, I take myself to have shown that asking whether someone is justified in believing that p is tantamount to asking whether they have responsibly allocated their epistemic resources in seeking to determine whether p . A positive answer to this question is in no way contingent on them having exhausted the available evidence bearing on p .

In this sense, my account clearly presumes that responsible epistemic resource allocation is *necessary* for justification: an agent cannot be justified in believing that p unless they have devoted sufficient epistemic resources toward establishing whether p . But at the same time, what counts as sufficient is clearly context-sensitive, in the sense that we can imagine two individuals (or two research groups) allocating the same amount of epistemic resources to p , but finding that one is justified, but the other is not. The difference, broadly speaking, would lie in what further opportunity costs must factor into their respective epistemic resource allocation decisions.

As such, I do not think we should expect a clear answer to the question of “how low one can go” in terms of epistemic resource allocation for some p and still retain justification for belief (perhaps short of establishing greater than chance probability). This is perhaps most easily seen in the case of subsidiary lines of inquiry. Suppose we determine that an agent has “done well” by the epistemic resource allocation problem but has not devoted an enormous amount of resources to some relatively peripheral question q . Nonetheless, in virtue of “having done well” by the epistemic resource allocation problem, we assume that they are justified in believing that q . What confers the justification? Is it some kind of deep epistemic holism according to which all relevantly connected propositions stand to “inherit” their justification from the “web of beliefs” as a whole?

In response, I do not think our commitment to epistemic holism needs to be quite so unconstrained. Following more traditional epistemological accounts, I continue to assume that questions of justification are naturally centered on a particular focal hypothesis p . The central strand of my criticism of these more traditional accounts comes to this: we cannot measure the justification *simply* in terms of how they stand with respect to the actual or available evidence regarding p , considered in isolation. Sometimes determining whether p requires epistemic investments also in subsidiary lines of inquiry into supporting propositions q , r , s , and so forth. However, given our cognitive limitations, there will be cases in which we remain justified in believing these propositions even though we have not exhausted the available evidence relating to them.

Imagine, then, an epistemic agent whose primary concern is p : they have devoted significant amounts of epistemic resources to inquiring into p directly, but will also have allocated a nontrivial amount of epistemic resources to q , as required to support their inquiry into p . What should we say about their justification for believing q , in light of the comparative shortness of the resources that they have allocated to this question as such? I think the correct answer is that they are justified in believing that q , *insofar as p is concerned*. But of course, q itself might at some point come under scrutiny in its own right: I do not think there is any inconsistency in saying that A is justified in believing q , insofar as p is concerned, while also allowing that they might not be justified in believing q , were q to become a focal point of inquiry. In this sense, “how much” epistemic resources one should allocate to a particular question is dependent on “how central” that question is in the context of one’s overall epistemic endeavors. This is precisely the point of the argument above: “how central” a question is in the context of one’s overall epistemic endeavors is just the question of how “informative” one should assume inquiry into that question might be.

So, to be sure, there remains a limited but relevant commitment to holism in my account: we cannot hope to assess an agent’s epistemic standing with respect to propositions on a one-by-one

basis. But when we determine that someone is justified in believing p , it is not the case that we also determine that they are justified *simpliciter* in believing every proposition on which p depends.

6. Concluding Remarks: Non-Ideal Epistemology as Normative Epistemology

This article has explored some motivations for shifting away from a conception of epistemic justification as a status that supervenes on the three-way relation between a subject, some body of evidence, and a particular individual proposition, so as to move in the direction of a conception of epistemic normativity centered on the subject's management of the inherent risks of epistemic resource allocation across several lines of inquiry. In the final section, I suggested, as a first pass, that responsible risk management here should be guided by a norm of allocating resources according to expected informativeness: roughly, that one should distribute one's epistemic resources in the manner which one should expect to yield the greatest amount of information overall.

What set us on this path was an acknowledgement of the epistemic significance of cognitive limitations, and in particular, the fact that human beings can never hope to form their beliefs in light of all the evidence. Obviously, this is not to say that evidence (and the pursuit of evidence) does not play a key role in our epistemic agency. But it does raise the question, what is the proper subset of the total evidence that one is required to consider in order to have a justified belief? I have argued that the apparently most promising answer—the subset of the evidence that is “available to us”—is not so promising after all. The problem is not simply that it is of limited help, since we do not always know what counts as “available.” More importantly, it also seems to be false: there will be lots of cases in which someone retains their justification even though they have knowingly neglected available evidence. It is in order to get a handle on this puzzle that I suggested we switch to considering not just how agents stand with respect to their evidence for particular lines of inquiry, but how they distribute their limited resources in the search for evidence across several lines of inquiry.

One might think that the problem here was that of making any concessions to “non-ideal epistemology” to begin with, by way of moving to recognize the epistemic significance of our cognitive limitations.³¹ It is only if we set our foot on that greased path that we will end up stymied by these puzzles. By contrast, if we staunchly maintain that the epistemic norm is what it is regardless of human imperfections, then we can steer well clear of these rocky waters.

But it is doubtful that this is a viable strategy in the end since the question of concern remains very much on the table. Even if there is some sense of “norm” on which the epistemic norm should make no concessions to human finitude, there is clearly another sense of “norm” which aims to fulfill precisely the role that this paper describes, namely that of providing guidance for cognitively limited agents engaged in a variety of epistemic pursuits, whether in science or in everyday life. How should such agents comport themselves, epistemically speaking, in recognition of their own finitude? This is a question that clearly calls out for a norm—and an epistemic norm, at that—even if there are other things one might also call by that name. Moreover, giving content to that norm should be part of the epistemologist's job description, even if it requires them to drop some of the idealizations that have long shaped theoretical inquiry in their field.³² In this article, I have

³¹See Begby (2021), McKenna (2023), and St. Croix (2024) for different developments of the idea of non-ideal epistemology. For counterpoints, see Carr (2021) and Greco (2023). Note, though (as I argue in detail in Begby, 2021: ch.3), that the claim animating non-ideal epistemology is not that idealizations have no role to play in normative epistemology, but rather that the idealizations in question must be relativized to a realistic picture of human cognitive capacities.

³²A more radical approach might argue that a “fully idealized epistemology” was never a conceptually stable notion to begin with. Here's one way to think about it: even in “ideal epistemology,” the epistemic norm is presumably supposed to proscribe the transition from evidence to belief. But consulting evidence is an inherently risky (i.e., “uncertain”) method of belief formation; only finite minds should ever consider deploying it. Therefore, any evidence-based epistemology is already a form of non-ideal epistemology. The only question is which idealizations we should hold onto and which we should abandon.

argued that we will be forced to drop one particular idealization, one that comes to us so naturally that we are rarely even brought to consider it as such, namely the idealization that cognitive agents are only ever engaged in a single line of inquiry at a given time. As I have argued, it is not just that this is not true; more importantly, it really could not be true. Nonetheless, this is the idealization that pushes us to think that epistemic justification must be a status that supervenes on the relation between a cognitive subject, some evidence base, and a particular, individual hypothesis.³³

References

- Ackert, L. T. (2006). The role of microbes in agriculture: Sergei Vinogradskii's discovery and investigation of chemosynthesis, 1880-1910. *Journal of the History of Biology*, 39, 373–406.
- Angier, N. (1990). "Great 15-Year Project to Decipher Genes Stirs Controversy." *The New York Times* June 5, 1990. <https://www.nytimes.com/1990/06/05/science/great-15-year-project-to-decipher-genes-stirs-opposition.html>.
- Begby, E. (2021). *Prejudice: A Study in Non-Ideal Epistemology*. Oxford University Press.
- Carr, J. R. (2021). Why ideal epistemology? *Mind*, 131(524), 1131–1162.
- Chandrasekhar, S. (1976). Verifying the theory of relativity. *The Royal Society Journal of the History of Science*, 30(2), 249–260.
- Dallman, J. (2017). When obstinacy is a better (cognitive) policy. *Philosophers' Imprint*, 17(24), 1–17.
- Dougherty, T. (2014). The 'ethics of belief' is ethics (period): reassigning responsibility. In J. Matheson & R. Vitz (Eds.), *The ethics of belief*. Oxford University Press.
- Dubilier, N., Bergin, C., & Lott, C. (2008). Symbiotic diversity in marine animals: harnessing the power of chemosynthesis. *Nature Reviews Microbiology*, 6(10), 725–740.
- Eddington, A. S. (1919). The total eclipse of 1919 May 29 and the influence of gravitation on light. *The Observatory*, 42, 119–122.
- Fantl, J., & McGrath, M. (2002). Evidence, pragmatics, and justification. *Philosophical Review*, 111(1), 67–94.
- Feldman, R., & Conee, E. (1985). Evidentialism. *Philosophical Studies*, 48(1), 15–34.
- Flores, C., & Woodard, E. (2023). Epistemic norms on evidence-gathering. *Philosophical Studies*, 180, 2547–2571.
- Friedman, J. (2017). Why suspend judging? *Nous*, 51(2), 302–326.
- Friedman, J. (2019). Inquiry and belief. *Nous*, 53(2), 296–315.
- Goldberg, S. (2017). The asymmetry thesis and the doctrine of normative defeat. *American Philosophical Quarterly*, 54(4), 339–351.
- Goldman, A. (1979). What is justified belief? In G. Pappas (Ed.), *Justification and knowledge*. Boston, MA: D. Reidel.
- Greco, D. (2023). *Idealization in Epistemology*. Oxford University Press.
- Hacking, I. (1983). *Representing & Intervening: Introductory Topics in the Philosophy of Science*. Cambridge University Press.
- Harman, G. (1973). *Thought*. Princeton, NJ: Princeton University Press.
- Harman, G. (1986). *Change in View: Principles of Scientific Reasoning*. Cambridge, MA: The MIT Press.
- Heesen, R. (2015). How much evidence should one collect? *Philosophical Studies*, 172(9), 2299–2313.
- Hume, D. (1748). *An Enquiry Concerning Human Understanding*. T. L. Beauchamp (Ed.). Oxford University Press, 1999.
- Keller, J. (2013). "Uncovering the Past Using the Future: How Lasers are Revolutionizing Archeology." *The Verge*, June 21, 2013. <https://www.theverge.com/2013/6/20/4445568/lasers-lidar-archaeology-detailed-topographical-maps>.
- Kelly, T. (2006/2014). "Evidence." *Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/entries/evidence/>.
- Kelp, C. (2021). Theory of inquiry. *Philosophy and Phenomenological Research*, 103(2), 359–384.
- Kennefick, D. (2011). Not only because of theory: Dyson, Eddington, and the competing myths of the 1919 eclipse expedition. In C. Lehner, J. Renn, & M. Schemmel (Eds.), *Einstein and the changing worldviews of physics*. New York: Springer Publishing.
- Kornblith, H. (1983). Justified belief and epistemically responsible agency. *Philosophical Review*, 92(1), 33–48.
- Kuhn, T. S. (1962). *The Structure of Scientific Revolutions*, 3rd ed. Chicago, IL: The University of Chicago Press.
- Lackey, J. (2008). *Learning from Words: Testimony as a Source of Knowledge*. Oxford University Press.
- Lycan, W. G. (1977). Evidence one does not possess. *Australasian Journal of Philosophy*, 55(2), 114–126.
- McKenna, R. (2023). *Non-Ideal Epistemology*. Oxford University Press.
- McWilliams, E. C. (2021). Evidentialism and belief polarization. *Synthese*, 198, 7165–7196.
- Nolfi, K. (2018). Why only evidential considerations can justify belief. In C. McHugh, J. Way, & D. Whiting (Eds.), *Normativity: Epistemic and practical*, Oxford University Press, 2018.
- O'Connor, C. (2021). "Science Isn't Here for Your Mommy-Shaming." *Nautilus*. <https://nautil.us/issue/100/outside/science-isnt-here-for-your-mommy-shaming>.

³³In developing this article, I have benefitted greatly from discussions with Holly K. Andersen, Yan Chen, Matt Maxwell, Bjørn Ramberg, Jacob Schwartz, and Jingyi Wang, as well as two anonymous reviewers of this journal and audiences at King's College London (May 2018), University of Oslo (May 2018) and the June 2022 meeting of the European Epistemology Network conference at the University of Glasgow.

- Pettigrew, R. (2016). *Accuracy and the Laws of Credence*. Oxford University Press.
- Pettigrew, R. (2011/2019). "Epistemic Utility Arguments for Probabilism." *Stanford Encyclopedia of Philosophy*. <https://plato.stanford.edu/entries/epistemic-utility/>.
- Rescher, N. (1989). *Cognitive Economy: The Economic Dimension of the Theory of Knowledge*. Pittsburgh, PA: The University of Pittsburgh Press.
- Sher, G. (2009). *Who Knew? Responsibility without Awareness*. Oxford University Press.
- Smith, H. (1983). Culpable ignorance. *Philosophical Review*, 92(4), 543–571.
- St. Croix, C.. (2024) Evidence in a non-ideal world: how social distortion creates skeptical potholes. In H. Hänel & J. Müller (Eds.), *The Routledge handbook of non-ideal philosophy*. London: Routledge.
- Torrey, E. F. (2019). "Was the Human Genome Project a Dud?" *The Dallas Morning News*, Oct. 13, 2019. <https://www.dallasnews.com/opinion/commentary/2019/10/13/was-the-human-genome-project-a-dud/>.
- Wamsley, L. (2020). "What Accounts for High Coronavirus Positivity Rates among Florida Kids?" *NPR*, July 17, 2020. <https://www.npr.org/sections/coronavirus-live-updates/2020/07/17/892427161/what-accounts-for-high-coronavirus-positivity-rates-among-florida-kids>.
- Whitcomb, D. (2010). Curiosity was framed. *Philosophy and Phenomenological Research*, 81(3), 664–687.
- Williamson, T. (2000). *Knowledge and its Limits*. Oxford University Press.
- Zimmerman, M. J. (2008). *Living with Uncertainty: The Moral Significance of Ignorance*. Cambridge University Press.
- Zollman, K. (2021). The theory of games as a tool for the social epistemologist. *Philosophical Studies*, 178, 1381–1401.