Synthese _####################_
https://doi.org/10.1007/s11229-025-04973-0

ORIGINAL RESEARCH

Interventions in inquiry

Jennifer Rose Carr[1]

Received: 9 August 2024 / Accepted: 21 February 2025

Abstract

A central question on the norms of inquiry—"zetetic norms"—concerns how they relate to epistemic norms. Both centrally involve learning about the world. Epistemic norms are thought to govern beliefs and credences, while zetetic norms extend to actions. A cluster of common views of epistemic norms are consequentialist: they hold that the fundamental source of epistemic value is true beliefs/accurate credences (*veritism*), or knowledge (*gnosticism*), and that rational beliefs are those that best promote epistemic value. It's natural, then to think of zetetic norms as extending epistemic consequentialism to non-doxastic acts. This paper argues against these forms of zetetic consequentialism, with a focus on zetetic veritism. The general argument is this: we can come to have true beliefs in two distinct ways: by *discovering* truths and by *creating* truths. Only the former is zetetically valuable. Zetetic veritism falsely entails that, in many ordinary cases, good inquiry requires creating truths: intervening into the world to make it predictable, and indeed, to make it match our predictions. These arguments generalize to gnosticism and other popular forms of zetetic consequentialism. I discuss an alternative view, *zetetic observationalism*, according to which the aim of inquiry is learning by pure observation: the facts that matter for inquiry are independent of our interventions. This theory avoids counterexamples to zetetic veritism, but has surprising consequences: that learning answers to questions under inquiry often has no immediate value for inquiry, and that epistemology and inquiry don't value accuracy in the same way.

Keywords Inquiry · Accuracy · Epistemic norms · Zetetic norms

This paper concerns inquiry and its norms. Inquiry, and norms of inquiry, are often thought of as the project of science and scientists. But we engage in inquiry through-

✉ Jennifer Rose Carr
j2carr@ucsd.edu

[1] University of California San Diego, La Jolla, CA, USA

out ordinary life, wondering about banal questions, taking banal steps toward uncovering their answers. I wonder whether I have a can of chipotles; I poke around the pantry to learn the answer. I wonder what the cat will do if I poke my finger in his mouth while he's yawning; I perform a simple but dangerous experiment. Conversations are acts of inquiry: they involve sharing information, narrowing down the space of possibilities in our epistemic common ground.[1] Conversations are governed by questions under discussion[2] and aim at collectively reasoning toward their answers. In general, we investigate our world, gather evidence, narrow down the space of epistemic possibilities, aiming to narrow our probability distributions around the actual world.[3]

A central question on the norms of inquiry—"zetetic norms"[4]—concerns how they relate to epistemic norms. Epistemic norms have been thought of as governing beliefs and credences, while the norms of inquiry extend to actions. But both seem to concern learning about the world.

A common view of epistemic norms is what I'll call "epistemic veritism": the view that the epistemic value of beliefs (credences, etc.) is a function of their accuracy—the more accurate, the better—and that rational beliefs are those that best conduce to promoting that value. This view is often simply called "epistemic consequentialism", because the hypothesis that accuracy is the unique source of epistemic value is so prevalent.[5] And since the goal of inquiry is to learn the answers to questions, it's natural to think of zetetic norms as guided in the same way by the pursuit of accuracy. That is, it's natural to think of zetetic norms as an extension of epistemic norms to non-doxastic acts: for example, to think of zetetic decision theory as an extension of epistemic decision theory, with the same axiology and rules, but applying beyond belief states to outward actions. I call this view **zetetic veritism.**

The veritist's accuracy-based axiology isn't the only candidate theory of fundamental epistemic value. Littlejohn (2015), for example, argues that knowledge, rather than truth, is the source of epistemic value; Littlejohn (2018) calls this view "gnosticism". In the zetetic literature, this view has many proponents, e.g. Whitcomb (2010), Kelp (2014), Kelp (2021), Sapir and van Elswyk (2021). Some epistemic and zetetic consequentialisms may prefer the gnostic, rather than veritist axiology;[6] still others may prefer some other alternative, perhaps sensitive to evidence, or conceptual awareness, or...

[1] Stalnaker (1970, 1978).

[2] See Roberts (1996) for an influential defense.

[3] For a detailed discussion of the mental states and activities involved in ordinary inquiry, see Friedman (forthcoming).

[4] This term was introduced into the literature in Friedman (2020).

[5] See e.g. Goldman (1979) and much of the reliabilist tradition, Ahlstrom-Vij and Dunn (2014); Singer (2018). Accuracy-first epistemology is a popular recent version of this view. Seminal works in this area include Joyce (1998, 2009), Greaves and Wallace (2006); Leitgeb and Pettigrew (2010a, b); Pettigrew (2016).

[6] Though Littlejohn (2015, 2018) argues that this axiology pairs better with nonconsequentialism.

This paper argues against these forms of zetetic consequentialism. I focus on zetetic veritism as an example case;[7] my arguments generalize straightforwardly to zetetic gnosticism, and less straightforwardly to some other forms of zetetic consequentialism. The general argument is this: we can come to have true beliefs in two distinct ways: by *discovering* truths and by *creating* truths. But only the former is zetetically valuable. Popular forms of zetetic consequentialism entail that in some cases, you're zetetically required to intervene into the world to make it predictable, and indeed, to make it fit your predictions. But the norms of inquiry, whatever they are, do not require you to alter the world to make it fit your predictions. So these forms of zetetic consequentialism are false.

The plan: in §1, I'll provide a bit of dialectical prologue to show what inspired the central questions of this paper and why they matter for inquiry and rational belief. In §2, I'll characterize zetetic norms and zetetic veritism. §3 presents our guiding counterexample to zetetic veritism, and §4 diagnoses the problem for zetetic veritism, visible in a range of other counterexamples. In §5, I discuss an alternative view, **zetetic observationalism,** according to which the aim of inquiry is learning by *pure observation*: the facts that matter for inquiry are those independent of our interventions. This theory both avoids and explains the counterexamples to zetetic veritism, but has surprising and possibly costly consequences: in particular, that in many cases, accuracy regarding answers to a question under inquiry has no value for inquiry, and epistemology and inquiry don't value learning in the same way.

1 A dialectical prologue

Before launching into the main arguments of this paper, I'm going to saddle you with a quick recounting of my own idiosyncratic backstory in approaching these questions.

I'm not an epistemic veritist in the sense above, for the typical reasons: *epistemic tradeoffs* and *self-verifying beliefs*. In cases of epistemic tradeoffs, one can take on a belief that one is in a position to know is false, in order to predictably gain enough true beliefs to counterbalance the falsehood.[8] Epistemic veritism entails—implausibly, on my view—that epistemic rationality sometimes requires making these tradeoffs:

Tradeoff case

You're a researcher in a close relationship with your coworker, and you've just received decisive evidence that they've been extremely cruel behind your back. If you believe this evidence, then you'll find it impossible to concentrate on your research with them as a coworker, and will want to leave your job. If you reject your evidence, then your research career will thrive. Believing the evi-

[7] Why focus on veritism? (i) The accuracy-first framework offers an off-the-shelf quantitative decision theory for gradational accuracy; (ii) doing so provides continuity with my previous work, and with Bayesian epistemology.

[8] Firth (1978/1998); Berker (2013a); for examples in the credal context, see Greaves (2013).

dence results in beliefs with lower expected accuracy than does rejecting your evidence. So epistemic veritism requires rejecting the evidence.

In cases of self-verifying beliefs, there are some belief states that, if adopted, *cause* their contents to be accurate; epistemic veritism holds—implausibly, on my view—that epistemic rationality requires adopting such belief states wherever possible.[9]

Self-verifying belief case

The greater your credence in q, the greater the objective chance that q, and you know this. So maximal and minimal credence in q both accuracy-dominate every other credence in q. So epistemic veritism prohibits any degree of uncertainty about q.

In response to objections from epistemic tradeoffs and self-verifying beliefs, some veritists bite the bullet, sometimes offering accompanying error theories of our intuitions.[10] Indeed, in the case of self-verifying beliefs, some reject that there's a bullet to bite. Many insist that it's *intuitively* irrational to turn down a "free true belief", and so hold that these examples support veritism.

The following response briefly appealed to me: there is a sense in which you oughtn't turn down free true beliefs. But the relevant 'ought' shouldn't be thought of as *epistemic*, but rather pragmatic. Suppose you happen to desire having the most accurate beliefs possible. As a matter of rational desire-satisfaction, you *pragmatically* ought to accept epistemic tradeoffs, self-verifying beliefs, and incoherent beliefs, whenever doing so maximizes expected accuracy. In doing so, though, you might end up with beliefs that are *epistemically* irrational. Epistemic norms are objective; your subjective desires—even the desire for accurate beliefs—are epistemically irrelevant.

Where the pursuit of accuracy is pragmatically rational, we can assess not only beliefs but also actions for accuracy-conduciveness. And so, you might think, we've moved from epistemic to *zetetic* norms. These norms seem to fit within the mold of pragmatic, instrumental rationality: they apply to voluntary, outward actions. And so it's *prima facie* appealing to view them as pragmatic norms for the effective pursuit of accuracy.

This picture seemed to me to have explanatory benefits. It seemed to offer a tidy explanation of the split intuitions on epistemic tradeoffs and self-verifying beliefs: one set of intuitions tracks epistemic norms; the other tracks related but distinct zetetic norms. And this picture offers a particularly nice explanation of an apparently

[9] See e.g. Berker (2013b); for examples in the credal context, see Carr (2017). Such cases relate closely with self-falsifying beliefs, where the beliefs that are most accuracy-conducive are often incoherent. If I know that q is true if and only if I don't believe q, then the most accurate belief state will involve doxastic neutrality on q and full-throated rejection of *not-q*. But this state is incoherent: if I fully reject *not-q*, then I ought to fully believe q. See Caie (2013); Greaves (2013).

But since the latter generate a clash between plausible epistemic norms—*accuracy* vs. *coherence*—they are sometimes taken to generate epistemic dilemmas. Self-verifying beliefs needn't generate any similar clash. Instead, they suggest a more fundamental problem for epistemic veritism.

[10] E.g. Pettigrew (2018); Singer (2018).

virtuous epistemic trade-off that many find most compelling: cases where scientists incorporate false idealizations into theories or models. These idealizations promote various goals in scientific inquiry—improving tractability, characterizing specific phenomena in isolation from complicating extraneous factors, etc.—and therefore promote more overall accurate belief states.

I've led you down a primrose path toward a prima facie easy solution. But there's a problem. Closely related plausible counterexamples arise for zetetic veritism; I present several in §3 and §4. So, on my view, we should reject veritism in both the epistemic and zetetic realms. Moreover, though the two forms of veritism generate similar problems, they seem to require different resolutions. The alternative to veritism that I discuss in §5 doesn't port over to the epistemic realm, and is orthogonal to the alternative to epistemic veritism that I've defended in other work—which doesn't port over to non-doxastic action.[11]

Before we see these arguments, let me introduce the main puzzle in a less idiosyncratic context.

2 Zetetic norms

2.1 What are zetetic norms?

What makes a norm a *zetetic* norm?[12] It's conventional wisdom that epistemic norms are distinct from moral and other pragmatic norms.[13] But so far, no conventional wisdom in the younger literature on zetetic norms has emerged. On my view, zetetic norms are unified, and distinguished from other pragmatic norms, by the kind of value or rightness they concern (e.g., the value of accuracy). Others, e.g. Thorstad (2024, forthcoming), hold that zetetic norms are unified by their object of evaluation ("acts of inquiry"), and that the value/rightness relevant to inquiry is simply moral value/rightness.

Why accept that there are distinctively zetetic norms?

First, there are cases where distinctively inquiry-based values/reasons seem to clash with moral and prudential values/reasons. Animal research ethics offers plenty of examples.

Second, if zetetic norms are unified by their object of evaluation—"acts of inquiry"—then these ought to form a distinctive class of acts. But it's hard to see how they can. Gathering evidence is a canonical act of inquiry. But everything we do involves gathering evidence. When I ate tacos today, I gathered evidence that the salsa had not gone bad. If eating tacos is thereby an act of inquiry, then anything is. Perhaps scientific inquiry comprises only certain forms of action, but (i) its

[11] In (Carr, 2017, 2019), I argue that the correct epistemic decision theory evaluates, not acts of possessing a credence function—events that interact causally with the world and can affect their own accuracy—but "free-floating" credence functions, construed as mathematical objects, which are causally inert and evaluable even at worlds where no agent adopts them.

[12] Under "norms", I include various normative categories, including "oughts", values, rational statuses, and so on.

[13] But see Cohen (2016), Rinard (2017).

boundaries are famously blurry, and (ii) this paper concerns inquiry understood more generally.

Third: at first pass, we might think of zetetic norms as a kind of generalization of distinctively epistemic norms. Zetetic norms have the same aims as epistemic norms—both aim at the truth—but zetetic norms apply beyond cognitive states to outward, voluntary actions. The relation between the two is more complicated,[14] but this loose picture seems like a good starting point.

2.2 Zetetic veritism

Given the loose picture above, let's explore the theory that the epistemic and the zetetic share an axiology. Accuracy-first epistemology—a popular view in contemporary epistemology—is **veritist**: it holds that epistemically rational belief states[15] are those that are most *accuracy*-conducive: belief states that avoid accuracy domination or maximize expected accuracy.[16] Note that on this usage, veritism comprises both an axiology (accuracy as value) and a deontic theory (expected value maximization).

Accuracy is generally interpreted in terms of a divergence between one's total belief state—typically a credence function—and the omniscient belief state at a world.[17] The most uncontroversial example of such a divergence is the **Brier score** of a credence function c: where W is the set of possible worlds, $\mathcal{F} \subseteq \wp(W)$ is an agenda of propositions, and v_w is the characteristic probability function of w, the local Brier score of c's credence in $q \in \mathcal{F}$ at w is the squared Euclidean distance between $c(q)$ and $v_w(q)$:

$$\mathfrak{b}(c, q, w) = (c(q) - v_w(q))^2$$

The Brier score is a measure of *inaccuracy*; the negative Brier score is a measure of *accuracy*. The negative additive Brier score of a credence function as a whole is a relatively uncontroversial veritist measure of the epistemic value of that credence function:

$$\mathfrak{e}(c, w) = -\sum_{q \in \mathcal{F}} (c(q) - v_w(q))^2 \qquad \text{(EPISTEMIC VALUE)}$$

In accuracy-first epistemology, epistemic value is identified with accuracy. The extension of accuracy-first epistemology to the zetetic treats the accuracy of one's future credences as determining zetetic value. An act φ is zetetically better than an act ψ just in case the expected accuracy resulting from φ is greater than that of ψ. Call this position **zetetic veritism**.

[14] See Friedman (2020).

[15] Under "belief states" I include credence functions, total sets of beliefs, credal sets, comparative confidence preorders, etc.

[16] Some canonical contributions: Joyce (1998, 2009); Greaves and Wallace (2006); Leitgeb and Pettigrew (2010a, b).

[17] For a generalization to binary beliefs, see Easwaran (2016).

(Note: many in the recent literature on inquiry accept that *knowledge*, not merely accuracy, is the aim of inquiry, or the locus of zetetic value, e.g. Kelp (2014, 2018, 2021); Sapir and van Elswyk (2021). Falbo (2023) argues that inquiry aims at epistemic improvement more generally. The arguments I present against zetetic veritism generalize straightforwardly to these alternative forms of zetetic consequentialism. I present my arguments in terms of accuracy for the sake of continuity with previous work (Carr, 2017, 2019), which this paper extends. I don't frame the paper as an argument against zetetic consequentialism generally for reasons discussed in §5. For arguments that the aim of inquiry is accuracy rather than knowledge, see Beddor (2023).)

We might define the relevant accuracies as that of one's credence function at a relevant future time (e.g., after an experiment, or a research program, is completed). Where "c" is a definite description of one's credence function at the relevant future time, c_w is the extension of "c" at a given w, and φ is the act one performs in w:

$$\mathfrak{z}(\varphi, w) = \mathfrak{z}(\mathsf{c}_w, w) = -\sum_{q \in \mathcal{F}} (\mathsf{c}_w(q) - v_w(q))^2 \qquad \text{(ZETETIC VALUE)}$$

Where $c(w \mid\mid \varphi)$ the probability weighting for w on the supposition that one performs φ on the appropriate decision theory, c's **expected zetetic value** for φ is the expected accuracy of one's future credences:

$$\mathbb{E}_c\mathfrak{z}(\varphi) = \sum_{w \in W} c(w \mid\mid \varphi)\mathfrak{z}(\mathsf{c}_w, w) \qquad \text{(EXPECTED ZETETIC VALUE)}$$

Alternatively, we might aggregate the accuracy of one's credences over a sequence of future times. Or we might interpret zetetic value as inhering not in individual accuracy, but *group* accuracy, aggregating the credence functions of groups or populations; see appendix A.1 for technicalities.

One point of contention is whether the zetetic value of an act should be relativized to some question whose answer the agent aims to learn: call this the **question under inquiry (QUI)**. In the Hamblin (1973) tradition, I'll represent a question as a set of propositions representing its possible answers: a partition over possibilities, $Q = \{q_1, \ldots, q_n\}$.[18] Each q_i is a possible answer to Q. The question-relativized variant of zetetic veritism specifies the agenda $\mathcal{F}$ in (ZETETIC VALUE) as the relevant Q: the only credences that are treated as relevant are credences in possible answers to Q.[19]

[18] Partitionality is a simplifying idealization. For natural language questions, there's controversy over whether answers must form a partition (Groenendijk and Stokhof, 1984; Ciardelli et al., 2018). And this matters for understanding inquiry: in cases where questions are nonpartitional, Salow and Ahmed (2019) argue that Good's (1962) theorem doesn't hold.

[19] I discuss question-relativized zetetic value at more length in §5. Unlike much of the recent literature on inquiry, I don't assume that question-relativity is a constitutive or defining feature of zetetic norms. I leave it an open question. First, if φing sheds light on propositions that aren't on the intended agenda, that may be valuable to the project of inquiry. The value at issue in scientific inquiry—which I take to be canonical inquiry and which is at stake in some of the motivating examples—concerns not merely known

For present purposes, it isn't necessary to choose among these formulations. In each case, according to zetetic veritism, an act φ is zetetically permissible just in case for all available acts ψ, $\mathbb{E}_{c\mathfrak{z}}(\varphi) \geq \mathbb{E}_{c\mathfrak{z}}(\psi)$.[20] For ease of exposition, I focus on *individual, question-neutral* zetetic value derived from credences *at some future time*.

3 A counterexample to zetetic veritism

3.1 The counterexample

Consider the following problem case for zetetic veritism:

Microecosystem

You're a biologist and have encountered an isolated microecosystem with its own peculiar plant and fungal species, which will change over time in unpredictable ways: various populations might overtake others, face resource scarcity or new infectious diseases, etc. So if you observe the microecosystem over time with minimal intervention (call this option "OBSERVE"), you'll face a great deal of uncertainty about its future, various of its hidden features, etc. But you have another option: you can set the microecosystem on fire, reducing it to ash, in such a way that it'll *predictably* remain nothing but ash forever after. (Call this option "INTERVENE".) If you do, your present and future beliefs about the microecosystem will have greater expected accuracy than if you merely observed.

Thesis: you are not required by the norms of inquiry to set the microecosystem on fire. More generally, setting the microecosystem on fire is not a zetetically better action than merely observing it.

So, the norms of inquiry permit actions that predictably generate less accurate beliefs than other available actions. Zetetic veritism is false.[21]

unknowns, but unknown unknowns and unconceived unknowns. Any non-trivial question-relativization would be artificial.

Second, in ordinary inquiry, the motivations for question-relativizing—for example, the agent's cognitive limitations or time-limitations—can typically be accommodated by using coarse-grained "small worlds" in the representation of the relevant decision problem and credence function. This maneuver, common in both pragmatic and epistemic decision theory, lets us side-step the issue of whether the relevant question-relativity is normatively fundamental, or merely a simplifying idealization, or a (bi-)contextualist feature of the language of zetetic evaluation, or the result of confounding non-zetetic goals, or... Focusing on question-neutral zetetic veritism also allows us to avoid the risk of tendentious intuition-mongering about what "the *real* question under inquiry" is in fictional examples.

[20] For those who are uncomfortable with talk of (non-all-things-considered) zetetic permissibility/requirement, the expected accuracy maximization condition can be restated in terms of reasons (one has most zetetic reason to φ) or goodness (it's zetetically best to φ).

[21] Note: a modified version of this example is equally a problem case for *epistemic* veritism. We simply ensure that the relevant options are epistemic acts: if you *believe* that you'll set the microecosystem on fire, then that belief will cause you to form the intention to do so, which in turn will cause you to do so. So this

3.2 Objections and modifications

Objection #1: Modal accuracy. INTERVENING may make your beliefs more accurate about nonmodal facts about the microecosystem, but your future beliefs will be *less* accurate about many modal propositions: propositions about the specific ways the microecosystem would have evolved over time had you not set it on fire. If you OBSERVE, you'll learn answers to this question and related ones; if you INTERVENE, you won't. And there's an asymmetry here: if you OBSERVE, your beliefs will still be accurate about counterfactuals involving how the system would have evolved if you'd INTERVENED, since by hypothesis, these are predictable *ex ante*.

Reply. If you interpret the case this way, tweak the example: you know in advance that if you merely observe the microecosystem, there'll be many facts about it that you'll never learn: it will be complex enough that many parts will be unobservable. It'll persist long after your death,[22] and its future thereafter will be similarly random and unpredictable. The scenario can be refined to one in which your expectation of the modal accuracy you gamble on gaining in choosing OBSERVE is outweighed by the nonmodal accuracy you can confidently expect to gain from choosing INTERVENE.[23]

Objection #2: Epistemically weighty propositions. OBSERVING doesn't only affect the accuracy of our beliefs about the system. It has the potential to improve our accuracy regarding general scientific laws. This may have a disproportionate impact in two ways: (i) by improving our accuracy wrt. epistemically important propositions, which might be afforded extra weight in their contribution to global accuracy[24]; and (ii) by improving our accuracy wrt. vast swaths of propositions entailed by those laws together with our evidence. And so OBSERVING may still maximize expected accuracy.

Reply. If you interpret the case this way, tweak the example: there are many, many microecosystems like this, all internally unpredictable in how they'll evolve over time. If the system you're considering INTERVENING on has anything important to reveal about more general scientific laws, it's highly likely that one of the other microecosystems will reveal that same information independently. But if you INTERVENE, then at least you can be sure of *this* microecosystem's future and its hidden depths.

With these refinements, of course, it might not feel so important that you *not* INTERVENE. But remember, the zetetic veritist doesn't just say that you're permitted to INTERVENE; it says you're *required* to. And this is implausible: the norms of inquiry do not require you to INTERVENE to make the microecosystem predictable.

epistemic intervention will predictably cause your present and future beliefs about the microecosystem to be more accurate. I discuss the consequences of such cases for epistemic decision theory and accuracy-first epistemology in Carr (2017).

[22] Or, if the relevant accuracy is not agent-relative, everyone's death.

[23] An annoying challenge for both sides: there may be infinitely many propositions in the agenda. This poses challenges for comparing the expected accuracy of one's options. But suppose we stipulate that the inquiry agenda for our microecosystem is finite. For some selections of agendas, OBSERVING will have greater expected accuracy; for others, INTERVENING will. To vindicate the idea that OBSERVING maximizes expected accuracy, one would need to justify prohibiting the second sort of agenda. I discuss this possibility in §5.

[24] See Levinstein (2019) for discussion.

Objection #3: Ethical or aesthetic confounds. Our zetetic intuitions about **Microecosystem** may be clouded by ethical qualms; perhaps plants, or indeed ecosystems, might have moral standing. Or perhaps our intuitions are clouded by aesthetic qualms.

Reply: The perk of this example is that it involves specifically scientific inquiry, often treated as a core case. But one can construct a wide range of examples that illustrate the same point—that zetetic veritism requires interventions to make the world predictable—and that aren't subject to the same worries. Section 4.2 canvases several examples.

Objection #4: INTERVENING isn't an act of inquiry. Maybe improved accuracy in **Microecosystem** provides the "wrong kind of reason" for inquiry. Horowitz (cited in Arpaly, 2014) raises the question: if eating a sandwich before conducting research predictably improves the epistemic value of your research, does that mean that we can have epistemic or zetetic reasons to eat a sandwich? Fleisher (2023) and Palmira (2024) argue that we can't, on the grounds that eating a sandwich isn't an act of inquiry. One might argue that INTERVENING also isn't an act of inquiry, and so zetetic value provides the wrong kind of reason for this act.

Reply: in §2.2, I argued briefly against the claim that there is a distinctive class of acts of inquiry. Even if you're unconvinced, though, the question remains: *why* isn't INTERVENING part of traditional zetetic practice? The arguments I provide below offer an explanation.

The **Microecosystem** case suggests two explananda:

(1) *Major explanandum: permissible non-maximization.* It can be zetetically permissible to choose acts that don't maximize the expected accuracy of one's future beliefs. If zetetic veritism is false, then what are the relevant norms—axiological or decision theoretic—for inquiry?
(2) *Minor explanandum: impermissible maximization.* It can sometimes be zetetically *impermissible* to choose acts that maximize expected accuracy.

4 Diagnosis

4.1 Discovering vs. creating truth

Rough thesis: The aim of inquiry isn't simply to acquire true beliefs or accurate credences. There are two ways to end up with true beliefs: by *discovering* the truth, or by *creating* the truth. Inquiry, and epistemic rationality, don't aim to create truths. The aim of inquiry is to *discover* truths—truths already, in some sense, out there.

Zetetic norms—like epistemic norms—have a *mind-to-world direction of fit*. The problem for zetetic veritism, as a theory of zetetic norms, is that it is neutral regarding direction of fit. It values acts that produce a match between beliefs and the world, but that match can be achieved both by shaping our beliefs to conform to the world and—as in the **Microecosystem** case—by shaping the world to conform to our beliefs.[25]

[25] See Berker (2013b), Carr (2017), and Konek and Levinstein (2019) for discussion.

The rough thesis above provides a straightforward diagnosis for why INTERVENE is not zetetically required, and may even be zetetically impermissible. Creating truths can make it impossible to discover truths.

Precisifying the rough thesis requires considering some subtleties. First, "discovering truths" is misleadingly binary: it suggests the result is knowledge. More reasonably, an aim of inquiry is to increase our gradational accuracy by bringing our beliefs or credences more in line with the world as we find it.

Second, "creating truths" is misleadingly binary: it suggests that the act is causally sufficient for the relevant true proposition. There are acts that raise the objective chance of some propositions without guaranteeing them, and indeed, there are acts that may lead to discovering truths but may also risk creating truths:

Microecosystem #2

Same as **Microecosystem**, but with the option of a mixed strategy: I'll flip a coin. If the coin lands heads, I'll OBSERVE. If the coin lands tails, I'll INTERVENE.

Note that if the distinction between discovering and creating truths is blurry, it doesn't follow that the aim of inquiry is blurry. The zetetic value of a borderline case of creating truths may be indeterminate, but it doesn't follow that it's indeterminate whether creating truths is a goal of inquiry.[26]

Third, "already out there" shouldn't be taken to mean that the relevant truths must concern the present or past, nor that the relevant truths only concern the inhuman world.[27]

So we'll replace this intuitive but potentially misleading terminology with technical terms: an **intervention into a question** Q is an act that causally alters the objective chances of some answers to Q, relative to a presupposed status quo.[28] In the simplest case, one can intervene into Q by making some $q_i \in Q$ true. But one can also intervene in ways that only raise (or lower) the chance of some q_i.

Note that whether some action counts as an "intervention" is *question-relative*: an act may constitute an intervention into Q and not an intervention into Q'. Indeed, an intervention into Q might be zetetically ideal for inquiry into Q'.[29] Note also that what counts as status quo, and therefore as an intervention, may be context-sensitive.

[26] Compare: in beer, I value hoppiness and disvalue dryness. Suppose my only options are a dry, hoppy IPA and a sweet, non-hoppy Belgian ale. It may be indeterminate which of these better satisfies my preferences, but it's not indeterminate whether I value non-hoppiness or disvalue sweetness: I determinately do neither.

[27] Indeed, inquirers may be the subjects of their own inquiry. But when they are, conforming their actions to their earlier predictions about those actions, thereby ensuring that those predictions are true, isn't itself zetetically valuable—even if it maximizes expected accuracy.

[28] Why objective chances? In mere observations of the world, we causally affect our own subjective probabilities of answers to our questions. But this still involves our belief states coming into conformity with the world, rather than the reverse.

[29] I don't assume that at most one question can be actively under inquiry at a time. Roberts (1996) argues that in discourse, questions under discussion are hierarchically organized, where answers to subquestions are evidentially relevant to superquestions; the same is plausible for inquiry outside of language.

The general point: inquiry doesn't require us to influence the facts so as to make those very facts more predictable to us, even if doing so has greater expected accuracy than the alternatives.

4.2 Example interventions into questions under inquiry

We can contrast intervention into Q with (pure) "observation" regarding Q: an **observation** regarding Q is an act that *isn't* an intervention into Q but has a non-zero subjective probability of altering the inquirer's future subjective probabilities for some $q_i \in Q$. Below are additional examples of observational and interventional acts relative to QUIs; in each, the *observational* act is listed first and the *interventional* act second.

Urn

You have an urn full of white and blue balls; you don't know how many of each. A sequence of random draws with replacement is about to occur. The QUI: for a randomly selected draw in the sequence, what's the objective chance that the ball drawn will be blue? Your only options are:

(1) WATCH: observe the draws and update by Bayes' rule.
(2) DYE: pour blue dye into the urn, dying all the balls blue, and infer that the objective chance of drawing a blue ball is 1.

Zetetic veritism requires that you choose DYE.

Dumplings

My partner and I are somewhat *too* accommodating of each other's preferences. I want to inquire into what my partner will want for dinner. (I'm first-order indifferent between dinner options.) But how I word the question may affect how they answer. I have two options:

(1) ASK-WHAT: I can ask: "What would you like for dinner?"
(2) ASK-WHETHER: I can ask: "Would you like dumplings for dinner?"

If I choose ASK-WHAT, there's a decent chance that I'll discover their preferences. But they might instead suggest their best guess about what I prefer. If I choose ASK-WHETHER, there's much stronger chance they'll infer that *I* want dumplings, and—because they prefer for my preferences to be satisfied—I will inadvertently (but predictably) alter their preferences so that they come to want dumplings.

Zetetic veritism requires that I choose ASK-WHETHER.

Birds

You're inquiring into the migratory patterns of a group of birds. The QUI is where these birds will be next January. Their migratory patterns seem complex and unpredictable. You don't have the means to track the birds. Your options are:

(1) RESEARCH: avidly birdwatch, monitor online wildlife webcams, keep track of bird-related social media, censuses, etc.
(2) LURE: build a bird sanctuary and lure the birds to it, offering them food and nesting accommodations so lavish that they'll predictably take up residence there.

If and only if you choose LURE, you can predict with great accuracy with respect to your question: next January, they'll be in your bird sanctuary.

Zetetic veritism requires that I choose LURE.

In each case, the interventional option is expected to result in some belief state with greater expected accuracy than the observational option. And so in each case, the interventional option is required by zetetic veritism. Each interventional option is, however, intuitively peculiar and doesn't seem to serve the relevant inquiry. But whether or not you share that intuition, and whether or not the interventional options are zetetically permissible, they certainly aren't zetetically *required.* These cases therefore present the same problem for zetetic veritism as **Microecosystem**.

4.3 Objections

Objection #5: Sometimes we inquire into our own future actions, and our future actions are trivially interventions into these questions. (*Will I have a beer tonight, even though I'm trying to cut down? Will I go to the gym tomorrow?*) Pursuing these inquiries can be zetetically valuable, even though interventions into the QUI are unavoidable.

Reply. First, notice that natural examples of inquiry into how one *will* behave in the near future often involve speculation about irrational (e.g., akratic) behavior: behavior that one *doesn't* regard as under their present control.

Second, distinguish between *inquiring into how one will act* vs. *deliberating about how to act*. Between inquiry into how one will act and ultimately acting, there's generically a period of deliberation on how to act. In deliberating on how to act, though, one shifts the QUI: the question is no longer how one *will* act, but how one *should* act. In inquiring into Q, one occupies a third-personal perspective, guided by evidence. In deliberating about how to intervene on Q, one occupies a first-personal

perspective, guided by pragmatic considerations. In this sense, perhaps, deliberation crowds out inquiry.[30]

But then, by ultimately acting, one doesn't intervene on the deliberative QUI: one can't influence the normative question of what one *should* do. (Indeed, this is plausibly why examples involving future irrationality come most naturally: they isolate the non-normative question of how one *will* act from normative confounds.)

Objection #6: Suppose I'm inquiring into the question: *what would happen if I were to* φ? Here, φing can be a zetetically valuable way of pursuing my inquiry.

Reply. Where Q is the question *what would happen if I were to* φ?, its answers are propositions expressed by counterfactuals of the form: ⌜*were I to* φ, *q would obtain*⌝. When you φ, you raise the chance that you φ (to 1). But that's not one of the answers to the QUI.

5 A possible alternative to zetetic veritism

Here, I want to discuss the prospects of the following hypothesis:

Zetetic observationalism a proposition q is zetetically relevant—i.e., the accuracy of the agent's future credence in q is a source of contributory zetetic value—only if q is *causally independent* of the agent's choices.

Zetetic observationalism has some benefits—as we'll see, it explains both the major and minor explananda from §3—as well as some costs. This paper sets out the theory without either endorsing or rejecting it.

5.1 Restricted veritism

From the lens of treating zetetic epistemology as an extension of epistemic decision theory, zetetic observationalism might look revisionary and potentially ad hoc. It requires rejecting a core commitment of accuracy-first epistemology:

Extensionality Let c range over propositions $\mathcal{F} = \langle q_1 \dots q_n \rangle$ and c^* over $\mathcal{F}^* = \langle q_1^* \dots q_n^* \rangle$. Relative to any admissible epistemic value function $\mathfrak{e}$ and worlds w, w^*, if $c(q_i) = c^*(q_i^*)$ and $v_w(q_i) = v_{w^*}(q_i^*)$ for all $i \leq n$, then $\mathfrak{e}_{\mathcal{F}}(c, w) = \mathfrak{e}_{\mathcal{F}^*}(c^*, w^*)$.[31]

[30] There's a relation between the problems for zetetic veritism I've discussed and the literature on the "deliberation crowds out prediction" (DCOP) thesis, defended in Spohn (1977); Levi (2007); Price (2007). Motivating examples for DCOP often involve self-frustrating decisions (e.g., the familiar Death in Damascus case from Gibbard and Harper (1978)). Self-gratifying cases are generally less concerning for practical decision theory, but in the epistemic case, self-verifying cases are arguably just as puzzling. Note that unlike proponents of DCOP, this paper is not committed to the claim that in deliberative contexts, it's irrational to hold any credence in propositions about which act one will choose.

[31] This is a notational variant of Extensionality as defined by Joyce (2009); see also Joyce (1998).

Extensionality requires that the epistemic value of a credence function is a function solely of credence assignments and truth value assignments to the propositions—not their practical usefulness, fruitfulness, or even specific propositional content.

On the other hand, within the literature on inquiry, it's commonly held that not all propositions contribute equally to zetetic value: number of blades of grass, contents of phonebooks, etc. Call this view **restricted veritism.** On one variant of this view, it's an objective (though perhaps contingent) matter which propositions are relevant for inquiry, and to what degree; on another, only possible answers to the question(s) under inquiry are zetetically relevant.

5.2 Observational accuracy

According to restricted veritism, only accuracy with respect to **zetetically relevant** propositions contributes to zetetic value. And which propositions are zetetically relevant is world-, agent-, and time-relative. A center $\acute{c}$ is an ordered pair of an agent and time. Where $\mathcal{F}^a \subseteq \wp(W)$ is the range of agent a's credence function, let $\mathcal{Z}^{\mathbb{c}} \subseteq \mathcal{F}^a$ be the set of zetetically relevant propositions at $\mathbb{c}$, under that description. $\mathcal{Z}^{\mathbb{c}}_w$ is $\mathcal{Z}^{\mathbb{c}}$'s value at w.[32]

Which propositions are zetetically relevant at $\mathbb{c} = \langle a, t\rangle$? Different variants of restricted veritism offer different answers: $\mathcal{Z}^{\mathbb{c}}$ may be restricted to theoretically "important" propositions or to the QUI. Zetetic observationalism imposes a necessary (but arguably not sufficient) condition: if at t, a can affect the objective chance of a proposition q, then $q \notin \mathcal{Z}^{\mathbb{c}}$, and at t, future gains in accuracy about q are not contributory sources of zetetic value at $\mathbb{c}$. Letting Φ be the partition of options available at $\mathbb{c} = \langle a, t\rangle$:

Causal independence condition for all $q \in \mathcal{F}^a$, $q \in \mathcal{Z}^{\mathbb{c}}$ only if, for all $\varphi \in \Phi$, $ch(q \mid\mid \varphi) = ch(q)$.

In English: q is zetetically relevant only if none of the agent's available options would influence the chance of q.

A variety of propositions satisfy the causal independence condition for zetetic relevance: among them, propositions about events that take place before t, noncontingent propositions, future contingent propositions that a cannot influence (e.g., in Minkowski spacetime, events outside a's upper light cone at t), and so on.

But what about our motivating examples, where an agent is inquiring into questions whose answers are not causally independent of the agent's present available actions? In these cases, according to zetetic observationalism, improving your accuracy wrt. answers to the QUI isn't zetetically valuable.

Importantly, however, zetetic observationalism doesn't entail that inquiry into these questions is valueless. Instead, the inquiry's value is derived from improvements in accuracy about other propositions.

Suppose an agent is inquiring into Q and the members of Q do not satisfy the causal independence condition. The agent can improve her accuracy with respect to

[32] See §5.4 and appendix A.2 for the resulting decision theory and its technicalities.

a related set of propositions that the causal independence condition doesn't exclude from $\mathcal{Z}^{\mathfrak{C}}$: propositions about how answers to Q do or do not depend on the agent's available acts. In the Urn case, e.g., it isn't valuable to learn: *this ball I've drawn is blue*; it is valuable to learn: *this ball I've drawn would've been blue whether or not I'd chosen* DYE.

Notice: if we assume that only answers to the QUI can be zetetically relevant—i.e. members of $\mathcal{Z}^{\mathfrak{C}}$—then in cases where answers to the QUI fail the causal independence condition, $\mathcal{Z}^{\mathfrak{C}}$ is empty. In that case, considerations of zetetic value can't decide between options. And so one might conclude that zetetic observationalism is incompatible with the question-relative notion of zetetic value. But the proposal I'm suggesting is entirely compatible with question-relative zetetic value: even in contexts where answers to the QUI are themselves zetetically irrelevant, the QUI may still play a central role in determining which propositions are zetetically relevant.

There are various options for how to derive zetetically relevant propositions from an agent's QUI(s), together with other features of context. One option involves *dependency hypotheses*, a familiar tool from pragmatic decision theory for representing how outcomes depend on an agent's presently available acts. Dependency hypotheses are, by design, causally independent of an agent's acts.[33] Where $Q_1, \ldots, Q_n$ are a's QUIs at t, let $\mathcal{S}$ be the maximal partition that's a refinement of the cover $\bigcup_{i=1}^{n} Q_i$.[34] A **zetetic dependency hypothesis** can be represented as a maximal consistent conjunction of non-backtracking counterfactuals of the form $\ulcorner \varphi_j \mathbin{\Box\!\!\rightarrow} q_i \urcorner$, where $\varphi_j \in \Phi$ and $q_i \in \mathcal{S}$.[35] A *maximal consistent conjunction* of these counterfactuals contains, for each $\varphi_j \in \Phi$, exactly one conjunct of the form $\ulcorner \varphi_j \mathbin{\Box\!\!\rightarrow} q_i \urcorner$ for some $q_i \in \mathcal{S}$. A dependency hypothesis itself is the propositional content of such a conjunction: $[[(\varphi_1 \mathbin{\Box\!\!\rightarrow} q_i) \wedge \ldots \wedge (\varphi_m \mathbin{\Box\!\!\rightarrow} q_j)]]$.

Zetetic dependency hypotheses satisfy the causal independence condition and so, on this hypothesis, are zetetically relevant. Improving one's accuracy wrt. these propositions is a source of contributory zetetic value.

There are other ways of representing the relevant dependency structures: for example, causal models, e.g. causal Bayes nets (Pearl, 2000; Spirtes et al., 2000) where Φ appears as an exogenous variable and QUIs are included among variables. So zetetic observationalism doesn't depend for its plausibility on any particular semantics for counterfactuals. What's important is that the relevant information capture dependency relations between available options at t and possible answers to QUIs at t, where these dependency relations aren't themselves dependent on the agent's options at t.[36]

[33] Dependency hypotheses are introduced in Lewis's (1981) causal decision theory; on the present hypothesis, they play a second, very different role within zetetic decision theory, as propositional loci of contributory zetetic value.

[34] In other words, assuming again that each Q_i is a partition, $\mathcal{S}$ is the meet of $Q_1, \ldots, Q_n$.

[35] For readability, in counterfactuals I use φ_j and q_i as variables both for propositions and sentences expressing those propositions.

[36] This excludes, for example, contents of some $\ulcorner \varphi_j \mathbin{\Box\!\!\rightarrow} q_i \urcorner$ counterfactuals from $\mathcal{Z}^{\mathfrak{C}}$: sometimes, if q_i is true, the agent can raise the chance of $[[\varphi_j \mathbin{\Box\!\!\rightarrow} q_i]]$ to 1 by $\varphi_j$ing.

5.3 Explanatory benefits

Zetetic observationalism straightforwardly explains our major explanandum: why it can be zetetically permissible to φ even when φing doesn't maximize expected accuracy. What about our minor explanandum: that it can sometimes be zetetically *im*permissible to ψ even when ψing maximizes expected accuracy? One might think zetetic observationalism makes interventions into QUIs zetetically *irrelevant*, and so unhelpful for the minor explanandum. But it doesn't: intervening into QUIs will impede improvements in future expected accuracy with respect to some zetetically relevant propositions.

Return to the Microecosystem case. Suppose you're inquiring into some Q about the future of the microecosystem: for example, how many plant species it will contain at some future time. If you INTERVENE, you reduce your opportunity to improve your future accuracy about hypotheses of the form $[[\textsc{observe} \boxdot\!\!\rightarrow q_i]]$ for various $q_i \in Q$. And you gain little or no accuracy about hypotheses of the form $[[\textsc{intervene} \boxdot\!\!\rightarrow q_i]]$ to compensate. After all, *ex ante*, you're in a position to accurately predict that if you were to INTERVENE, the number of plant species in the microecosystem would be zero. You can expect to gain little or no accuracy on that proposition by INTERVENING rather than OBSERVING: your degree of accuracy in propositions about what would happen if you INTERVENE will be high either way.

Another example: return to the Urn case above. *Ex ante,* you can confidently predict that whatever you choose, you'll end up with accurate credences about $[[\textsc{dye} \boxdot\!\!\rightarrow ch(\text{blue}) = x]]$, for any x: accurate high credence if $x = 1$, and accurate low credence for all other $x \in \mathbb{R}$. But while you'll acquire incremental evidence relevant to $[[\textsc{draw} \boxdot\!\!\rightarrow ch(\text{blue}) = x]]$ for each x if you DRAW, and therefore increase your expected accuracy wrt. these propositions, you won't acquire any such evidence if you DYE.

One might expect, on these grounds, that interventions are universally zetetically disvaluable. But intuitively, some acts that count as interventions into a system are zetetically valuable, and zetetic observationalism helps to explain that fact. In the Microecosystem case, for example, it might be zetetically valuable to introduce a new species into the Microecosystem to see how it interacts with other species and with abiotic components of the ecosystem. This form of intervention helps us gain accuracy about a variety of questions: questions about the causal impact of introducing that species into the environment, questions about properties of the Microecosystem before the introduction of the new species, etc.

Notice, though, that these are questions whose answers satisfy the causal independence condition, and therefore that we can't presently intervene on: the former question's answers characterize the causal dependency structure of our present acts, and the latter's answers are propositions about the past. In other words, introducing a new species is an intervention into the *microecosystem*, but not necessarily an intervention into the *question under inquiry*. And the zetetic benefits it affords are related to learning truths that are causally independent of our actions. Of course, we also learn the answer to the question: *will this particular new species ever enter into the micro-*

ecosystem?. But because we learn this merely by making one of its answers true, that increased accuracy contribution does not affect zetetic value.[37]

5.4 Decision theory for zetetic observationalism

Zetetic observationalism requires certain modifications on §2.2's decision theory, because whether propositions are zetetically valuable is world-relative. Whether $q \in \mathcal{Z}^{\mathbb{c}}_w$ depends on whether $\mathbb{c}$ can influence the chance of q at w. As a result, the cardinality $|\mathcal{Z}^{\mathbb{c}}|$ is world-relative. The calculation of the zetetic value of a credence function needs modification's to avoid penalizing worlds where there are more zetetically relevant propositions to be imperfectly accurate about. For technicalities, see appendix A.2.

Moreover, an agent can be uncertain whether she can influence the chance of q, and so uncertain about the contributory value of accuracy wrt. q. So even when q is *in fact* zetetically irrelevant, it can still influence the *expected* zetetic value of an act.

5.5 Possible costs of zetetic observationalism

5.5.1 Learning answers to QUIs often isn't valuable for inquiry

If zetetic observationalism is true, then in a wide range of ordinary contexts, *learning the answer to QUIs isn't valuable for inquiry.* This is so even if your QUI is zetetically important. That's not especially intuitive.

There is, however, zetetic value in learning a variety of related propositions: for example, propositions about whether an answer to the QUI would be true if you did or didn't intervene into the QUI. And unions or intersections of such propositions might be equivalent to answers to your QUI—which may make this unintuitive consequence tolerable.

It might be thought that the question-relative zetetic veritist can avoid this possible cost by providing an alternative diagnosis of the motivating examples for zetetic observationalism: perhaps, in these examples, the *real* QUIs are ones whose answers are not under the agent's control.

One benefit of zetetic observationalism, though, is that it doesn't require any stipulations about the QUI in a context. The QUI may be, in many contexts, up to the agent (both descriptively and normatively); but the causal independence condition is not. The constraint is compatible with various intuitive characterizations of the QUI—e.g., as a question the agent has consciously wondered about—even in cases where such characterizations result in the prediction that the QUI's answers are zetetically irrelevant.

The constraint does have ramifications for determining which questions are zetetically valuable to adopt as one's QUI. But interestingly, it needn't motivate a dispreference for questions whose answers the agent can control. Suppose Q's answers fail the causal independence condition, but Q determines a set of zetetically relevant propositions $Q^\star$ (e.g. a set of zetetic dependency hypotheses about how Q's answers

[37] Thanks to Ryan Doody for example and discussion.

do or do not depend on the agent's actions). Then, in selecting a QUI, it needn't matter whether the agent chooses Q (which she can control) or $Q^\star$ (which she can't): the same set of propositions, $Q^\star$, are zetetically relevant either way.[38]

5.5.2 Epistemic and zetetic value aren't unified

If zetetic observationalism is the correct theory of zetetic norms, it won't work for *epistemic* norms. So zetetic norms cannot be understood as a generalization of epistemic norms. After all, accuracy-first epistemology derives basic coherence norms (probabilism for credences, consistency for beliefs) from accuracy. Suppose accuracy wrt. zetetically irrelevant propositions makes no contribution to *epistemic* value. Then incoherence in our attitudes toward these propositions is no worse than coherence. The resulting theory—epistemic observationalism—won't be able to explain why incoherence toward these propositions is epistemically rational.

One might think that we can recover coherence norms for such propositions by looking at disjunctions or conjunctions of zetetically relevant propositions. After all, it's often accepted that $\ulcorner(p \boxrightarrow q) \wedge (\neg p \boxrightarrow q)\urcorner$ is equivalent to q.[39] So one might think we'd be able to recover accuracy constraints on a zetetically irrelevant q from accuracy constraints on $[[p \boxrightarrow q]]$ and $[[\neg p \boxrightarrow q]]$ if both are zetetically relevant. But even if they are, their intersection need not be. After all, if q is under the agent's control, so is any proposition equivalent to it. And so accuracy-based considerations governing beliefs about the conjuncts won't require them to cohere rationally with beliefs about the conjunction.

Indeed, this example shows that epistemic observationalism would allow incoherence affecting zetetically *relevant* propositions. For example, suppose $q \in \mathcal{Z}^{\mathfrak{c}}$ is equivalent to some union $\bigcup(p_1, \ldots, p_n)$ where $p_1, \ldots, p_n \notin \mathcal{Z}^{\mathfrak{c}}$. Epistemic observationalism won't predict any irrationality where for some p_i, $c(q) < c(p_i)$—i.e., in effect, having lesser confidence in a disjunction than in one of its disjuncts.

Epistemic norms govern *all* beliefs—not just beliefs about propositions we have no control over, and not just beliefs that are zetetically relevant. It's irrational to believe a contradiction about the number of blades of grass in the lawn even if it's irrelevant to your inquiries. So zetetic observationalism has a potential cost: it means that epistemic and zetetic norms are surprisingly disunified, in their axiology or deontic theory or both.

6 Conclusion

The central claim of this paper is that the norms of inquiry don't always require us to maximize the expected accuracy of our future beliefs. In cases where we can best improve our accuracy by forcing the world to conform to our predictions, or to become more predictable, the norms of inquiry needn't require us to do so. The facts that matter for inquiry are those we can only discover, not control. The problem for

[38] Thanks to an anonymous referee for raising this question.

[39] Assuming a strong centering constraint on counterfactuals.

the accuracy-first project in inquiry is similar to its problems in epistemology: that it ignores belief's mind-to-world direction-of-fit, and so sometimes demands that we fit the world to our minds. The positive zetetic normative theory discussed in §5 avoids this problem; but the theory may disappoint if we want a more unified treatment of the direction of fit problem for both the zetetic and the epistemic. Whether such a treatment is forthcoming might depend on whether zetetic norms (governing action) and epistemic norms (governing belief) have fundamentally different decision theories,[40] or values, or both—or whether they'll be revealed as unified by the so-far-undiscovered correct theories of each.[41]

A Appendix

A.1 Zetetic value of sequences of credence functions

Throughout we idealize to the finite case. Where $\mathbf{t}_w$ is the sequence of relevant times in w and $\mathbf{c}_w = \langle \mathrm{c}_w^1, \ldots, \mathrm{c}_w^{|\mathbf{t}_w|} \rangle$ is your sequence of credence functions over relevant times at w, the aggregated accuracy of $\mathbf{c}$ at w might be its mean accuracy over relevant times:

$$\mathfrak{z}(\mathbf{c}_w, w) = \frac{1}{|\mathbf{t}_w|} \sum_{\mathrm{c}_w^i \in \mathbf{c}_w} \mathfrak{z}(\mathrm{c}_w^i, w)$$

Averaging allows us to relativize the cardinality of relevant times by world without penalizing "longer" worlds.

The expected zetetic value of one's act is given by one's expectation of the average accuracy of one's future credence functions over relevant times:

$$\begin{aligned} \mathbb{E}_c\mathfrak{z}(\varphi) &= \sum_{w \in W} c(w \mid\mid \varphi) \cdot \mathfrak{z}(\mathbf{c}_w, w) \\ &= \sum_{w \in W} \left(c(w \mid\mid \varphi) \cdot \frac{1}{|\mathbf{t}_w|} \sum_{\mathrm{c}_w^i \in \mathbf{c}_w} \mathfrak{z}(\mathrm{c}_w^i, w) \right) \end{aligned}$$

A similar maneuver allows for a generalization of zetetic value to aggregate the accuracy of the credence functions of multiple agents.

[40] See Carr (2017) and Konek and Levinstein (2019) for discussion.

[41] Thanks to Ryan Doody, Melissa Fusco, Richard Pettigrew, David Thorstad, Christopher Willard-Kyle, and audiences at the WIP Inquiry Group (2021) and the Inquiry Conference in Glasgow (2022) for valuable discussion.

A.2 Observational zetetic value

On the account in §5, the zetetic value of a credence function c is not its global accuracy over $\mathcal{F}^a$, but rather over $\mathcal{Z}^{\mathfrak{c}}_w$. If, following §2.2, we compute this additively, though, then worlds where more propositions are zetetically relevant will be penalized: anything less than perfect accuracy wrt. additional propositions will contribute more inaccuracy.

Another way to see this point: worlds with fewer relevant propositions are rewarded, because additive aggregation will treat *not contributing zetetic value* as equivalent to *contributing zetetic value 0*. Recall that because inaccuracy is divergence from the truth, zero divergence means maximal accuracy. But irrelevant propositions shouldn't be equal in value, irrespective of one's credence, to perfect accuracy about relevant propositions. Instead, they should be excluded from affecting zetetic value entirely.

So we need two modifications to (ZETETIC VALUE): (i) the relevant range of credences is $\mathcal{Z}^{\mathfrak{c}}_w$, not $\mathcal{F}^a$; (ii) accuracy over $\mathcal{Z}^{\mathfrak{c}}_w$ is *mean* accuracy over its members, not additive:

$$\mathfrak{z}(c, w) = -\frac{1}{|\mathcal{Z}^{\mathfrak{c}}_w|} \sum_{q \in \mathcal{Z}^{\mathfrak{c}}_w} \mathfrak{b}(c, q, w)$$

If the zetetic value of an act aggregates relevant accuracy across a sequence of future times, then these too will be averaged:

$$\begin{aligned} \mathfrak{z}(\mathbf{c}_w, w) &= \frac{1}{|\mathbf{t}_w|} \sum_{\mathfrak{c}^i_w \in \mathbf{c}_w} \mathfrak{z}(\mathfrak{c}^i_w, w) \\ &= -\frac{1}{|\mathcal{Z}^{\mathfrak{c}}_w| \cdot |\mathbf{t}_w|} \sum_{\mathfrak{c}^i_w \in \mathbf{c}_w} \sum_{q \in \mathcal{Z}^{\mathfrak{c}}_w} \mathfrak{b}(\mathfrak{c}^i_w, q, w) \end{aligned}$$

Funding Not applicable.

Declarations

Conflict of interest Not applicable.

References

Ahlstrom-Vij, K., & Dunn, J. (2014). A defence of epistemic consequentialism. *Philosophical Quarterly, 64*(257), 541–551.

Arpaly, N. (2014). Epistemology and sandwiches. https://theviewfromtheowlsroost.com/2017/10/29/epistemology-and-sandwiches/.

Beddor, Bob. (2023). Inquiry beyond knowledge. *Philosophy and Phenomenological Research, 109*(1), 330–356.

Berker, S. (2013). Epistemic teleology and the separateness of propositions. *Philosophical Review, 122*, 337–393.
Berker, S. (2013). The rejection of epistemic consequentialism. *Philosophical Issues, 23*(1), 363–387.
Caie, M. (2013). Rational probabilistic incoherence. *Philosophical Review, 122*(4), 527–575.
Carr, J. R. (2017). Epistemic utility theory and the aim of belief. *Philosophy and Phenomenological Research, 95*(3), 511–534.
Carr, J. R. (2019). A modesty proposal. *Synthèse, 198*(4), 3581–3601.
Ciardelli, I., Groenendijk, J., & Roelofsen, F. (2018). *Inquisitive semantics*. Oxford University Press.
Cohen, S. (2016). Theorizing about the epistemic. *Inquiry: An Interdisciplinary Journal of Philosophy, 59*(7–8), 839–857.
Easwaran, K. (2016). Dr. Truthlove: or, how i learned to stop worrying and love Bayesian probabilities. *Noûs, 50*(4), 816–853.
Falbo, A. (2023). Inquiring minds want to improve. *Australasian Journal of Philosophy, 101*(2), 298–312.
Firth, R. (1978/1998). The Schneck Lectures, Lecture 1: Epistemic Utility. In J. Troyer (ed.), *In Defense of Radical Empiricism: Essays and Lectures*, 317–33. Lanham, MD: Rowman & Littlefield. First delivered as a lecture at the University of Arizona, 1978.
Fleisher, W. (2023). Intellectual courage and inquisitive reasons. *Philosophical Studies, 180*(4), 1343–1371.
Friedman, J. (2020). The epistemic and the zetetic. *Philosophical Review, 129*(4), 501–536.
Friedman, J. (forthcoming). Zetetic epistemology. In *Towards an expansive epistemology: Norms, action, and the social sphere.*, Routledge
Gibbard, A., & Harper, W. L. (1978). Counterfactuals and two kinds of expected utility. In A. Hooker, J. J. Leach, & E. F. McClennen (Eds.), *Foundations and applications of decision theory* (pp. 125–162). D. Reidel.
Goldman, A. I. (1979). What is justified belief. In G. Pappas (Ed.), *Justification and knowledge* (pp. 1–25). D. Reidel.
Good, I. J. (1962). Subjective probability as the measure of a non-measurable set. In P. Suppes, E. Nagel, & A. Tarski (Eds.), *Logic methodology and philosophy of science* (pp. 319–329). Stanford University Press.
Greaves, H. (2013). Epistemic decision theory. *Mind, 122*(488), 915–952.
Greaves, H., & Wallace, D. (2006). Justifying conditionalization: Conditionalization maximizes expected epistemic utility. *Mind, 115*(459), 607–632.
Groenendijk, J., & Martin, S. (1984). On the semantics of questions and the pragmatics of answers. In F. Landman & F. Veltman (Eds.), *Varieties of formal semantics: Proceedings of the Fourth Amsterdam Colloquium* (pp. 143–170). Foris.
Hamblin, C. L. (1973). Questions in Montague English. *Foundations of Language, 10*(1), 41–53.
Joyce, J. M. (1998). A nonpragmatic vindication of probabilism. *Philosophy of Science, 65*(4), 575–603.
Joyce, J. M. (2009). Accuracy and coherence: Prospects for an alethic epistemology of partial belief. In F. Huber & C. Schmidt-Petri (Eds.), *Degrees of Belief* (Vol. 342, pp. 263–297). Springer.
Kelp, C. (2014). Two for the knowledge goal of inquiry. *American Philosophical Quarterly, 51*(3), 227–32.
Kelp, C. (2018). Inquiry and the transmission of knowledge. *Philosophy and Phenomenological Research, 99*(2), 298–310.
Kelp, C. (2021). *Inquiry, Knowledge, and Understanding*. Oxford University Press.
Konek, J., & Levinstein, B. A. (2019). The foundations of epistemic decision theory. *Mind, 128*(509), 69–107.
Leitgeb, H., & Pettigrew, R. (2010). An objective justification of Bayesianism I: Measuring inaccuracy. *Philosophy of Science, 77*(2), 201–235.
Leitgeb, H., & Pettigrew, R. (2010). An objective justification of Bayesianism II: The consequences of minimizing inaccuracy. *Philosophy of Science, 77*(2), 236–272.
Levi, I. (2007). Deliberation does crowd out prediction. In J. Josefsson T. Ronnow-Rasmussen, B. Petersson & D. Egonssson (Eds.), *Hommage à Wlodek: Philosophical Papers Dedicated to Wlodek Rabinowicz*, https://www.fil.lu.se/hommageawlodek/site/papper/LeviIsaac.pdf.
Levinstein, B. A. (2019). Imprecise epistemic values and imprecise credences. *Australasian Journal of Philosophy, 97*(4), 741–760.
Lewis, D. (1981). Causal decision theory. *Australasian Journal of Philosophy, 59*, 5–30.
Littlejohn, C. (2015). Who cares what you accurately believe? *Philosophical Perspectives, 29*(1), 217–248.
Littlejohn, C. (2018). The right in the good: A defense of teleological non-consequentialism in epistemology. In K. Ahlstrom-Vij & J. Dunn (Eds.), *Epistemic consequentialism* (pp. 23–47). Oxford University Press.

Palmira, M. (2024). Higher-order evidence and the duty to double-check. *Noûs, 58*(3), 799–824.
Pearl, J. (2000). *Causality: Models reasoning and inference*. Cambridge University Press.
Pettigrew, R. (2016). *Accuracy and the laws of credence*. Oxford University Press.
Pettigrew, R. (2018). Making Things right: The true consequences of decision theory in epistemology. In H. Kristoffer & J. D. Ahlstrom-Vij (Eds.), *Epistemic consequentialism* (pp. 220–240). Oxford University Press.
Price, H. (2007). The effective indexical. Unpublished manuscript.
Rinard, S. (2017). No exception for belief. *Philosophy and Phenomenological Research, 94*(1), 121–143.
Roberts, C. (1996). Information structure in discourse: towards an integrated formal theory of pragmatics. In Jae-Hak Yoon & Andreas Kathol (Eds.), *Papers in semantics (Working Papers in Linguistics 49*, vol. 6 67, The Ohio State University.
Salow, Bernhard, & Ahmed, Arif. (2019). Don't look now. *British Journal for the Philosophy of Science, 70*(2), 327–350.
Sapir, Y., & van Elswyk, P. (2021). Hedging and the ignorance norm on inquiry. *Synthese, 199*(3–4), 5837–5859.
Singer, D. J. (2018). How to be an epistemic consequentialist. *Philosophical Quarterly, 68*(272), 580–602.
Spirtes, P., Glymour, C., & Scheines, R. (2000). *Causation, Prediction, and Search*. MIT Press.
Spohn, W. (1977). Where Luce and Krantz do really generalize savage's decision model. *Erkenntnis, 11*(1), 113–134.
Stalnaker, R. (1978). Assertion. *Syntax and Semantics (New York Academic Press), 9*, 315–332.
Stalnaker, R. C. (1970). Pragmatics. *Synthese, 22*(1–2), 272–289.
Thorstad, D. (2024). *Inquiry under bounds*. Oxford University Press.
Thorstad, D. (forthcoming). Norms of inquiry. *Philosophical Topics*.
Whitcomb, D. (2010). Curiosity was framed. *Philosophy and Phenomenological Research, 81*(3), 664–687.