

Epistemic Insurance Policies

Conner Schultz

Abstract

We often cannot tell whether our beliefs have been shaped by motivated reasoning, implicit bias, echo chambers, polarization, misinformation, pernicious ideology, and the like—or so I argue. It typically appears to us as though we are in the epistemic good case, even when we are in the bad case. This places us in positions of epistemic risk: for all we can tell, we might well be in the bad case. What should we believe under such conditions? One response is to go “externalist”: keep your beliefs if you are in the good case; otherwise, give them up. I argue, however, that the externalist strategy should be rejected: by effectively instructing us to ignore the risk that we are in the bad case, externalism counsels an objectionable form of intellectual negligence. In place of externalism, I defend a framework for the kind of epistemic norms we should accept, which I term “epistemic insurance policies.” The core idea is simple: just as we ought to take out ordinary insurance to protect ourselves against practical risks, we should adopt epistemic norms that safeguard us in the event that we are in a bad epistemic position. I argue that, although they impose costs when we are in the good case, these costs function as premiums: they are worth paying, all things considered. I conclude with broader remarks about how the insurance framework unifies familiar norms across different domains and clarifies their shared rationale.

Introduction

You read two studies on a politically charged question. One confirms what you already believe; the other doesn’t. You examine their arguments and conclude, with apparent rationality, that the confirming study is right. Yet you also know that humans are prone to motivated reasoning—that our evaluation of evidence is often covertly shaped by what we antecedently desire to be true.

Can you tell whether you’ve fallen prey to motivated reasoning? Usually not—or so I argue. This predicament generalizes to other dimensions of our nonideal epistemic agency and our epistemically unhealthy environments: we often cannot tell whether our beliefs are

shaped by bias, misinformation, echo chambers, or bad ideology. More to the point: we often cannot tell whether we are in a good or bad epistemic position. I argue that such *epistemic risk* demands a response from us. In particular, it calls for distinctive epistemic norms—“Epistemic Insurance Policies”—that insulate us from the harms of being in a bad epistemic position.

To give a preview of my argument, the analogy with ordinary insurance is instructive. We buy home or health insurance because, while we hope to never need it, we can’t rule out the possibility that disaster will strike, and it’d be *negligent* not to protect ourselves against that risk. Likewise, even when our epistemic position *seems* good, ignoring the risk that it isn’t is epistemically negligent. Epistemic insurance policies are norms designed for precisely this predicament: they carry costs when all goes well, but safeguard us when things go badly.

The plan is as follows. In §1 I argue that we often cannot tell whether we are in a good or bad epistemic position. In §2 I outline one response to this dilemma—the suggestion that we should go “externalist.” In §3 I argue that the externalist response faces the joint problems of unfollowability and epistemic risk. In §4-5 I outline and defend the framework of epistemic insurance. In §6 I demonstrate the power of this framework in unifying and fending off a common objection to a variety of epistemic norms.

§1 The Good Case and the Bad Case

Let’s begin with something that has a skeptical sound to it: the distinction between the good case and the bad case.¹ In external world skepticism, the good case is the ordinary one in which your perceptual appearances accurately represent the world around you, while the bad

¹ The terms “good case” and “bad case” were popularized by Williamson (2000).

case is the skeptical scenario in which you—a brain in a vat—are radically deceived. Across both cases, your perceptual appearances are identical: in both cases, it seems to you as if you are in the good case, even when you are not.² Indeed, most epistemologists hold that you *cannot tell* which case you’re really in. The sense in which you “cannot tell” is that your belief is *insensitive*: you would believe you are in the good case even if you weren’t.³

I borrow this distinction not to revive skepticism, but to apply it to our epistemic lives more broadly: to the distinction between being in a good *epistemic position*—where things are roughly as they seem—and a bad one—where things aren’t.

It’s a truism that we’re nonideal epistemic agents in degraded epistemic environments.⁴ We’re prone to bias and motivated reasoning; we forget evidence; we fall into echo chambers; we polarize; we absorb ideology that justifies unjust systems of oppression; we place trust in untrustworthy people; apparently rational peers routinely disagree; and so on.

We can design good and bad cases along any of these dimensions:

Good Case	Bad Case
Your belief did not result from cognitive bias.	Your belief resulted from cognitive bias.
Your belief did not result from an echo chamber.	Your belief resulted from an echo chamber.
Your belief did not result from bad ideology.	Your belief resulted from bad ideology.

² This claim is routinely accepted but not uncontroversial, cf. Williamson (2000).

³ The notion of sensitivity was popularized by Nozick (1981).

⁴ Cf. Blake-Turner (2020) for more on the idea of epistemic environments.

These examples are only rough approximations. Along any given dimension, we might quibble about what really counts as the good or the bad case.⁵ Furthermore, the good case along any one axis need not be good *tout court*: a belief formed outside an echo chamber, for instance, might still be unjustified all things considered. The point is only that, along each dimension, there is something epistemically better about being in the good case and something epistemically worse about being in the bad one. Many other dimensions could be added—for misinformation, polarization, disagreement, and so on—but these suffice to illustrate the pattern.⁶

Hume famously quipped that it is easy to dismiss skeptical hypotheses—the bad cases of skepticism—as soon as we leave the study. Not so for these more *immanent* bad cases.

None of us are immune to the bad case. We know, in general, that we’re susceptible to bias, motivated reasoning, echo chambers, ideology, and misinformation. Yet we usually cannot tell which case we’re in. As above, the relevant sense in which we “can’t tell” is that our beliefs are *insensitive* to which case we’re actually in: we would believe we’re in the good case even if we weren’t.

Consider motivated reasoning. It rarely *appears* to us that we’re engaging in it; on the contrary, it typically appears as though our reasoning is impartial.⁷ Psychological work on the “introspection illusion” and “bias blind spot” supports this: we recognize bias in others more readily than in ourselves.⁸ If I knew I had arrived at my belief by motivated reasoning, I could not coherently maintain that belief, since I would see it as transparently ill-founded.⁹

⁵ See, for instance, the discussion of Lackey (2021) in §2.

⁶ The epistemic import of polarization is contentious. See Kelly (2008), McWilliams (2021), and Dorst (2023) for more on the controversy.

⁷ Cf. Kunda (1990).

⁸ Cf. Nisbett & Wilson (1977); Pronin, Lin, & Ross (2002); Pronin, Gilovich, & Ross (2004); Wilson (2002). See also Ahlstrom-Vij (2013) and the references therein.

⁹ Cf. Carter & McKenna (2020).

Of course, I don't *always* believe that I've avoided motivated reasoning. Sometimes I observe myself from a third-person perspective, notice that I have an interest in a particular answer, and infer that I reasoned in a motivated way.¹⁰ Yet such self-diagnoses are rare, and even our third-personal evaluations are themselves susceptible to motivated reasoning. Thus, it seems plausible that I (often) cannot tell whether I have reasoned in a motivated way.

Similar points apply to the other epistemic dimensions of interest. Echo chambers, for example, need not *exclude* opposing evidence outright; on the contrary, they typically include but *discredit* it in a biased way.¹¹ From within one, it's easy to see oneself as an honest inquirer who's merely judged the opposing side to be unconvincing. The opacity of motivated reasoning compounds the problem: if I cannot tell whether I'm reasoning in a motivated way, then I cannot tell whether I'm dismissing opposing evidence on good grounds.

Importantly different considerations apply to bad ideology. Ideology works (partly) by systematically limiting what evidence is available for certain propositions.¹² Yet the incompleteness of one's evidence is, of course, difficult to detect from within. Furthermore, unlike motivated reasoning and echo chambers (which exploit our limitations as non-ideal agents), the opacity of ideology poses a problem even for *ideal* agents, for there's nothing about perfect rationality that enables one to see when the evidential field itself has been restricted.

Generalizing, our epistemic situation often resembles the skeptical one. Even when things go well, it appears to us as though they are going well; and when things go poorly, it

¹⁰ Cf. Worsnip (2023).

¹¹ Cf. Nguyen (2020a; 2020b).

¹² This idea traces back to Marx's claim that the ruling class controls the "means of mental production," and has been developed by Althusser (1971), Gramsci (1971), and Chomsky & Herman (1988), among others in critical theory and standpoint epistemology (cf. Geuss 1981; Intemann 2010; Shelby 2003; Haslanger 2017; and the essays in *Analyzing Ideology*).

still appears that way. In either case, the appearances are functionally identical. In this respect, we often cannot tell whether we are in the good case or the bad case.

The question is: which epistemic norms should we accept when we cannot tell whether we're in the good case or the bad case?

§2 The Externalist Response

Drawing again on the literature on skepticism, one answer is go *externalist*: if you're in the good case, keep your beliefs; if you're in the bad case, give them up.¹³ Call this "the externalist norm." Crucially, the externalist norm doesn't require you to be able to *tell* which case you are in: all that matters for whether you ought to keep or give up your beliefs is which case you're really in.

Many philosophers have endorsed instances of the externalist norm in response to particular topics in social and applied epistemology.

First, Lackey (2021) on echo chambers. Though she doesn't use this terminology, Lackey distinguishes "good" echo chambers, which reliably produce true beliefs in their members, from "bad" echo chambers, which peddle falsehoods. She then endorses this norm: if you're in a good echo chamber, stay in it; otherwise, leave.

Second, Titelbaum (2015) on peer disagreement. Given his "Fixed Point Thesis" that "mistakes about the requirements of rationality are mistakes of rationality," (p. 253) he holds that if one's initial attitude toward p was the rational response to one's evidence, then

¹³ Cf. Williamson (2000). In the skepticism debate, externalism is ordinarily formulated as a thesis about knowledge or justification, not directly about what one *ought* to believe.

recognizing a peer's dissent provides no reason to change one's attitude. Hence: if you were antecedently a rational party to the disagreement, keep your belief; otherwise, give it up.

Third, White (2010) on irrelevant influences. Many of our beliefs are due to factors unrelated to their truth—say, where we were born. Yet irrelevantly influenced beliefs can still accord with one's first-order evidence. For White, what matters is not whether a belief is irrelevantly influenced, but whether that influence has led the belief away from one's evidence. Hence: if an irrelevantly influenced belief accords with your evidence, keep it; otherwise, give it up.

Fourth, Srinivasan (2020) on bad ideology. Ideological beliefs, she argues, are pervasively false and “have the function of sustaining, and are in turn sustained by, systems of social oppression” (p. 407). Such systems limit and distort our evidential access, often in ways undetectable from within. Still, some agents may possess reliable, implicit sensitivities—e.g. to racist or sexist cues—despite these distortions. Hence, a form of reliabilism: if a belief is produced by a reliable mechanism, keep it; if it's produced by an *unreliable* (e.g. ideologically infected) one, give it up.

All of these norms share the form of the externalist norm. As before, we have a distinction between the good case and the bad case. Like the externalist norm, each of these norms maintains that what you ought to do or believe depends on which case you're *really* in, not on whether you can *tell* which case you're in.

These are just four examples of externalist norms; there may be others. Indeed, we can construct any number of externalist norms in epistemology wherever a distinction between a good case and a bad case can be drawn. For example: if a piece of news is genuine

(good case), believe it; if it's fake (bad case), don't. If an expert is reliable (good case), trust them; if they're not (bad case), don't.

Before proceeding, a clarification. My use of "externalist" is stipulative and might not perfectly align with its standard use in debates about knowledge or justification. There, internalism is typically understood either as the view that justification supervenes on one's non-factive mental states, or as the view that justification is determined by what's accessible upon reflection; in either case, externalism is understood as the denial of the relevant claim.¹⁴ Some of the norms above, such as Lackey's and Srinivasan's, fit the traditional understandings of externalism well, since they make normative standing depend on features of one's environment, like the reliability of one's belief-forming methods. Others, such as Titelbaum's and White's, don't clearly do so.¹⁵

Still, all share an externalist spirit: in each case, what one ought to believe depends not on how things seem from the inside, but on how things actually are. Take Titelbaum. Whether one is the antecedently rational party to a disagreement is typically something one cannot tell: in either case, it will *seem* as if one is. Since Titelbaum's norm makes what one ought to believe depend on which case one is *really* in, it shares the spirit of externalism even if it does not strictly satisfy standard definitions.

¹⁴ Cf. Conee & Feldman (2001), Srinivasan (2020), and Johnson-King (2022) for discussion of the supervenience distinction; cf. Goldman (1999) for the access distinction. Access-internalism and supervenience-internalism are not necessarily incompatible.

¹⁵ Indeed, in his follow-up paper, Titelbaum (2019) explicitly rejects that his view is externalist. Though I cannot defend the point here, I find internalism an awkward fit for his view.

§3 The Joint Problems of Followability and Epistemic Risk

The attraction of externalism is clear: it honors the thought that being in the good case is itself epistemically good. But externalism faces two deep problems.

The first is the **problem of followability**. Externalism says: if you're in the good case, keep your beliefs; if you're in the bad case, give them up. If, as I've argued, we're often unable to tell which case we're in, then we're often unable to tell whether to keep or give up our beliefs. Therefore, the externalist norm is unfollowable.

Externalists sometimes respond that followability is normatively irrelevant. Drawing on Williamson's (2000) "anti-luminosity" argument, Srinivasan (2015) argues that no norm is perfectly followable, since there are no luminous conditions—that is, conditions such that one is always in a position to know whether they obtain. So if norms must be followable, then few—if any—norms survive.¹⁶ I submit that anti-luminosity is hard to resist. Yet followability comes in degrees: all else equal, a more followable norm is preferable to a less followable one.¹⁷ Furthermore, if followability were irrelevant, the externalist should simply adopt the truth norm ("believe p iff p ").¹⁸ But this is a move few are willing to make, precisely because we want epistemic norms to be guiding.

Ethicists and epistemologists sometimes distinguish primary from secondary norms, but this distinction offers no refuge for the externalist.¹⁹ If the externalist norm is meant to be primary, then, as just indicated, the externalist should simply adopt the truth norm as primary, since it's typically thought that primary norms need not be followable.²⁰ If, instead,

¹⁶ See also Hawthorne & Srinivasan (2013).

¹⁷ In another paper (m.s.), I argue that Srinivasan's conception of followability is too strong, and that a better conception does not require luminosity.

¹⁸ Though, epistemologists do often accept something like the truth norm as belief's *standard of correctness*.

¹⁹ Cf. McHugh (2011) on this distinction.

²⁰ Strictly speaking, the truth norm is an *instance* of the externalist norm, where the good case is a true belief.

the externalist norm is meant to be secondary (and thus followable), then its routine unfollowability makes externalism a poor candidate for that role, and we should seek an alternative. Therefore, the externalist norm is ill equipped for either role.

Much more could be said about the problem of followability.²¹ But I submit that this problem points to a second, deeper one: the **problem of epistemic risk**.²² Two facts are clear:

- (1) We often cannot tell whether we're in the bad case.
- (2) We know that many people are.

Together, these put us in positions of epistemic risk: for all we can tell, we ourselves might be in the bad case. The externalist response is to *ignore* this risk. Yet (previewing my argument in §5) ignoring risk seems to be an inappropriate response to risk: indeed, ignoring risk is *negligent*. Externalism, therefore, counsels epistemic negligence.

Worse, earnest attempts to follow the externalist norm *increase* our risk. Since we tend to assume we're in the good case, the de facto result of attempting to follow the externalist norm is that we'll almost always keep our beliefs, and thereby entrench ourselves more deeply in the bad case, if that's where we started.²³ To wit, trying to follow the externalist norm predictably deepens the risk of the very epistemic peril it purports to address.²⁴

²¹ Cf. Weatherston (2019) for a sophisticated response to the problem of unfollowability.

²² For more on epistemic risk, see Pace (2010), Riggs (2003), and Pettigrew (2022).

²³ Cf. Worsnip (2019; 2023) for a similar argument. See also Horowitz (2013) for the view that, to maintain one's beliefs, one must take oneself to be in the good case.

²⁴ Cf. Nguyen (2020b) for an analogous claim about "runaway" echo chambers.

For these reasons, the problem of epistemic risk can be understood as a sharpened extension of the unfollowability objection: it's precisely because externalism is unfollowable that attempts to follow it leave us epistemically worse off.

§4 Epistemic Insurance Policies

The lesson, then, is that we need epistemic norms that can both guide us and mitigate epistemic risk. The framework of *epistemic insurance policies* is designed to meet these twin demands.

Epistemic Insurance Policies: Norms for agents who are unable to tell whether they're in a good epistemic position, but which will insure them in case they're in a bad epistemic position.

Epistemic insurance policies are analogous in important respects to ordinary insurance policies—like health, home, or car insurance—and the analogy helps explain why we should, as it were, buy into them. Like ordinary insurance, they're designed to protect us when things go wrong. That said, the analogy is imperfect, and I discuss its limits below.

Epistemic insurance policies share two structural features with their practical counterparts. First, they aim to *improve* your epistemic position if you're in the bad case. Just as car insurance compensates you after a crash, or health insurance covers you after injury, an epistemic insurance policy “compensates” for the bad epistemic case—say, by counteracting the effects of an echo chamber or other distorting influence. In this way, they protect us against the epistemic bad case.

Second, epistemic insurance policies *worsen* your epistemic position, to some degree, if you're in the good case. They carry *costs* when things are going well: like any insurance premium, you pay without necessarily benefiting. Following an epistemic insurance policy may, when you're already well-situated, incur epistemic costs, but paying those costs in the good case is precisely what insures you against the bad case.

To see these two structural features in action, consider a real-life epistemic insurance policy (I consider more in §6):

CONCILIATE: Agents in a disagreement with a peer should conciliate to some degree.²⁵

CONCILIATE can be thought of as an epistemic insurance policy. Suppose you have a high credence in a true proposition p , and enter into a disagreement with someone who has a low credence in p . By conciliating, you lower your credence in a true proposition, thereby worsening your epistemic position with respect to p . This is the “buy-in” cost of conciliating. But now suppose, instead, you have a high credence in a *false* proposition p , and your interlocuter has a low credence in p . By conciliating, you lower your credence in a false proposition, thereby improving your epistemic position with respect to p . This is the “pay-out” benefit of conciliating.

These two structural features are the main respects in which epistemic insurance parallels ordinary insurance. Still, the analogy is imperfect, and drawing out the differences helps clarify the framework.

An insurance policy might help with the bad case in three distinct ways. First, it might help *prevent one from entering* the bad case—for instance, a policy on echo chambers might

²⁵ This is just a rough-and-ready version of conciliationism. Cf. Christensen (2007) and the essays in *The Epistemology of Disagreement: New Essays* for more on this debate.

reduce the likelihood of joining one. Second, it might help one *exit* the bad case—say, by providing heuristics for recognizing an echo chamber once inside it. Third, it might help by *compensating* for the damages caused by the bad case. Ordinary insurance works this third way: health insurance does not prevent or cure injury, but cushions its costs.

By contrast, epistemic insurance primarily functions in the first and second ways, but it can also mirror compensation in structure: it can make being in the bad case less bad without necessarily removing one from it. For instance, while one may be unable to escape an echo chamber, epistemic insurance might still mitigate its harms by priming one to not rely on the deliverances of one’s echo chamber in one’s reasoning.

Two further disanalogies are worth noting. First, ordinary insurance is necessarily social—its effectiveness depends on pooling costs and risks—whereas epistemic insurance can in principle operate individually. Second, in ordinary insurance, one’s risk profile is often salient (hence, a healthy person might rationally opt out), but in the epistemic case, it rarely is: since we cannot tell whether we’re in the good or the bad case, we cannot reliably gauge our own epistemic risk.

§5 The Argument

I previewed the argument for epistemic insurance in §3—viz. that externalist responses to risk are negligent.²⁶ The intuitive argument can be gleaned by analogy with regular insurance.

²⁶ Some distinguish negligence from recklessness, but I will treat them similarly and use “negligence” throughout. Cf. Shiffrin (2016) for discussion.

Consider this externalist policy: take out fire insurance on your house for the coming year only if your house will *actually* catch fire. That is not only *unhelpful* as advice (since it's unfollowable), but is intuitively the *wrong* advice. For even if your house will *not* catch fire, it's negligent not to insure it. No matter whether your house *will* catch fire, you should be averse to the risk.

The same reasoning applies to epistemic risk. Consider this externalist policy: take out epistemic insurance only if you *are* in the bad case. That's not only unhelpful as advice (for the reasons given in §3), but is intuitively the wrong advice. As far as you can tell, you might well be in the bad case, and failing to insure yourself against this epistemic risk is epistemically negligent.

There are, tacitly, three ingredients to this argument, to which I now turn: (a) the disvalue of the bad case, (b) its likelihood, and (c) the appropriate attitude toward risk. The presentation presumes a background in risk-weighted epistemic decision theory.²⁷ However, using epistemic decision theory doesn't commit me to epistemic consequentialism—the idea that whether we ought to believe p depends on whether believing p would be optimific.²⁸ As in ethics, we can understand epistemic decision theory as a “consequentialized” formalization of non-consequentialist views.²⁹

5.1 Epistemic Axiology: How Bad Is the Bad Case?

Taking out health insurance is (often) rational in part because the bad case is bad. Not only is breaking your femur extremely painful, but paying for surgery and physical therapy out of

²⁷ Cf. Greaves (2013).

²⁸ Cf. Berker (2013) for powerful objections to epistemic consequentialism.

²⁹ Cf. Drier (2011) on consequentializing and Sylvan (2020) on epistemic deontology.

pocket can be in the thousands. Though health insurance has a cost, the cost is worthwhile for removing the exorbitant out-of-pocket damage of paying for surgery. In that way, insurance makes the bad case less bad.

Just how bad is the *epistemic* bad case? Following James, we can distinguish between two goals of belief: believing truly and avoiding error.³⁰ These goals can come apart. Whenever one believes, one risks error; when one suspends, one avoids error but foregoes the possibility of believing truly.

How should we weigh these goals? Some epistemologists think there are no distinctly epistemic criteria by which to weigh these goals, and that, instead, their value is practical or moral.³¹

I disagree: I maintain that avoiding error is *epistemically* more important than believing truly—i.e. avoiding error is more epistemically disvaluable than believing truly is valuable. Three considerations support this claim.

First, this claim enjoys historical precedent. There is a long history of philosophers tacitly assuming this claim—for instance, Socrates’ preference for suspension of judgment over false belief and Descartes’ demand for indubitable foundations. Furthermore, infallibilism—the view that knowledge must be infallible—indicates a priority for error avoidance. While such views may be *too* demanding, the core intuition that error is especially epistemically bad remains.

Second, the relative marginal (dis)utilities of truth and falsity differ. Adding one more true belief to an already extensive stock of truths yields little additional epistemic value,

³⁰ Cf. James (1896). See also Pace (2011), Kelly (2013), and Johnson King & Babic (2020).

³¹ Cf. Dandelet (2024) for this view. See also Easwaran (2016) and Easwaran & Fitelson (2015) for more.

whereas adding one more false belief to a body of falsehoods considerably increases epistemic disvalue. By accumulating false beliefs, we become increasingly disconnected from reality, and since we reason from what we already believe, each falsehood makes further error more likely. One might object that analogous points apply to the “runaway good case,” where true beliefs generate further truths. Yet intuitively, a runaway bad case is far worse than a runaway good case is good. There is, as it were, something sadder about this sort of radical divorce from reality.³²

Third, the asymmetry between the values of truth and error helps explain why we should not believe propositions that are only marginally more probable than their alternatives.³³ Suppose a fair coin has a 0.5 probability of landing heads. Should you believe it will land heads? No. Yet if believing truly and avoiding error were equally valuable, you should be indifferent between believing and suspending, since each would realize the same expected epistemic value. Additionally, if the probability were 0.50000001, you would be required to believe that it will land heads. Intuitively, however, we should still refrain. The best explanation for why is that avoiding error carries greater epistemic weight than believing truly.³⁴

Yet even if we cannot say that one goal is more *epistemically* valuable than the other, moral and pragmatic considerations may affect their relative weight. It is often pragmatically and morally valuable to avoid error.³⁵ Consider passing judgment on a loved one’s character. By judging, you gain the chance at a true belief, but risk the far greater moral and practical

³² Cf. Wolf (2003) on insanity.

³³ This point connects to debates over whether belief is weak. On “weak belief” views, one can sometimes believe p with credence of less than 0.5, but not when p and not- p are equally likely. Cf. Hawthorne et al. (2016) and Holguín (2022).

³⁴ Cf. Dandelet (2024) for a discussion of this point, in support of a very different conclusion.

³⁵ Of course, suspending judgment be practically and morally bad, too; cf. Basu (2019); Ichikawa (2024); Crewe & Ichikawa (2021); Gordon-Smith (forthcoming).

cost of a false one, especially when the judgment is negative. The same holds, more pressingly, in political, moral, and economic domains: false beliefs in these areas can produce serious practical and ethical harms. This is particularly evident for high-stakes decision-makers, but it also applies to citizens in a democracy tasked with electing them.

One might object that false beliefs can be—and often are—good for us in ways that outweigh the badness of error. It may seem, for instance, that the benefits of some beliefs—community, purpose, belonging—outweigh their accuracy. Yet these goods, I submit, are often parasitic on truth. Many who leave dogmatic or insular communities later judge those experiences as disvaluable precisely because they rested on falsehoods. Even when some belief confers social or psychological benefit, we still care deeply that our communities are getting things right.

5.2 How Likely Is the Bad Case?

How likely is the bad case? In matters of health, the likelihood of misfortune depends on many factors: one's genetic dispositions, existing conditions, and lifestyle. A person with a chronic illness, or a competitive snowboarder, faces higher risk than a healthy homebody, all else equal.

Similarly, the likelihood that you are in the epistemic bad case depends on factors such as intelligence, the health of your informational environment, upbringing, education, and personality. Assigning a precise probability—especially in the first person—is difficult. But thinking yourself *above* the bad case is plainly hubristic. We know that we are imperfect cognitive agents operating in degraded epistemic environments; we know, too, that many others are in the bad case; and we know that we often cannot tell whether we are. Taken

together, these facts make it plausible that the probability of being in the bad case is significant—at least, significant enough to warrant attention.³⁶

5.3 Which Attitude Toward Risk Should We Have?

In her risk-weighted decision theory, Buchak (2013) argues that individuals differ not only in how they evaluate outcomes, but in how much they *care* about avoiding the worst ones. Some are risk-seeking, others risk-neutral, others risk-averse. Two agents might agree on the utility of the worst outcome (and its probability), yet the risk-averse one simply cares more about steering clear of it.

I submit that, at least in some contexts, we ought to be risk-averse—or at least not risk-seeking. While a person may permissibly take on high risks for themselves (as in skydiving), it seems wrong to do so when others depend on them. A parent, for instance, shouldn't walk their child across a rickety bridge even if the chance of collapse is low: they ought to care more about avoiding the bad case. At the very least, they should not be risk-seeking. Likewise for belief. When others rely on our epistemic performance, it may be impermissible to be risk-seeking with our beliefs.

A further reason for risk-aversion arises from the uncertainty of the probability itself. As noted above, we struggle to assign precise credences to how likely we are to be in the bad case. That very uncertainty is itself troubling. When we cannot quantify the risk, it is rational—and, indeed, epistemically responsible—to treat it as serious. Being risk-averse in

³⁶ The point remains even once we distinguish objective from subjective probabilities. The objective probability that people like us are in the bad case is fairly high, and we all know this to be the case. So even if we cannot know the exact numbers, we all have plenty of evidence that the objective probability is high, and it would be a rational mistake—and, indeed, hubristic—not to adjust one's subjective probabilities to this fact.

such cases amounts to acting as if the bad case were fairly likely. The literature on the precautionary principle captures this idea: when the probabilities (or disvalues) of bad outcomes are obscure, it is “better to be safe than sorry.”³⁷

5.4 *Completing the Argument*

We can now summarize the argument. First, avoiding error is often more valuable—epistemically, morally, and prudentially—than believing truly. Second, the likelihood of being in the bad case is significant. Third, we ought not be risk-seeking; especially when the probabilities are unclear, we should even be risk-averse. Taken together, these points indicate that it is usually rational, and often obligatory, to take out epistemic insurance: to adopt norms that safeguard us against the bad case even at some cost in the good case.

§6 Unifying Norms as Insurance Policies

Let me now demonstrate the power of the epistemic insurance framework.

Several epistemic norms already defended in the literature are natural candidates for being epistemic insurance policies. Framing them this way not only clarifies their rationale, but also helps *defuse common objections*. In particular, many epistemic norms face the objection of imposing costs. The insurance framework, however, *reframes* these costs as the premiums we pay to safeguard against epistemic risk.

³⁷ Cf. Rawls (1971) on the maximin rule and Sandin (2007) on the precautionary principle.

I have already showed how CONCILIATE can be framed as an insurance policy. CONCILIATE is a norm governing our *attitudes*. But other, more Zetetic insurance policies may concern, in the first instance, how one gathers and evaluates evidence:

DOUBLE-CHECK: Agents should double-check their work before believing.³⁸

An opponent might object that double-checking is worthwhile only if one *has* erred in evaluating one's evidence, and wasteful otherwise. An externalist counter to DOUBLE-CHECK yields: double-check if you have erred; otherwise, do not. But such advice is epistemically negligent: since we often cannot tell whether we have erred, the possibility of having erred is precisely the kind of epistemic risk we should guard against. Framed as an insurance policy, DOUBLE-CHECK gains plausibility. It functions as a safeguard against error; it is costly in the good case—it wastes time and thereby carries epistemic opportunity costs—but worthwhile all things considered.

So too for another evidence-gathering norm:

DIVERSIFY SOURCES: Agents should consume media from a diverse array of sources across the political spectrum.³⁹

This, too, incurs costs: if one already consumes only unbiased, reliable sources, diversification potentially exposes one to *worse* information, and adds time and cognitive effort. Yet refusing to diversify is epistemically negligent in light of our pervasive risk of misinformation and echo chambers. As an insurance policy, DIVERSIFY SOURCES mitigates these risks by (a)

³⁸ Cf. Woodard (2024). Her view is that double-checking *while* believing is permissible.

³⁹ Cf. Worsnip (2019). It's worth mentioning here that some decidedly non-externalist norms also constitute poor responses to epistemic risk, such as Rini's (2017) subjectivist norm for "partisan epistemology", to which Worsnip's norm is posed as an alternative.

reducing the likelihood of entering an echo chamber, (b) facilitating exit from one, and (c) cushioning the epistemic damage when one cannot escape entirely.

The framework of epistemic insurance provides a roadmap for creating and evaluating further norms beyond just these. Indeed, this framework also offers natural rivals to the other externalist norms discussed earlier. Against White’s norm on irrelevant influences: if you reasonably suspect that a belief is irrelevantly influenced (for example, “suspiciously convenient”), reduce confidence in it.⁴⁰ Against Srinivasan’s norm on ideology: if you reasonably suspect that a belief is ideological, reduce confidence—or, perhaps, ask “whose interests does this belief serve?” and withhold belief until you can justifiably answer.

These are just a small sample. Other extant norms that plausibly function as insurance policies include redundancy norms in testimony, higher-order evidence norms; complexity-reduction norms, calibration norms, humility and fallibility norms, norms of intellectual playfulness, and more. Importantly, appreciating the scope of epistemic insurance policies shows that such policies do not merely call for suspending judgment or tempering credences. Many operate across different dimensions of inquiry. Prioritizing error avoidance through these varied norms need not paralyze belief formation: indeed, by insulating ourselves against epistemic risk, we often place ourselves in a position where forming outright beliefs is, in fact, rational.

Finally, unifying these norms as insurance policies reveals their shared structure and rationale. Each imposes a cost in the good case, but that’s precisely the point. Their costs are not bugs, but features: one buys in even in the good case so that the norm pays out in the bad

⁴⁰ Cf. Worsnip (2023).

case. From this perspective, the question is not whether norms should have costs *at all*, but whether the costs of a given policy are worth paying, all things considered.

Final Remarks

Our norms should be tailored to fit our situation as non-ideal agents in non-ideal circumstances in particular times and places. If our culture's intellectual space lies on a spectrum between excessive confidence and excessive skepticism, then at present we lean too far toward the former. The remedy for this illness is a kind of intellectual caution: a deliberate recalibration toward humility, self-scrutiny, and open-mindedness. In this way, the call for epistemic insurance is not conservative, but, in Marx's words, a call for "the ruthless criticism of all that exists."

References

- Ahlstrom-Vij, K. (2013). Why We Cannot Rely on Ourselves for Epistemic Improvement. *Philosophical Issues*, 23(1), 276-296.
- Althusser, L. (1971). Ideology and Ideological State Apparatuses (Notes Towards an Investigation) (B. Brewster, Trans.). In *Lenin and Philosophy and Other Essays*. London: New Left Press. (English translation)
- Basu, R. (2019). What We Epistemically Owe to Each Other. *Philosophical Studies*, 176(4), 915-931.
- Berker, S. (2013). Epistemic Teleology and the Separateness of Propositions. *Philosophical Review*, 122(3), 337-393.
- Blake-Turner, C. (2020). Fake News, Relevant Alternatives, and the Degradation of Our Epistemic Environment. *Inquiry*, 1-21.
- Buchak, L. (2013). *Risk and Rationality*. Oxford, GB: Oxford University Press.

- Carter, J. A., & McKenna, R. (2020). Skepticism Motivated: On the Skeptical Import of Motivated Reasoning. *Canadian Journal of Philosophy*, 50(6), 702-718.
- Celikates, R., Haslanger, S., & Stanley, J. (Eds.) (forthcoming). *Analyzing Ideology*. Oxford University Press.
- Chomsky, N. & Herman, E. (1988). *Manufacturing Consent*. Pantheon Books.
- Christensen, D. (2007). Epistemology of Disagreement: The Good News. *Philosophical Review*, 116(2), 187-217.
- Christensen, D., & Lackey, J. (Eds.). (2013). *The Epistemology of Disagreement: New Essays*. Oxford University Press.
- Conee, E. & Feldman, R. (2001). Internalism Defended. *American Philosophical Quarterly*, 38(1), 1-18.
- Crewe, B. & Ichikawa, J. (2021). Rape Culture and Epistemology. In J. Lackey (ed.), *Applied Epistemology*. New York, NY: Oxford University Press.
- Dandelet, S. (2024). Epistemic Rationality and the Value of Truth. *Philosophical Review*, 133(4), 329-365.
- Dorst, K. (2023). Rational Polarization. *Philosophical Review*, 132(3), 355-458.
- Dreier, J. (2011). In Defense of Consequentializing. In M. Timmons (Ed.), *Oxford Studies in Normative Ethics, Volume 1*. Oxford University Press.
- Easwaran, K. (2016). Dr. Truthlove or: How I Learned to Stop Worrying and Love Bayesian Probabilities. *Noûs*, 50(4), 816-853.
- Easwaran, K. & Fitelson, B. (2015). Accuracy, Coherence, and Evidence. *Oxford Studies in Epistemology*, 5, 61-96.
- Geuss, R. (1981). *The Idea of a Critical Theory: Habermas and the Frankfurt School*. Cambridge: Cambridge University Press.
- Goldman, A. (1999). Internalism Exposed. *Journal of Philosophy*, 96(6), 271-293.
- Gordon-Smith, E. (forthcoming). The Ethics of Disbelief. *Australasian Journal of Philosophy*.
- Gramsci, A. (1971). *Selections from the Prison Notebooks* (Q. Hoare & G. Nowell Smith, Eds. & Trans.). International Publishers.
- Greaves, H. (2013). Epistemic Decision Theory. *Mind*, 122(488), 915-952.
- Haslanger, S. (2017). Racism, Ideology, and Social Movements. *Res Philosophica*, 94(1), 1-22.

- Hawthorne, J., Rothschild, D., & Spectre, L. (2016). Belief is Weak. *Philosophical Studies*, 173(5), 1393-1404.
- Hawthorne, J., & Srinivasan, A. (2013). Disagreement Without Transparency: Some Bleak Thoughts. In D. P. Christensen & J. Lackey (Eds.), *The Epistemology of Disagreement: New Essays* (pp. 9-30). Oxford: Oxford University Press.
- Holguín, B. (2022). Thinking, Guessing, and Believing. *Philosophers' Imprint*, 22(1), 1-34.
- Horowitz, S. (2013). Epistemic Akrasia. *Noûs*, 48(4), 718-744.
- Ichikawa, J. (2024). *Epistemic Courage*. Oxford: Oxford University Press.
- Intemann, K. (2010). Years of Feminist Empiricism and Standpoint Theory: Where Are We Now? *Hypatia*, 25(4), 778-796.
- James, W. (1896). "The Will to Believe." In *The Will to Believe and Other Essays in Popular Philosophy*. New York: Longmans, Green, and Co.
- Johnson King, Z. (2022). Radical Internalism. *Philosophical Issues*, 32(1), 46-64.
- Kelly, T. (2008). Disagreement, Dogmatism, and Belief Polarization. *Journal of Philosophy*, 105(10), 611-633.
- Kelly, T. (2013). Evidence Can Be Permissive. In M. Steup & J. Turri (Eds.), *Contemporary Debates in Epistemology*.
- Kunda, Z. (1990). The Case for Motivated Reasoning. *Psychological Bulletin*, 108, 480-498.
- Lackey, J. (2021). Echo Chambers, Fake News, and Social Epistemology. In S. Bernecker, A. K. Flowerree, & T. Grundmann (Eds.), *The Epistemology of Fake News*. New York, NY: Oxford University Press.
- McHugh, C. (2011). What Do We Aim At When We Believe? *Dialectica*, 65(3), 369-392.
- Nguyen, C. T. (2020a). Echo Chambers and Epistemic Bubbles. *Episteme*, 17(2), 141-161.
- Nguyen, C. T. (2020b). Cognitive Islands and Runaway Echo Chambers: Problems for Epistemic Dependence on Experts. *Synthese*, 197(7), 2803-2821.
- Nozick, R. (1981). *Philosophical Explanations*. Cambridge, MA: Harvard University Press.
- Pace, M. (2010). The Epistemic Value of Moral Considerations: Justification, Moral Encroachment, and James' 'Will To Believe'. *Noûs*, 45(2), 239-268.
- Pettigrew, R. (2022). *Epistemic Risk and the Demands of Rationality*. Oxford, UK: Oxford University Press.
- Pronin, E., Gilovich, T., & Ross, L. (2004). Objectivity in the Eye of the Beholder: Divergent Perceptions of Bias in Self Versus Others. *Psychological Review*, 111(3), 781-799.

- Pronin, E., Lin, D. Y., & Ross, L. (2002). The Bias Blind Spot: Perceptions of Bias in Self Versus Others. *Personality and Social Psychology Bulletin*, 28(3), 369-381.
- Rawls, J. (1971). *A Theory of Justice*. Cambridge, MA: Harvard University Press.
- Riggs, W. D. (2003). Balancing Our Epistemic Goals. *Noûs*, 37(2), 342–352.
- Rini, R. (2017). Fake News and Partisan Epistemology. *Kennedy Institute of Ethics Journal*, 27(S2), 43-64.
- Sandin, P. (2007). Common-Sense Precaution and Varieties of the Precautionary Principle. In T. Lewens (Ed.), *Risk: Philosophical Perspectives* (pp. 99-112). London; New York.
- Shelby, T. (2003). Ideology, Racism, and Critical Social Theory. *Philosophical Forum*, 34, 153-188.
- Shiffrin, S. (2016). The Moral Neglect of Negligence. In D. Sobel, P. Vallentyne, & S. Wall (Eds.), *Oxford Studies in Political Philosophy*. Oxford University Press.
- Srinivasan, A. (2015). Normativity Without Cartesian Privilege. *Philosophical Issues*, 25(1), 273-299.
- Sylvan, K. L. (2020). An Epistemic Non-Consequentialism. *The Philosophical Review*, 129(1), 1-51.
- Titelbaum, M. (2015). Rationality's Fixed Point. *Oxford Studies in Epistemology*, 5.
- Titelbaum, M. (2019). Return to Reason. In M. Skipper & A. Steglich-Petersen (Eds.), *Higher-Order Evidence: New Essays*. Oxford University Press.
- Weatherson, B. (2019). *Normative Externalism*. Oxford, UK: Oxford University Press.
- McWilliams, E. (2021). Evidentialism and Belief Polarization. *Synthese*, 198(8), 7165-7196.
- White, R. (2010). You just believe that because.... *Philosophical Perspectives*, 24(1), 573-615.
- Williamson, T. (2000). *Knowledge and Its Limits*. New York: Oxford University Press.
- Wilson, T. D. (2002). *Strangers to Ourselves: Discovering the Adaptive Unconscious*. Harvard University Press.
- Wolf, S. (2003). Sanity and the Metaphysics of Responsibility. In G. Watson (Ed.), *Free Will* (2nd ed). Oxford: Oxford University Press.
- Woodard, Elise (2024). Why Double-Check? *Episteme*, 21(2), 644-667.
- Worsnip, A. (2019). The Obligation to Diversify One's Sources: Against Epistemic Partisanship in the Consumption of News Media. In J. Saunders & C. Fox (Eds.), *Media Ethics, Free Speech, and the Requirements of Democracy* (pp. 240-264). Routledge.

Worsnip, A. (2023). Suspiciously Convenient Beliefs and the Pathologies of
(Epistemological) Ideal Theory. *Midwest Studies in Philosophy*, 47, 237-268.