

CHAPTER 3: THE CASE FOR EPISTEMIC OPTIONS

Elizabeth Jackson, Saint Louis University

ABSTRACT: This chapter defends the existence of epistemic options: cases in which a person can rationally adopt more than one belief-attitude toward a proposition, given their evidence. While most philosophers reject such options, I argue that epistemic options are both real and relatively common, and that they often extend beyond bi-optionality to include tri-optionality (belief, disbelief, and withholding). After briefly reviewing existing arguments for epistemic options and their limitations, I develop my positive case in two stages. First, I present a range of intuitive examples across domains—future predictions, past events, theoretical inquiry, and the interpretation of others and oneself—where the evidence fails to rationally determine a single attitude. Second, I provide two theoretical arguments for epistemic options: from the complexity and heterogeneity of evidence, and from the idea that common patterns of reasoning, particularly inference to the best explanation, do not always select a unique best explanation, but instead often leave multiple, equally good explanatory options. In such cases, it can be rational to adopt any of the competing attitudes. Together, these arguments support a permissivist view on which epistemic options are not rare anomalies but a pervasive feature of our lives.

KEYWORDS: epistemic permissivism; epistemic options; epistemic rationality; evidence; evidential balance; evidential weight; inference to the best explanation

Note to MEW readers: *I am in the early stages of working on a book on epistemic permissivism—this is a chapter in the book. If you're curious, the table of contents tentatively looks like this:*

Chapter 1: Introduction

Chapter 2: Epistemic Ties and Epistemic Options

Chapter 3 (this chapter): The Case for Epistemic Options

Chapter 4: Objections to Permissivism

Chapter 5: Voluntary Belief

Chapter 6: Encroachment

Chapter 7: Permissivism about Religious Belief

Chapter 8: Faith, Voluntary and Rational

Chapter 9: The Problem of Evil

Chapter 10: Pascal's Wager

Chapter 11: Conclusion

This chapter is a work in progress, so I appreciate your willingness to read through ideas that are still somewhat half baked! I would like to expand and improve especially the arguments in sections 4 and 5, so suggestions there are most welcome. I am also wondering if section 3 is too long or repetitive, so feedback on that would be useful as well. I know there are a lot of objections I don't address here, and I plan to devote the next chapter to objections, including arbitrariness, accuracy, immodesty, and the idea that withholding is the rational attitude in the cases in question. Happy to discuss any of those at MEW, of course, but wanted to note that I do plan to address many common objections later in the book even though they aren't mentioned here. Thank you again for taking this time to read this and for your feedback!

1. Introduction

You send a message to a close friend—something that matters—and two days pass without a reply. Part of you starts to think they're upset with you; after all, your last conversation was slightly tense, and the silence feels pointed. But there are other explanations. They've been busy before, and you

know they sometimes withdraw when work gets overwhelming. Then again, you notice they've been active elsewhere—liking posts, responding to others—which makes the silence harder to dismiss. As you reflect on the situation, you can feel yourself pulled in different directions. You could believe that they're intentionally avoiding you. You could believe that there's an innocent explanation. Or you could withhold judgment altogether, leaving the matter unsettled in your mind.

Intuitively, your evidence leaves open more than one rational response. Cases like this are not rare or peripheral—they arise both throughout theoretical inquiry and ordinary life. My goal in this chapter is to convince you that, in such cases, you have *epistemic options*. Recall:

S has **epistemic options** =_{df} at a time, S can rationally adopt more than one belief-attitude to a proposition, given her evidence.

This goal is divided into more modest and more ambitious claims. The more modest version is that there are at least some cases of **bi-optionality** (involving exactly two attitudes, one of which is withholding). More ambitiously, I hope to convince the reader that (i) there are also cases of **tri-optionality** (involving all three attitudes), and (ii) such cases are relatively common.

A defeasible reason to posit epistemic options comes from the analogy between ethics and epistemology. These are two central normative domains: ethics governs our actions, and epistemology governs our beliefs. Because inquiry in these domains often proceeds independently, we can sometimes make progress in one area by drawing on insights from the other. While we should not force symmetry—belief and action differ in important ways—it is reasonable to begin with a presumption of similarity unless we have good reason to reject it. And it is uncontroversial that options arise in ethics: it is permissible to eat eggs or cereal for breakfast, to mow the lawn or take a nap, to study medicine or study law. Given that ethical options are both widely accepted and commonplace, we have defeasible reason to expect that epistemic options may arise as well.

Despite this, the idea of epistemic options has been widely rejected. Impermissivists deny them outright, but even many permissivists reject them, allowing only interpersonal or diachronic permissivism (as discussed in the previous chapter). I contend, however, that there's a strong case for the existence of epistemic options. Given the defeasible symmetry between epistemology and ethics, we have good reason to take this possibility seriously, despite its relative unpopularity.

This chapter develops that case. **Section 2** reviews existing arguments for epistemic options and explains their limitations. **Section 3** presents a range of intuitive examples, both to give a sense of the cases of interest, and to give an “argument from cases” so to speak. **Sections 4 and 5** then offer two positive arguments, both of which support the claim that epistemic options are relatively common and often tri-optional. The first appeals to the limitations of evidence: I argue that evidence is often limited, complex, and difficult to integrate, and so does not always rationally determine a single attitude. The second draws on inference to the best explanation, arguing that there is not always a unique best explanation of the evidence, but sometimes several equally good ones. **Section 6** concludes. The next chapter considers and responds to objections to these arguments.

2. Existing Arguments for Epistemic Options

This section briefly reviews existing arguments that motivate epistemic options. Section 2.1 discusses Matthew Kopec’s self-justifying belief argument, Section 2.2 discusses Blake Roeber’s “evidence trickle” argument, and Section 2.3 discusses Richard Pettigrew’s epistemic risk argument. While part of my aim in discussing these arguments is to illustrate the myriad of ways epistemic options can be motivated, my primary goal is to show how my view—and my defense of it—differs from these existing approaches. As we will see, each of these arguments is limited in important ways: they tend to apply only to relatively rare cases and/or support at only bi-optionality, rather than tri-optionality.

2.1 Kopec on Self-Justifying Beliefs

Some beliefs are such that, when formed, they generate evidence for themselves. In cases where such beliefs produce sufficient evidence, they can be *self-justifying*: upon forming the belief, one thereby acquires adequate evidence for it, and so the belief is justified. While the existence of such beliefs is controversial (see Antill 2019; Marxen 2020), there is some tentative empirical support for them (Milkman 2021: ch. 6). Several authors have argued that self-justifying beliefs support permissivism (Kopec 2015; Raleigh 2017; Drake 2017). Consider the following case (modified slightly from Kopec & Titelbaum 2016: 197):

Rain: God makes the following proposal: If you believe that it will rain in LA tomorrow, then God will make sure it rains in LA tomorrow. But if you disbelieve it will rain in LA tomorrow, then God will make sure it doesn’t.

Intuitively, in this case, if you believe it will rain in LA tomorrow, you thereby have sufficient evidence to justify that belief, and it is rational. The same goes for disbelief. And perhaps if you withhold belief, there remains enough uncertainty—about the forecast and/or about God’s actions—that withholding is also rational.

Rain is, of course, a philosopher’s case. There are (arguably) more realistic examples: believing that a date will go well might calm your nerves and help bring that about; similarly for believing you will win a basketball game, give a good talk, or recover from an illness. Such cases are interesting and deserve further exploration, both philosophically and empirically. They also represent one of the few existing attempts to defend tri-optionality in the literature.

However, such cases are quite limited. Beliefs that generate evidence for themselves are already rare, and cases in which they generate *sufficient* evidence for rationality are rarer still. As a result, even if self-justifying beliefs provide a technical counterexample to impermissivism, they support epistemic options only in a narrow range of cases (for an impermissivist response, see Greco & Hedden 2016: 371).

2.2 Roeber on the Trickle of Evidence

Evidence can come in small quantities. As a result, our evidence for a proposition can gradually improve or deteriorate over time. This observation forms the basis of another argument for epistemic options, developed by Roeber (2019, 2020), which I will call the evidence trickle argument. Consider the following case (modified slightly from Stapleford & Jackson 2026: 158):

Sign: You are walking toward a sign. At the beginning, you cannot make out what it says at all, and so you should withhold belief. At the end, you can clearly read that the sign says “Exit 31.” Somewhere in between, your evidence equally supports believing that the sign says “Exit 31” and withholding belief.

Or similarly, when watching a sporting event, you may begin by withholding belief about which team will win. As one team gradually pulls ahead, your evidence shifts until, at some point, it clearly supports the belief that they’ll win. Roeber’s idea is that, in cases like these, there is an intermediate stage at which the evidence supports belief and withholding equally well. In such “equipollent cases,” you have the option to either believe or withhold.

Roeber’s cases are intuitive and illuminating. However, they are also limited in important ways. First, the evidence trickle argument supports bi-optionality—between belief (or disbelief) and withholding—but does not generate genuinely tri-optional cases. Second, cases in which the evidence is perfectly balanced between belief-attitudes are likely to be rare. Even in the examples above, the evidential situation quickly resolves as one moves closer to the sign or as the game progresses. Thus, while such cases may again provide a technical counterexample to impermissivism, they do not support the existence of widespread epistemic options.

2.3 Pettigrew on Epistemic Risk

Some of us like to take risks; others prefer the safer option. Sometimes, different attitudes toward risk can be simultaneously rational—e.g. I want to try the brand new restaurant, whereas you want to go to our regular spot. Pettigrew (2016, 2022), following William James, applies this insight to epistemology. Belief and disbelief are epistemically risky, as each offers the possibility of getting things right but also getting things badly wrong. Withholding, in contrast, avoids both extremes: it cannot err but does not contribute to accuracy.

Working within an epistemic utility framework, Pettigrew’s basic idea is this: rational belief is determined by decision-theoretic principles that encode, among other things, different attitudes toward epistemic risk. More risk-averse principles tend to favor withholding in uncertain cases, while more risk-tolerant principles tend to favor belief (or disbelief). On some such principles, both believing and withholding belief are rationally permissible responses to the same evidence, depending on how one weighs the competing epistemic goals of truth and error-avoidance.

It’s not clear that Pettigrew explicitly endorses epistemic options regarding *belief* (i.e. synchronic, intrapersonal, belief permissivism)—his 2016 paper endorses belief permissivism, but it’s not clear whether he endorses intrapersonal permissivism (his discussion of risk-averse and risk-seeking agents on page 257 suggests he’s focused on the interpersonal case). His 2022 book specifically defends permissivism about credence. Nonetheless, his account could be extended to support epistemic options insofar as, for a single agent with fixed evidence, there are different rationally admissible ways of balancing epistemic risk that yield different permissible attitudes.

Quanbeck and Worsnip (2025: 8–9) more explicitly appeal to epistemic risk to defend epistemic options. They argue that, in some cases, we are not merely choosing between belief-attitudes given a fixed risk profile; rather, sometimes, we are “permitted to either {have a risk-seeking attitude, believe} or {have a risk-averse attitude, suspend}” (p. 9), thereby choosing *both* a risk-attitude and a first-order

attitude at once. Thus, while the epistemic risk framework naturally supports interpersonal permissivism, it also supports intrapersonal permissivism and even synchronic epistemic options.

Such arguments, however, will not yield tri-optionality. The central contrast in such frameworks is between taking on epistemic risk (by believing or disbelieving) and avoiding it (by withholding). As a result, the epistemic options are either (withhold, believe) or (withhold, disbelieve). The framework does not, on its own, yield cases in which all three attitudes are simultaneously permissible. Put another way, because belief and disbelief play structurally similar roles as risk-taking attitudes, changes in one's risk-attitudes cannot permit, at once, belief that p and disbelief that p .

There's a further question of when—i.e. across what range of cases—it is permissible to alter one's risk attitudes. On Pettigrew's framework, permissivism arises only when different attitudes toward epistemic risk are themselves rationally admissible. But, in the intrapersonal case, it is unclear when these risk-attitudes should be treated as fixed background features, and when they are elements that are themselves subject to rational choice in a way that generates epistemic options. If they are largely fixed, then the framework will yield relatively few epistemic options, since a given risk profile will typically determine a unique rational attitude. If, by contrast, they are revisable, then epistemic options may be more widespread—but this requires an account of when such revisions are rationally allowed. (Note that Quanback and Worsnip 2025: 9 emphasize only needing a single case for permissivism to be true, which suggests that they are not necessarily committed to widespread epistemic options.) For our purposes, then, the limitation is twofold: the view does not yield tri-optionality, and it leaves unclear whether the epistemic options it generates are rare or widespread.

I now turn to my own case for epistemic options. I'll motivate that epistemic options are more commonplace and motivate tri-optionality—not just bi-optionality.

3. Cases of Epistemic Options

This section provides categorized, concrete cases of epistemic options. I'll focus on cases of tri-optionality, but I'll also note how slight tweaks on these cases motivate bi-optionality. The cases I'll focus on fall roughly into five categories: future predictions, past events, big theoretical questions, interpreting others, and interpreting oneself. Of course, the point isn't that any proposition that falls into one of these categories is automatically a case of epistemic options—it will depend on the thinker and their evidence—but these categories tend to be good places to look for cases of epistemic options.

There's some overlap between these categories; for example, some big questions (e.g. in science) involve future forecasting. Interpreting the behavior of others (e.g. facilitating a disagreement between your kids) might involve deciphering past events. I don't necessarily take these categories to be exhaustive, either. Nonetheless, part of the goal of this section is to begin to answer Forrester's question (2024) about where permissive cases arise and how to identify such cases. A brief preview of my answer is that we should look for epistemic options in cases where the evidence is complex and qualitatively diverse—and we'll see these cases shares this feature. Finally, I acknowledge that the reader might not share my intuitions about *all* the following cases or categories. Even if not, I hope that the reader finds at least *some* of the cases compelling. My theoretical arguments in sections 4 and 5 also further motivate that epistemic options exist and are relatively common.

3.1 Future Predictions

The first category of cases of epistemic options involves future predictions. Consider the following cases:

Startup Investment: You are deciding whether to invest a large amount of money into a startup. The upside case is compelling: rapid growth, a strong founding team, and a product that appears to solve a real problem. Several experienced investors have already committed, which gives you reason to believe the company will succeed. But there are warning signs. The market is crowded, the company is not yet profitable, and some of the growth metrics may be inflated or unsustainable. Trusted advisors disagree—one calls it a rare opportunity, another a classic hype-driven failure. As you assess the evidence, you can see a strong case for success and a strong case for failure, while also recognizing that the evidence is difficult to weigh overall. It seems rational for you to believe the company will succeed, to believe it will fail, or to withhold belief.

Quarterback: You are part of a front office deciding whether to draft a quarterback in the first round. The analytics team is strongly in favor: his efficiency metrics are elite, and historically, players with his profile succeed at a high rate. But the film complicates things. Much of his success comes from clean pockets behind an exceptional offensive line, and when plays break down, he looks hesitant. The intangible evidence is similarly mixed. His coaches praise his leadership and discipline, while scouts disagree about his demeanor—some see confidence, others overconfidence. As you weigh these different kinds of evidence, you can see a compelling case that he will succeed, and a compelling case that he will fail. At the same time, the evidence is difficult to integrate, and it feels equally reasonable to withhold belief. In this situation, it seems rational for you to believe he will be a franchise quarterback, to believe he will not, or to withhold judgment altogether.

This structure arises in a wide range of forward-looking contexts: investment decisions, predictions about election outcomes, climate projections, medical diagnosis and treatment, and long-term technological forecasting (e.g. the future of AI). These cases often involve heterogeneous evidence—such as statistical models, expert testimony, background assumptions, narrative projections—that do not combine in a straightforward way. Especially in complex contexts and further off predictions, the evidence can support competing predictions without clearly settling which is correct.

Slight variations on these cases can also motivate **bi-optionality**. To see this, we strengthen the evidence in one direction while preserving its qualitative diversity. For instance, in the startup case, suppose that in addition to the existing positive indicators, the company has just secured a major contract with a well-established partner, and early revenue numbers exceed expectations. These developments strengthen the case for success. At the same time, the earlier concerns remain: the market is still crowded, profitability is uncertain, and some of the growth may be unsustainable. In such a case, it may no longer seem rational to *disbelieve* that the company will succeed, given the weight of the positive evidence. However, it can still be rational either to believe that the company will succeed or to withhold judgment, given the remaining uncertainties and the difficulty of integrating the evidence. In this way, the same general structure—complex, heterogeneous evidence—can yield bi-optionality.

3.2 Past Events

We also often have epistemic options when trying to determine what happened in the past. This may concern events in the distant past (e.g., evolutionary mechanisms or cosmology) or the more recent past (e.g., unsolved crimes or mysteries, which are commonly discussed in the context of permissivism; see Rosen 2001). Again, the key feature is the presence of complex and qualitatively diverse evidence. Consider two cases:

True Crime: You are serving on a jury in a murder trial. An eyewitness places the defendant at the scene, and his testimony aligns in part with the timeline, giving you reason to believe he is guilty. But doubts emerge: his story has shifted across interviews, and there is evidence of investigative pressure. The physical evidence is similarly ambiguous—the defendant’s DNA is present, but explainable given a prior relationship—and an alternative suspect is introduced with a plausible motive. As you deliberate, you find that both guilt and innocence organize the evidence in a coherent way. Given the complexity of the evidence, withholding belief also seems reasonable. It seems rational for you, then, to believe the defendant is guilty, to believe he is not, or to withhold judgment.

Historical Interpretation: You are reading about a controversial political decision made decades ago. Some historians argue that the leader acted out of genuine concern for national security, citing internal memos and consistent patterns in prior decisions. Others argue that the decision was self-interested, pointing to suspicious timing, selective use of intelligence, and the political advantages that followed. The sources themselves are mixed—some are firsthand but biased, others more neutral but incomplete. As you consider the evidence, each interpretation organizes the available information in a coherent and plausible way. It seems rational for you to believe the decision was justified, to believe it was self-serving, or to withhold judgment.

Comparable cases arise in investigative journalism, unsolved mysteries, accident reconstruction, and even everyday disputes about what happened (e.g. which child started a fight, whether one player fouled another). In these contexts, the evidence is often incomplete, filtered through testimony, and structured by competing narratives. Different explanations can fit the same body of evidence, making it difficult to determine which, if any, is uniquely rational to accept. And again, similar cases can support bi-optionality when some significant, but not decisive, evidence is added to one side (e.g. adding a second, corroborating eyewitness), but the qualitative diversity of evidence is preserved. (All of the later cases in this section can be similarly modified to support bi-optionality, but I will stop mentioning this to avoid redundancy.)

3.3 Big, Theoretical Questions

We also often find epistemic options when it comes to “the big questions,” e.g. questions in philosophy, including ethics, epistemology, metaphysics, and aesthetics, but also questions in other disciplines, especially ones at the “edge” of a research program, including politics, economics, law, and science. These questions are often highly controversial, even among academics, and many (especially in philosophy) have been debated as far back as our written records go. Consider:

God: You're reconsidering your views about the existence of God. On the one hand, you are drawn to a picture of the world as dependent on another being, and you find some of the classical arguments for God's existence compelling. Testimony from people you trust about religious experience adds further weight. On the other hand, the problem of evil presses strongly against theism, and naturalistic explanations of the world are satisfying and simple. As you reflect, you notice that you can "see" the world in two different ways: as self-sufficient or as dependent on something beyond it. Neither perspective feels forced on you by the evidence. Instead, both remain live possibilities, and it seems rationally open to you to believe that God exists, to believe that God does not exist, or to withhold judgment (inspired by van Inwagen 1994; see chapter 7 for a chapter-length defense of the claim that theistic beliefs are often ones where epistemic options occur.).

Quantum Mechanics: You are interested in how to interpret the results of quantum mechanics. On the one hand, some interpretations posit a single, determinate reality in which quantum states collapse upon measurement. You find the idea of a single, determinate universe intuitive, but the idea that quantum states collapse upon measurement is puzzling. On the other hand, other interpretations deny collapse and instead posit a branching universe in which all possible outcomes occur. These views offer a more unified account of the theory, but at the cost of a highly counterintuitive ontology. As you reflect, you find that the empirical data underdetermines which interpretation is correct, and that each framework has its own theoretical virtues and costs. It seems epistemically open to you to believe that reality is determinate, to believe that it is fundamentally branching, or to withhold belief about the correct interpretation.

Many philosophical issues fall in this category, such as the nature of consciousness, personal identity, the status of moral claims, the permissibility of eating non-human animals, and the like. This can also include competing interpretations of scientific theories (e.g., climate models or evolutionary explanations). This category also involves legal, economic, and psychological controversies, where the evidence is complex and multifaceted. Legal theorists may disagree about whether a constitutional provision protects a particular right or how precedent applies in hard cases. Economists may dispute whether inflation is driven by monetary policy or supply shocks, whether raising the minimum wage increases unemployment, or how to explain a recession. Psychologists debate how to interpret conflicting replication results, or whether observed effects—e.g. the success of a therapy—reflect genuine causal mechanisms or instead arise from placebo effects or are mere correlations. While such disagreement is often taken to support interpersonal permissivism, these cases also support intrapersonal permissivism, as a single researcher may be pulled in different directions when faced with inconclusive evidence. Some of these cases can also spill over into the previous categories—for instance, when framed as questions involving the past or future—but they are unified here in being theoretical or "big" questions.

3.4 Interpreting Others

A central part of living in community is interpreting the beliefs, desires, and intentions of others. But these attitudes are often difficult to discern—we lack direct access to other minds, and people do not always fully reveal themselves. In such cases, we are often faced with epistemic options:

Dishwasher: After a tense morning, you come home and notice that your spouse has not done something they know matters to you: unloading the dishwasher. One possibility is that this was intentional, a small act of retaliation in light of the earlier disagreement. But there is a familiar alternative: they have forgotten this task before, even in neutral circumstances, and may simply have been distracted. The surrounding details—their tone earlier, how they left the house—can be read either way. As you reflect, you find yourself pulled in three directions. You could believe the omission was intentional, believe it was accidental, or withhold belief altogether. Given the ambiguity of the evidence, each of these attitudes seems rationally available.

Workplace Interaction: During a meeting, a colleague responds briefly to one of your ideas and quickly moves on. One possibility is that they are dismissive—that they do not take your contribution seriously. Another is that they are simply pressed for time or focused on keeping the meeting on track. You recall that they have been supportive in the past, but also that they can be curt under pressure. Their tone, body language, and the broader context admit of multiple readings. As you reflect, it seems epistemically open to you to believe they were dismissive, to believe they were not, or to suspend judgment about their intentions.

This pattern appears throughout ordinary social life: interpreting tone, intention, sincerity, and emotion. Because the evidence often consists of subtle behavioral cues, background knowledge, and only partial access to others' mental states, it is frequently ambiguous and open to competing interpretations. Similar dynamics arise in cases like the unanswered message from the introduction, perceived slights, or unclear social signals more generally. In such cases, multiple belief-attitudes can remain rationally available.

3.5 Interpreting Oneself

Even though we have privileged access to our own mental states, interpreting our own behavior and attitudes is not always straightforward. While there may be additional barriers to interpreting others, our own beliefs, desires, motives, and intentions are not always introspectively obvious. Consider the following cases:

New Job: You begin to reflect on your own motivations for a recent decision—say, taking a new job. On one interpretation, you chose it for the right reasons: it aligns with your values, offers meaningful work, and reflects careful judgment. But another interpretation is also plausible: you were drawn by status, external pressure, or a desire to avoid risk. The evidence available to you—your past reasoning, your emotional responses, the explanations you gave to others—is partial and difficult to interpret. Some considerations support the more charitable interpretation, others the more critical one, and many can be read either way. As you reflect, it seems open to you to believe that your motivations were genuine, to believe that they were not, or to suspend judgment about your own reasons.

Happiness: You begin to wonder whether you are genuinely happy in your current life. On the one hand, there is evidence that you are: your life is stable, you have meaningful relationships, and you often report satisfaction when asked. On the other hand, you notice recurring dissatisfaction, a sense of restlessness, and moments where

your engagement feels forced. You also consider the possibility that you have adapted to your circumstances or are rationalizing them. The evidence—your feelings, behaviors, and reflections—is fragmented and difficult to interpret. It seems epistemically open to you to believe that you are genuinely happy, to believe that you are not, or to suspend judgment.

Other introspective cases include assessing whether one sincerely holds a belief, whether one's motivations are self-interested or principled, and whether one's actions reflect one's values. These cases are distinctive in that the evidence is first personal but nonetheless opaque and can be difficult to evaluate. Even with privileged access, it can be unclear how to interpret one's own mental states. Thus, even first-personal evidence can render epistemic options.

We've now seen five categories of cases where epistemic options tend to arise. Now, to further motivate the existence and ubiquity of epistemic options, I provide two theoretical arguments: from the limitations of evidence and from "good enough" explanations.¹

4. Evidence: Limited, Complicated, Heterogeneous

Sometimes, a body of evidence straightforwardly supports a proposition (or its negation). Most people's evidence supports that the earth is round, $1+1=2$, and Paris is the capital of France. Generally, when our evidence supports a proposition, we should believe it (or disbelieve its negation) (although, since our evidence supports an infinite number of propositions, this principle should probably be constrained in certain ways). In other cases, the evidence for p and for not- p is tied (or nearly tied)—for example, the evidence that a fair coin landed heads or that a Democrat will win the next US election. Similar cases include if one expert says p and the other says not- p when they are roughly equally reliable, or a case where you have no evidence at all regarding p . In tie cases, the evidence for p and for not- p is parallel and symmetrical, and the probability of p on your evidence is normally precise and right around 0.5. In near-tie cases, the evidence slightly supports a proposition over its negation, such as tossing a coin very slightly weighted toward heads, a weak, secondhand rumor, or a quick glance at a faraway object in the distance. In cases of both ties and near ties, you should withhold belief.

All of these are cases where a body of evidence "points" in a particular direction, and thus rationalizes a corresponding attitude.² However, contra White (2005: 447; 2014: 314), not all evidence does this. Some bodies of evidence do not "point" in a single direction and thus do not rationalize a single corresponding attitude. Or so I will argue.

Before we proceed, note that my argument here does not rely on the idea that permissivism is true because rational thinkers can be mistaken about what their evidence supports. While this is a potential route to permissivism (see Kopec and Titelbaum 2016: 190), this is not the route I take here. Rather, rational thinkers may always be fully aware of what their evidence supports, yet still find themselves with epistemic options.

¹ These arguments draw on my arguments for permissivism in Stapleford & Jackson (2026) but are updated and expanded.

² However, even if every body of evidence always "points" in a particular direction—i.e. for every body of evidence and every proposition, there's a precise epistemic probability—that doesn't get us impermissivism. For example, the evidence trickle argument for permissivism (§2.2) is consistent with the evidence always pointing a particular direction.

Talk about evidence “pointing” is admittedly a bit metaphorical. A more formal way to frame the issue is whether *a proposition has a precise epistemic probability on the evidence*—or at least, precise enough. In cases where (dis)believing *p* is required, *p* normally has a precise high (low) epistemic probability on the evidence; in cases with withholding belief on *p* is required, the epistemic probability of *p* is normally around 0.5.

Such cases contrast with cases where the evidence is complex, multifaceted, and qualitatively diverse. In such cases, the evidence for *p* and for not-*p* is *not* perfectly parallel and is difficult to cleanly add together (Tucker & Jackson 2026 suggest that, in such cases, the evidence for *p* and for not-*p* is on a par, rather than tied). For example, the considerations for *p* might be testimonial or intuition-based, whereas the considerations for not-*p* might be inductive or perceptual. Consider the Quarterback case above: the analytics strongly support his future success, but what you perceive on film is a bit more pessimistic. The extent to which the analytics support his future success is partially undermined by other considerations (e.g. a strong offensive line). Coaches and scouts disagree and point out different aspects of his performance and demeanor. And so on. In cases where the evidence is complicated, rather than tied (or nearly tied), the probability of the target proposition on the evidence is not 0.5; it is imprecise, and, when represented as a probabilistic interval, straddles 0.5 but the upper and lower bounds are reasonably distant (i.e. both somewhat far away from 0.5). Intuitively, in such cases, the evidence does not “point” in a particular direction. Then, it does not require belief, it does not require disbelief, and it does not require withholding. Further, I maintain that, in such cases, belief is permissible, disbelief is permissible, and withholding is permissible. And similar cases where the epistemic probability of *p* is imprecise but largely on one side of 0.5 will yield bi-optionality.

More formally, an argument for epistemic options from the nature of evidence might look like this:

1. Possibly, the evidence regarding a proposition *p* is complicated and qualitatively diverse.
2. If the evidence regarding *p* is complicated and qualitatively diverse, then it doesn’t “point” in a single direction (i.e. the epistemic probability of *p* is imprecise and spread over a relatively wide interval).
3. If the evidence regarding *p* doesn’t “point” in a single direction, then there are epistemic options.
4. Therefore, there are epistemic options.

Premise 1 is supported by the myriad of cases discussed and categorized in the previous section (§3); I suspect most readers will find it plausible. In favor of **Premise 2**, recall the discussion above, where we noted that, while evidence can support *p* (rationally require belief), support not-*p* (rationally require disbelief), and be tied or nearly tied (rationally require withholding), not all bodies of evidence and propositions fall into one of these three categories. When there’s a good bit of evidence in favor of *p* and other evidence in favor of not-*p* and that evidence is heterogeneous, the evidence on each side doesn’t cleanly “add together” or “balance out.” In such cases, the evidence does not point in a single direction—there is not a specific epistemic probability of *p* on the evidence. Those sympathetic to the existence of imprecise epistemic probabilities will likely find this inference plausible (cf. Horowitz et al. 2024). (NOTE: I would like to say more here and perhaps draw on some of the imprecise probability literature. Suggestions welcome!)

Premise 3 suggests that, in cases where the evidence does not point in a particular direction, we have epistemic options. Say you are trying to figure out whether *p*. However, following your

evidence does not take you to p , nor does it take you to not- p , nor does your evidence for p and not- p balance out and require withholding. What attitude should you take? Well, some might suggest that withholding is still required in such cases. While I'll respond to this suggestion in more detail in the next chapter, note that there's a principled difference between such cases and tie and near-tie cases, where we all agree withholding is rationally required. Further, I worry such a response licenses too much skepticism—assuming limited and complicated evidence is relatively pervasive—and skepticism where skepticism isn't required—while withholding in the face of complicated evidence seems fine, as it doesn't seem rationally required to lack beliefs about controversial or complex matters. (See the next chapter for more on this.)

Another possibility is that in cases where your evidence doesn't point to one attitude, you're in an *epistemic dilemma* where belief, withholding, and disbelief are all irrational, so there is no rational attitude you can take to the proposition; some formulations of impermissivism allow for this possibility (see Feldman 2007 and Magnus 2020). Note, interestingly, that this response concedes that all three attitudes can be rationally on par; it just gives them all an equally negative evaluation instead of an equally positive one. In response, note that the cases I have in mind are importantly different than the cases where epistemic dilemmas are traditionally posited, such as cases where a thinker cannot understand a proposition (Matheson 2011: 361) or a paradox such as the liar paradox (e.g. “this sentence is false”; see Caie 2012: 19). Positing epistemic (and moral) dilemmas also has the cost of denying ought-implies-can style principles (cf. McConnell 2018).

Forget these dark dilemmas! Come to the light side. The glass is half full, not half empty. Instead of dishing out irrationality to everyone, why not dish out rationality? If all our options have the same status vis-à-vis rationality, I vote we call them all equally good, not equally bad. The latter requires us to deny ought implies can and to posit a new class of epistemic dilemmas. The idea that, in cases where the evidence doesn't clearly point in one direction, we have epistemic options, is plausible, especially relative to alternatives (e.g. requiring withholding or positing dilemmas). It is also an intuitive reading of the cases from §3. This is compelling support for Premise 3.

5. Good Explanations

Inference to the best explanation (IBE) is a common inference pattern. According to IBE, we should choose (or believe) the hypothesis or theory that best explains that data (i.e. the available evidence). For example, if the cookie jar is empty and your child has crumbs all over his face, you should conclude he ate the cookies. Or if a patient has a cough, fever, fatigue, and opacity on an x-ray, the doctor should conclude they have pneumonia. While it is controversial how often we should appeal to IBE (especially in fields like science), most agree that it's at least sometimes a rational inference pattern (see Lipton 2000).

In standard cases of IBE, there is a single best explanation that stands out above the others. But note that there are cases that are very structurally similar to cases of IBE, but there is not one single best explanation. Instead, there are several equally good explanations of the evidence. In such cases, I propose that a form of reasoning similar to IBE can sometimes rationalize inferring to *one of the best explanations*. (I'll focus on cases where there are two “sufficiently good” explanations. I am open to this working with perhaps three explanations, but likely not more.)

Lipton (2000: 5) lists five explanatory virtues: “scope, precision, mechanism, unification and simplicity. Better explanations explain more types of phenomena, explain them with greater precision, provide more information about underlying mechanisms, unify apparently disparate phenomena, or simplify our overall picture of the world.” Different explanations might have some of these features but lack others; there might also be tradeoffs between these virtues (often, explaining more phenomena comes at the cost of simplicity, for example). Given these diverse virtues that often involve tradeoffs, it makes sense that we would often find ourselves faced with situation where there are multiple explanations that are quite good, and perhaps good in different ways, but there isn’t one that is determinately best. Different explanations may excel along different dimensions without any one of them being overall superior. Assuming the explanations on the table are mutually exclusive, in such cases, thinkers may have epistemic options: if the explanations are good enough, they can rationally believe either of them.

We see this structure in several of the cases above. Consider the God case, for example. The thinker is faced with two live possibilities—theism and atheism—that both explain the evidence decently well. The features that make each a good explanation might be quite different: for example, theism might explain certain bits of data, such as why the world seems designed or the longing for ultimate meaning, purpose, or eternal existence. Atheism might better explain apparently pointless evil, and it might also have the virtue of being simpler. Both explanations do an adequate job of making sense of a complex body of evidence. They do this well enough that adopting, and thereby believing, either one can be rational.

A similar structure applies in the cases of interpreting behavior. For example, when we interpret the behavior of others to try to determine their beliefs, desires, motives, etc., we are often faced with multiple explanations that all make sense of the data. In the dishwasher case, both the idea that your spouse’s action was intentional and the idea that it was forgetful explain parts of your dataset: the former fits with the tense morning, but the latter fits with past behavior. Because both explanations adequately account for the complicated evidence, believing either can be rational.

More formally, an argument for epistemic options from this extension of IBE might look like this:

1. It can be rational to believe the best explanation of the evidence.
2. Sometimes there are two significantly good, mutually exclusive explanations of the evidence.
3. If it can be rational to believe the best explanation of the evidence, then, when there are two significantly good explanations of the evidence, it can be rational to believe either.
4. Therefore, it can be rational to believe either explanation; i.e. there are epistemic options.

Premise 1 simply amounts to the idea that sometimes, IBE-reasoning results in rational beliefs. While the exact scope and application of IBE is controversial, the idea that it can, at least in certain cases, lead us to rational beliefs is plausible. **Premise 2** is supported by the earlier discussion in this section: there are often two competing explanations, each with different explanatory virtues, that both do quite well with respect to the evidence, where neither is a clear winner. Here, an explanation counts as “significantly good” insofar as it exhibits a meaningful degree of relevant explanatory virtues. (This is intentionally somewhat vague, as I don’t want to take a stand on the exact threshold of goodness here, except that mutually exclusive explanations can both meet it at the same time.)

Premise 3 is likely to be the most controversial, so I will mention a few lines of support for it. First, it is supported by a range of intuitive cases. For example, it does not seem irrational for a scientist to believe that an asteroid impact caused the extinction of the dinosaurs, even while recognizing that other explanations of the available evidence may be comparably good. Likewise, it does not seem irrational for someone to believe in God while acknowledging that the relevant evidence—for instance, apparent design or reports of miracles—might also admit of plausible naturalistic explanations. Second, notice what the premise does and does not claim. It does not say that any explanation is fit to be believed. Rather, the explanation in question is already among the strongest explanations of the available evidence, and no competing explanation is explanatorily superior. Believing a weak or poorly supported explanation would be irrational, as would believing a subpar explanation when a better one is available. But when a theory is among the best available explanations of one's evidence, and there is no uniquely best explanation, sometimes it is permissible to believe it.

Third, when we engage in traditional IBE, the explanation that we infer to meets a certain threshold of explanatory virtue. An explanation doesn't have to be maximally virtuous for us to rationally believe it. Compare this to the case where there are two significantly good explanations. Sometimes, those two explanations exhibit explanatory virtues that are on a par with the inferred explanation in cases of traditional IBE (yes, of course they both cannot be the best explanation, but they can nonetheless have a high degree of virtue). While we aren't inferring to a single, best explanation, we are still inferring to an explanatory with quite a bit of explanatory power and virtue. Finally, while we'll again cover this objection in more detail in the next chapter, withholding isn't rationally required in such cases. This would force us to withhold belief on quite a lot of things, both theoretical and everyday, as there are often multiple sufficiently good explanations of our evidence.

6. Conclusion

I've argued that epistemic options are both real and relatively common. I motivated this with a range of intuitive cases: from past events, future predictions, theoretical "big" questions, interpreting others, and interpreting oneself. Then, I offered two arguments for epistemic options. The first appealed to the nature of evidence: when evidence is limited, heterogeneous, and difficult to integrate, it doesn't determine a unique rational attitude. The second appealed to inference to the best explanation: explanatory considerations do not always identify a uniquely best explanation, and when several explanations are comparably good but mutually exclusive, it can be rational to adopt any one of them.

These arguments may leave the reader with unanswered questions and objections. When multiple belief-attitudes are rationally available, isn't holding one attitude (as opposed to another) arbitrary in a problematic way? How does this all fit with the idea that rational beliefs are supposed to connect us to truth? Why shouldn't we just suspend judgment when the evidence is complicated or when there's not one clear, best explanation of the evidence? The next chapter will turn to these objections.

References

- Antill, Gregory. (2019). "Evidence and Self-Fulfilling Belief." *American Philosophical Quarterly* 56(4): 319–330.
- Caie, M. 2012. Belief and indeterminacy. *Philosophical Review* 121: 1–54.
- Drake, J. (2017). Doxastic permissiveness and the promise of truth. *Synthese* 194 (12):4897–4912.
- Feldman, R. (2007). Reasonable Religious Disagreements. In *Philosophers without Gods*, ed. L.M. Antony, 194–214. Oxford: Oxford University Press.
- Forrester, Paul (2024). A New Argument for Uniqueness about Evidential Support. *Episteme* 21 (4):1265–1286.
- Greco, Daniel & Brian Hedden. (2016). "Uniqueness and Metaepistemology." *The Journal of Philosophy* 113(8): 365–395.
- Horowitz, Sophie, Sinan Dogramaci, and Miriam Schoenfield. (2024). "Are You Now or Have You Ever Been an Impermissivist? A Conversation Among Friends and Enemies of Epistemic Freedom." In *Contemporary Debates in Epistemology*, 3rd ed., edited by Blake Roeber, Ernest Sosa, Matthias Steup, and John Turri, pp. 103–125. Hoboken, NJ: Wiley-Blackwell.
- Kopec, Matthew (2015). A Counterexample to the Uniqueness Thesis. *Philosophia* 43 (2):403–409.
- Kopec, M. & Titelbaum, M. (2016). "The Uniqueness Thesis." *Philosophy Compass*, 11(4), 189–200.
- Lipton, P. (2000). Inference to the best explanation. In W. H. Newton-Smith (Ed.), *A Companion to the Philosophy of Science* (pp. 184–193). Blackwell.
- Marxen, Chad (2020). On the epistemic rationality and significance of self-fulfilling beliefs. *Synthese* 199 (1-2):4243–4260.
- Magnus, P.D. (2020). "Evidential Nihilism." *Analysis* 80(4): 674–683.
- Matheson, J. 2011. The case for rational uniqueness. *Logos & Episteme* 2: 359–73.
- McConnell, T. 2018. Moral dilemmas. In *The Stanford Encyclopedia of Philosophy*, ed. E.N. Zalta. <https://plato.stanford.edu/archives/fall2018/entries/moral-dilemmas/>.
- Milkman, Katy. (2021). *How to Change: The Science of Getting from Where You Are to Where You Want to Be*. New York: Portfolio/Penguin.
- Pettigrew, Richard (2016). Jamesian epistemology formalised: An explication of 'the will to believe'. *Episteme* 13 (3):253–268.
- Pettigrew, Richard (2022). *Epistemic Risk and the Demands of Rationality*. Oxford, UK: Oxford University Press.

Quanbeck, Z. & Alex Worsnip. (2025). “A Permissivist Alternative to Encroachment.” *Philosophers’ Imprint* 25(1): 1–22.

Raleigh, Thomas (2017). Another Argument Against Uniqueness. *Philosophical Quarterly* 67 (267):327–346.

Roeber, Blake. (2019). “Evidence, Judgment, and Belief at Will.” *Mind* 128(511): 837–859.

Roeber, Blake. (2020). “Permissive Situations and Direct Doxastic Control.” *Philosophy and Phenomenological Research* 101(2): 415–431.

Rosen, Gideon. (2001). “Nominalism, Naturalism, Epistemic Relativism.” *Philosophical Perspectives* 15: 69–91.

Stapleford, Scott & Jackson, Elizabeth (2026). *Must Beliefs and Evidence Agree? A Debate*. New York: Routledge.

Tucker, Chris & Jackson, Elizabeth (2026). Parity and the Permissivism Puzzle: A Defense of Epistemic Options. *Philosophy and Phenomenological Research* 113 (1):233–247.

White, Roger. (2005). “Epistemic Permissiveness.” *Philosophical Perspectives* 19: 445–59.

White, Roger. (2014). “Evidence Cannot Be Permissive.” In *Contemporary Debates in Epistemology*, 2nd ed. Eds. M. Steup, J. Turri, and E. Sosa. Hoboken: John Wiley & Sons, 312–23.