

The Game of Belief

Barry Maguire

Stanford University

Jack Woods

University of Leeds

1. Introduction

It is exceedingly plausible that there are distinctively epistemic reasons that provide normative support for or against beliefs.¹ Call this thesis *evidentialism*. For instance, the fact that you see a walrus on the ottoman is a reason to believe there is a walrus on the ottoman. A stronger thesis, *anti-pragmatism*, maintains that epistemic reasons are the *only* kinds of reasons for belief. This thesis is significantly less attractive than eviden-

Special thanks to two anonymous referees for extraordinarily helpful feedback and to Sinan Dođramaci, who gave comments on a predecessor at the Southern California Epistemology Workshop at the Pacific APA in 2017. Thanks to Derek Baker, Rima Basu, Selim Berker, Darren Bradley, Michael Bratman, Chris Cowie, Jorah Dannenberg, Edward Elliot, Matti Eklund, Chris Howard, Nadeem Hussain, Jessica Isserow, Jessica Keiser, Benjamin Kiesewetter, Matt Kotzen, Krista Lawlor, Taylor Madigan, Meghan T. Maguire, Susanne Mantel, Milan Mossé, Austen McDougal, Erum Naqvi, Ram Neta, David Plunkett, Andrew Reisner, Susanna Rinard, Kate Ritchie, Léa Salje, Geoff Sayre-McCord, Karl Schafer, Daniel Star, Pekka Väyrynen, Robbie Williams, and to participants at the Pacific APA in 2017; the Normativity Workshop, NYU Abu Dhabi in 2016; the Workshop on Epistemic Consequentialism at the University of Kent, June 2015; the Centre for Ethics and Metaethics (at Leeds); and a graduate seminar at Stanford in 2017.

The order of the authors' names is alphabetical. Both authors contributed equally. Contractions are all Jack's.

1. For simplicity, we assume all epistemic reasons are evidential; for some arguments against this see Wright (2004), Talbot (2014), and Schroeder (2015). Throughout, we will be neutral between different varieties of evidentialism and different ways of distinguishing epistemic reasons.

tialism since a range of different cases provide support for the idea that we have practical reasons for belief.

Here are some examples. Epistemically well-supported beliefs—or as we'll say, *correct*² beliefs—are often very useful. Believing that the ice will hold your weight, or that the snake is friendly, or that your meeting is in Tokyo, will help you move around the world more effectively. But correct beliefs are often perfectly useless. There is little point in having correct beliefs about how many times the letter “e” appears in this paragraph. If beliefs take up limited storage space, then you have at least one reason *not* to have beliefs about these matters (Harman 1986). There may also be reasons to conform to the most useful of several equally well-evidenced doxastic states (Greaves 2013).

Correct beliefs might also be unpleasant or harmful: for instance, concerning how sausages are made, or the probability that your marriage will end in divorce. And incorrect beliefs might be beneficial or otherwise warranted. We plausibly have reasons to think better of our friends than the evidence would suggest (Stroud 2006; Keller 2004; Way and McHugh 2016; Crawford 2019), reasons to have certain beliefs that enhance our “self-esteem” (Kelly 2003), and reasons to be more optimistic than the evidence suggests about your chances of recovering from some challenging disease (Reisner 2008; Rinard 2015).

There are also cases where we have practical reasons to believe a certain way independently of the balance of evidence. These include believing that everyone is capable of significant moral improvement (Preston-Roedder 2013), and that there is no correlation between IQ and being the member of an oppressing class (Gendler 2011). Some beliefs are morally wrongful or unjust (Basu 2018). In a range of more or less fanciful cases, you can be offered a positive or negative incentive for being in some doxastic state (see, for example, Reisner 2008; Way 2012); Pascalian (Pascal [1670] 1995) or Jamesian (James 1896) reasons to believe that God exists may also fall into this category.³

Now, perhaps you are sufficiently confident, pragmatic, and robust, as are all your friends, that many of these claims do not apply to

2. Some readers might prefer “fittingness” to “correctness” here. We prefer the language of “correctness,” largely because it generalizes more satisfyingly to a range of actions. For further discussion, see Chappell (2012) and Howard (2019a). It is important that our usage is consistent with, but does not entail, that a belief is correct only if it is true. The evidentialist intuitions we want to preserve can be developed either way.

3. For additional cases, see especially Chignell 2016, Marušić 2013, McCormick 2014, and Williams and Saunders, forthcoming.

you. Perhaps your friendships and your pride and your optimism are steadied by resolute commitment to the best epistemic standards. Still, they might not have been. It is possible that your happiness, or the value of your relationships, or the strength of your motivation to do this or that depend on your beliefs, but not in ways that correlate impeccably with whether those beliefs are correct. These appear to be psychological possibilities that support alternative ethical possibilities. But anti-pragmatism is supposed to be necessarily, perhaps even conceptually, true. So any of these possible scenarios will yield counterexamples.

Anti-pragmatists need to explain all these cases away. One common strategy is to argue that putative practical reasons for beliefs are, instead, reasons for other actions or attitudes that bear on these beliefs: for *wanting to have the belief*, or for *attempting to bring about the belief*, or perhaps for *behaving as though one has the belief*.⁴ Since it's common to refer to epistemic reasons for belief as the "right kind of reasons" and others as the "wrong kind of reasons," this strategy is often called *wrong kind of reasons skepticism* (Way 2012).⁵ As Way explains (2012: 511): "This is a kind of error theory. It holds that people mistakenly think that incentives for attitudes are reasons for those attitudes because they confuse reasons for attitudes with reasons to want or bring about an attitude." This may be a little strong. But in any event, an error theory needs an argument. Anti-pragmatism itself has been supported by three related arguments, which share the following form. First, some feature is identified that distinguishes epistemic reasons from putative practical reasons for beliefs. Second, it is argued that only considerations with this distinguishing feature are reasons for those beliefs. The conclusion is drawn that all reasons for belief are epistemic reasons (modulo the qualifications in footnote 1).

In the *followability* argument (Shah 2006; Kelly 2003; Kolodny 2005; Raz 2013), it is maintained that epistemic reasons for beliefs, but not putative practical reasons for beliefs, are followable in the sense that one can come to be in that state for that very reason. It is further main-

4. There is little consensus among skeptics concerning *which* other states putative wrong kinds of reasons support. For an impressive taxonomy of proposals, see Berker, forthcoming.

5. Cf. Rabinowicz and Rønnow-Rasmussen 2004; Danielsson and Olson 2007. For some theoretical challenges, see Reisner 2008; for some replies, see Way 2016. Strictly speaking, anti-pragmatism isn't committed to any particular alternative view about putative practical reasons for beliefs. Their distinctive thesis is the denial that there are any nonepistemic reasons for beliefs.

tained, roughly, that a consideration is a reason for a state only if one could be in that state for that reason. Hence, according to these theorists, there are no practical reasons for beliefs.

In the *transmission* argument, Jonathan Way (2012) claims that putative practical reasons for beliefs transmit to states that will promote those beliefs, while epistemic reasons for beliefs do not transmit to states that will promote those beliefs. For instance, a practical reason to believe p yields an instrumental reason to take a pill that will result in your believing that p , but an epistemic reason to believe that p does not. Generalizing, it's claimed that reasons for desires and actions transmit to states that promote those desires and actions; reasons for beliefs do not. Way claims that the best explanation of the fact that practical reasons for doxastic states transmit, but epistemic reasons don't, is that practical reasons are reasons for actions bearing on beliefs, rather than reasons for the beliefs themselves.

Here we're primarily concerned with a third argument that addresses the *weighing* of epistemic and practical reasons (Feldman 2000; Berker, forthcoming). It is quite plausible, in general, that facts about what states you ought to be in are explained by weighing the reasons for and against the alternatives. This explanatory structure is familiar in the domain of practical normativity, and it is an increasingly popular way to think about the domain of epistemic normativity.⁶ If epistemic and practical reasons both bear on what you ought to believe, then it is natural to expect that what you ought to believe is explained by weighing all these reasons together. But epistemic and practical reasons seem just too different to be weighed against each other. It seems like a category mistake to ask whether a particular epistemic reason to believe p is weightier than a particular practical reason not to believe p —akin to asking whether my balcony is longer than your minuet. Since it seems clear that epistemic reasons bear on what you ought to believe, this line of thought suggests that practical reasons do not.

This line of thought is most often presented as an intuitive argument.⁷ Recently, Selim Berker has argued against pragmatism by noting a structural difference between the weighing behaviors of episte-

6. See Schroeder 2010, 2011. For some important further differences between the weighing behavior of reasons for beliefs and actions, see Berker, forthcoming.

7. See, for example, Steglich-Peterson 2011; Chappell 2012; and Papineau 2013; though see Reisner 2008; Steglich-Peterson 2009; Kelly 2003; and Feldman 2000 for complications.

mic reasons and practical reasons. Roughly, the situation is this: when we have equally balanced practical reasons for exclusive alternative states, and no other relevant practical reasons, then any of the alternatives is permissible. But when we have equally balanced epistemic reasons for one state (say, the belief that p) and another (say, the belief that not p), it is not the case that either is permissible. In such a case, one is required to suspend belief concerning p . Berker argues that the best explanation of this difference is that practical reasons are not reasons for beliefs at all, but for belief-related activities, as the “wrong kind of reasons skeptic” suggests.

It is not our intention to *refute* wrong kind of reason skepticism in this essay.⁸ Our reluctance to accept the skeptical approach is based on a methodological commitment to saving the appearances where possible. We wish to pursue an account that takes the intuitive responses to these cases at face value. It seems worthwhile to look for an explanation of the proliferation of cases that vindicates the idea that these are all cases of practical reasons for beliefs, or, in other words, one that vindicates *pragmatism*.

But we also want to vindicate *evidentialism*—or at least, to vindicate many of its key features.⁹ After all, evidentialism is also extremely intuitive, as Kelly (2003) emphasizes. We maintain that each of these three arguments for evidentialism correctly identifies a feature of epistemic reasons for belief. However, each mistakenly infers from the fact that epistemic reasons have these distinctive features that there are no practical reasons for belief. Our plan is to introduce a general framework that incorporates both epistemic reasons and practical reasons for beliefs and that explains how they interact.

8. In fact, we find it plausible that in some familiar cases, putative nonepistemic reasons for belief are really reasons for actions rather than beliefs. For instance, the fact that the madman threatens to shoot you unless you believe that p is, in the first instance, a reason to get the madman not to shoot you. Perhaps the best way to do this is to believe that p . Probably not. The best response to the gunman demanding that you believe that p may instead be to pull the trapdoor latch. This view about some cases is consistent with thinking that there are plenty of other cases in which there are nonepistemic reasons for belief. In our view, extreme incentive cases are not the most helpful in thinking about epistemic normativity. It is partly for this reason that we address a range of different cases.

9. Of course, we will not vindicate anti-pragmatism or “authoritative evidentialism”—the thesis that epistemic reasons are authoritative reasons (see section 5.2). We deny that these are necessary conditions for a satisfactory theory of these matters. We also mention a couple of more substantive concessions to certain intuitions in the conclusion.

We start with the conjecture that the distinction between epistemic reasons and practical reasons for belief is an instance of a more general distinction between considerations bearing on the correctness conditions for activities and considerations bearing on participation in activities other than by bearing on the correctness conditions (Schroeder 2010).¹⁰ In particular, this distinction applies not just to psychological states but also to actions in the context of activities. Mark Schroeder has argued that this is one compelling interpretation of the general distinction between the “right” and “wrong” kinds of reasons.¹¹ There seem to be right and wrong reasons for doxastic states, affective states, conative states, and also for actions in the context of activities. There are right and wrong kinds of reasons to move your knight to f4 or to give a student an “A” for their paper. The right kinds of reasons to give a student an “A” concern the quality of the essay; examples of the wrong kinds of reasons to give an “A” might be that you were paid, threatened, or sympathetic to their recent life hardships. The “right” kinds of reasons to move your knight are facts about whether the move will increase your chances of winning the game in the context of the relevant rules and conventions; again, the “wrong” kinds of reasons include threats, incentives, and sympathy. Quite generally, wherever there is an activity-specific standard of correctness, there is a “right”/“wrong” reasons distinction.

There are interesting questions about followability, transmission, and weighing in all of these cases. But the inference to skepticism about the “wrong” kind of reasons for actions, in the context of activities, is extremely unattractive. An account of the interaction of right and wrong kinds of reasons for actions in the context of activities provides a more attractive model with which to explain epistemic normativity. This model also can explain the followability and transmission differences noted above, capturing a wide range of both evidentialist and pragmatist insights. The generality of our account is furthermore advantageous since it offers a

10. See also Sharadin (2013); Danielsson and Olson (2007); and Howard (2018).

11. We take no stand on the question of whether this is the right or only interpretation of the right/wrong distinction. We suspect there are various distinctions in the ballpark (see the “trapdoor” point in footnote 8). Our conjecture is just that there are structural analogies between standards-based and other reasons bearing on actions in activities and beliefs, and that reflection on other instances might be fruitful in thinking about the ethics of belief. See Danielsson and Olson (2007) for an account of “right” and “wrong” types of reasons along the lines we’ve sketched here. For a useful general summary, see Gertken and Kiesewetter 2017.

framework for thinking about the interaction of right and wrong reasons across a range of cases.¹²

Let us quickly sketch our proposal. In our view, correctness-based reasons, like the distinctively epistemic considerations bearing on beliefs, are normative reasons, but they are not *authoritatively* normative reasons.¹³ Authoritative reasons are those that play a particular role in explaining what you *just plain ought* to believe.¹⁴ We conjecture that all and only the *practical* reasons are the authoritative reasons.¹⁵ We are neutral concerning which considerations are practical reasons, for instance as between value-based, desire-based, duty-based, or reasons-fundamentalist approaches. In our thinking about cases, we generally default to a simple value-based view, according to which practical reasons are facts about the promotion of valuable states of affairs. We often have practical reasons to have correct beliefs. And here, epistemic reasons play a key role. The interaction of epistemic reasons explains which belief is correct in a given situation.

This is analogous to other activities. Chess-based reasons—that this move will win material, or that that move will cost you a tempo—are not authoritative either. They don't explain what you just plain ought

12. We don't mean to deny there are interesting questions about how far or straightforwardly our account generalizes. For one example of some related complexity, see G. A. Cohen's remarks on differences between public and private activities (2008: 135).

13. We're using *authoritative* where we might have used *genuine*, *substantive*, *generic*, or *robust* (compare McPherson 2018; Wodak 2018; Woods 2018). We prefer this terminology because it avoids conveying that the correctness-based reasons are unreal or unimportant. They are real and often important. It is just that they don't play the reasons role in explaining what one just plain ought to do. They play the reasons role in explaining which state would be correct given the operative standards. We'll say more about all this in section 4.

14. We borrow this terminology from McLeod (2001). This is roughly the same concept as Reisner's (2018) "non-domain specific all things considered ought," and the "all things considered ought" in Maguire and Lord (2016); see also Rinard's "guidance-giving should" (2019), McPherson's "deliberative ought" (2018), and Woods's (2018) "generic ought." We discuss this more in section 4.

15. Of course, we deny that practical reasons are all reasons for action. This way of thinking about things supplants the traditional "theoretical"/"practical" reasons divide—which we think is all to the good. As will become clear, our distinction between "correctness-based reasons" and "practical reasons" tracks the distinction between *internal* standards of activities and sources of support for participation in those practices. But we are not here identifying practical reasons with the latter. There might be a practical reason for something other than participation in an operative standard-bearing activity, for example, to scratch your nose. If the reader prefers, they can substitute "value-based reasons" or even "desire-based reasons" for "practical reasons" throughout.

(henceforth: ought) to do. Practical reasons do that. But we sometimes have practical reasons to make correct chess moves—where by this we mean: make the strategically best move compatible with the rules of the game. The interaction of chess-based reasons explains which move is correct. This structure preserves a key normative role for epistemic reasons for beliefs, and a key normative role for practical reasons for beliefs. Each plays a distinctive role in the explanation of what we ought to believe.

The structure of the article is as follows. In section 2, we'll discuss several features of the interaction of correctness-based and practical reasons for action-involving activities. In section 3, we'll show that these features carry over to the interaction of correctness-based and practical reasons for belief; we'll also draw out some general morals for the theory of interaction of different kinds of reasons. In section 4, we discuss the key notions of authoritativeness and the just plain ought in more detail. In section 5, we compare our proposal favorably with some neighboring accounts.

2. Weighing Right and Wrong Reasons for Actions

Suppose you are playing chess with a friend. Passing by, I threaten to punch you on the nose if you advance your knight to f4.¹⁶ What ought you to do? Can you weigh the positional advantage in advancing your knight against the prudential disadvantage in getting a punch on the nose? Yes and no. There is no chess scale on which positional advantage can be compared with the painfulness of a punch on the nose. Positional advantages in chess can be weighed against disadvantages in material (pieces on the board), tempo (development per number of moves), and so on. Positional advantages cannot, as chess is normally understood, be weighed against punches on the nose.

But you aren't completely at sea either. You can distinguish two questions. How important is it that you win this game? And to what extent would failing to advance your knight here affect your chances of winning? The first is a question about your reasons to play this game (competitively). The second is a chess-specific question about the overall merits of moving your knight to f4, in terms of increasing your chances of winning fairly relative to the moves available to you. A modern chess engine like Stockfish will even put a number on this for you. Stockfish, as you would

16. Apologies.

expect, does not consider the merits of your being punched on the nose if you play Nf4.¹⁷

To answer the original question about whether to play the move and take the punch, you need to answer the two questions we distinguished. That will give you the importance of this very move as a function of *both* its activity-specific significance and the value of playing competitively. You need to think about the harm involved in being punched on the nose, and also about what is at stake in playing the game. You can then weigh the practical reason to avoid the punch against the practical reason to try as hard as you can to win the game (which, of course, involves conformity to activity-specific standards). When you are playing the world championship, and knight to f4 is a much stronger move than any alternative, you take the punch. When you are playing your friend to pass the time and nothing especially hangs on winning, you don't take the punch.

Similar remarks apply to positive incentives to make suboptimal moves. Perhaps it would be better in some particular situation to bend the rules a little to provide encouragement to a learner, or if your opponent needs cheering up. Of course, it is not the move as such that cheers up your opponent. They don't glow with pleasure whenever someone plays pawn to a3. They are cheered up by the prospect of winning or by the opportunity to make a good move. For you to have most reason to do this, of course, the expected value of cheering your opponent up will need to be greater than the expected disvalue not just of disappointment but also of the dissembling that is involved, along with the risk of their finding out. This latter, in turn, is a function of the probability of their finding out and the probability of their reacting badly if they did. Perhaps they would appreciate the gesture.

That a move available to you is correct by the standards of good chess doesn't necessarily entail that you have any *authoritative* reasons to make it. To see this, consider some grisly scenario in which people are killed every time you make a move in chess, in proportion to the internal merits of the move—for example, the higher the Stockfish number, the more people killed. Even more will be killed if you stop playing altogether. Suppose you hate playing chess anyway. Perhaps nothing in such a grisly scenario really speaks in favor of playing Stockfish-better moves—there are no authoritative reasons to do so. But that doesn't mean that those

17. The Stockfish merits of one move depend on your playing other future moves that are also Stockfish-optimal. If you aren't likely to do so, then the Stockfish-optimal move might not actually be the best move for you to play. We ignore this complication.

moves aren't still Stockfish-better, and hence better supported by the activity-specific standards of good chess. The moves you really ought to play in this scenario are still the most incorrect by the standards of good chess.

Similar remarks apply to a wide range of activities. For some size of a charitable donation, or some intensity of threat, you really should give a bad paper an "A." This doesn't make it academically responsible to do so—it is the height of academic irresponsibility—but still, sometimes needs must.

In these cases, there will be some further explanation of why the standards of the activity are thus and so—for instance, why it is impermissible to castle in check, or impermissible to consider effort in assigning grades. There will also be some further explanation for why these standards are *operative* in a context—for instance, why, in some other case, the standards of draughts and not chess apply, even though you are playing with chess pieces (since you couldn't find any draughts). The explanations of these facts will sometimes appeal to practical reasons: for instance, that the impermissibility of castling in check enables more interesting attack play, or that we thought draughts would be more relaxing than chess.

The important thing for current purposes is to take care in specifying what is supported by which reason: whether a rule is to be part of an operative standard for an activity, whether to engage in particular activities that are partially constituted by such rules, or whether to make some specific move within the context of this activity. Even if the status of a rule, or its operativity in a context, is explained by some practical reason, it doesn't by any means follow that that reason favors making some particular move in the game (compare Rawls 1955: 26).

3. Weighing Right and Wrong Reasons for Beliefs

As with chess, there appear to be intelligible questions about what one ought to believe in cases where there are both epistemic reasons and practical reasons bearing on having the belief. Let us run through the different kinds of examples that we started with to illustrate how this analogy plays out.

Let whether p be of immediate interest and bear on a number of other important inquiries. The simplest practical reasons against having a correct belief concerning p are *clutter avoidance* reasons—that doing so would take up some storage and processing power. Clutter avoidance is

easily outweighed by considerations like immediate practical need. In such cases, one will have most practical reason to have a correct belief concerning whether p . In these cases, the epistemic reasons carry the day. There is a clear analogy here with cases where you have most reason to play whatever chess move would be most correct by the standards of good chess. But in both cases, there is an important explanatory difference between the claim that (1) you ought to be in some state just because it is correct, and the claim that (2) you ought to be in some state because it is correct and you have most reason to be in the correct state.

In other cases, there are no practical reasons to have any belief concerning whether p . In such cases, plausibly, the clutter avoidance reasons carry the day. This result is explanatorily independent of the question of which attitude toward p would be epistemically correct. That latter fact is explained by all the epistemic reasons for and against in the normal fashion—that is, by the epistemic reasons bearing on whether p . But although some doxastic attitude toward p is epistemically correct even in cases of useless trivia, it's not the case that you ought to have this attitude, because you ought not to have any attitude toward p whatsoever.¹⁸

Very well, you might say, but supposing that you *do* have some doxastic attitude toward p , wouldn't it be better (in some sense) for that attitude to be correct rather than incorrect? Wouldn't it still be worse to believe confidently that not p in the face of evidence for p than to believe that p , even if a correct belief about p is very unlikely to be of any use to you? Well, yes, it would be worse in one important respect: it would be epistemically worse. It would be less correct. But if there really is nothing to be practically said for having a correct belief concerning p , and, moreover, nothing practically to be said in favor of having the correct belief rather than another, then, you have no more authoritative reason to have either belief than the other.¹⁹

There can also be practical reasons to have beliefs that are incorrect but still sensitive to the evidence. The fact that believing that you will recover will increase the chances that you will recover is, plausibly

18. Assuming, with Friedman (2013), that suspending judgment on whether p is a doxastic attitude toward p (though see Archer, forthcoming), plausibly in such cases you ought not suspend judgment concerning p either.

19. The second condition will not be met if there is a lightweight *standing* practical reason to believe correctly. Since we are officially neutral about what the practical reasons are, we have not ruled this out.

enough, a practical reason to believe that you will recover.²⁰ But if the expected likelihood of recovery is sufficiently remote, then the fact that believing you will recover will increase your chances is, at best, a rather weak practical reason to believe you will recover. Plausibly, you ought to embrace your fate and cherish your remaining time with your loved ones instead. This is a lively issue among hospice workers concerned about when one ought to encourage faith in recovery, and when acceptance of death. It is presumed in these debates that the brute facts about your chances do not settle these questions.²¹ To take a related example, there is a lively debate about the role of statistics in sentencing. It is not assumed in this debate that the general significance of evidence correlates exactly with its epistemic weight (compare Colyvan and Hedden, forthcoming).

Here is another case in which practical considerations can seemingly modify the weight of an epistemic reason. Assume that some social group suffers from testimonial injustice: in virtue of an individual's membership in this group, which we assume is epistemically irrelevant, their testimony is given less weight than it merits (by the lights of the pertinent epistemic standards). We might adopt a policy of epistemic affirmative action, whereby, as a short-term corrective, conscientious epistemic agents give the testimony from members of this group more weight than they judge, having considered their implicit biases, that it merits. We aren't saying this is a good policy on *ethical* grounds. We are simply saying that it appears to be conceptually unproblematic.²²

There are related cases that are more familiar in moral philosophy. Suppose that you are an epidemiologist offering advice. There is data bearing on the question that was uncovered by evil scientists using gruesome methods.²³ Perhaps—at least in low-stakes cases—you should

20. See also Greaves (2013) for similar cases not involving health, as well as cases where our confidence that we'll perform our action exactly matches the objective chance we'll perform it. The latter sort of cases are ones where we're guaranteed to have epistemically correct beliefs, but where there's a strong reason to think that practical considerations intervene to tell us which one of these correct beliefs to have.

21. For one relevant survey, see Ceyhan et al. (2018).

22. It is an interesting further question exactly how this modification works. In particular, does this put pressure on the incommensurability of practical and epistemic reasons? We think not. Full discussion would take us too far afield. We conjecture that this kind of modification is available in a minority of cases and involves a kind of pretending that the practical reason is an epistemic modifier.

23. Although, concerning the Nazi experiments in particular, we note that "in the use of concentration camp inmates and prisoners of war as human guinea pigs very little, if any, benefit to science was achieved" (Shirer 1991: 979).

ignore such evidence. Similar cases involve having “promised to forget” or “promised to remember” certain things—if we’ve promised to forget something, then we have practical reason to not factor that into our evidence for and against certain things, for instance, whether to trust one’s nephew around the family silverware (Marušić 2013).

There may also be cases in which what counts as the operative epistemic standard—that is, the rule taking one from epistemic reasons to (epistemic) correctness—is sensitive to nonepistemic considerations. Perhaps the choice between inductive rules with different “caution parameters” is sensitive to nonepistemic attitudes toward risk, in ways analogous to adding castling to the rules of chess (Field 2000; Buchak 2013). Even assuming that practical considerations cannot affect the output of the epistemic rule given a body of evidence, practical considerations can still affect what counts as the relevant evidence and what counts as the appropriate epistemic rule.

We are moving through all these cases quickly, of course. We do not mean to deny that alternative interpretations of these cases are available. Perhaps the promise doesn’t affect what you ought to believe but which beliefs you ought to rely on in practical reasoning—though we’re skeptical that this is easily done. Or perhaps the promise affects which beliefs you ought to desire to have, or try to have. We are not arguing that anti-pragmatists, whether they are “wrong kind reason sceptics” or not, do not have ways to try to explain these cases.

Rather, our point is that our account offers an attractive and plausible way to explain this broad range of data. Practical reasons can interpose themselves into the explanation of what you ought to believe at a number of stages along the way: whether to have any doxastic attitude toward a proposition, whether to have some particular doxastic attitude, whether to have the correct doxastic attitude, whether to have the attitude that would be correct if you ignored some of your evidence, whether to have some attitude that is related to but distinct from the correct doxastic attitude, and so on.

The crucial thing is to keep track of three things: (1) Which state a reason supports, in the first instance. (2) Whether that reason transmits to support for other states, and if so, which ones. (3) The difference between direct interaction in a reasons explanation of domain-specific standards and mutual participation in the full explanation of what one just plain ought to do. Let us elaborate.

On (1). A practical reason to have a *correct belief* that p cannot be weighed against an epistemic reason to believe that p . Rather, the epistemic reasons bearing on p interact only with each other. Suppose these reasons support believing that p . This fact, together with the practical reason to have a correct belief concerning p , yields a reason to believe that p . This practical reason to believe that p can, separately, weigh against other practical reasons to believe that p , for instance, that believing p might hamper your performance.

On (2). As mentioned above, there seems to be an *asymmetry of transmission*. Practical reasons to be in a correct state transmit, as it were, inwards, to reasons to be in the state that is correct. But correctness-based reasons for a state do not transmit, as it were, outwards, to reasons to be in the state that is correct. That was clear in the *grisly* chess game. It is also clear in a nongrisly chess game. Chess-based reasons to castle do not give you any reasons to move this horse-shaped piece of wood several feet in front of you, unless you have distinct practical reasons to play chess.²⁴

On (3). The epistemic domain is explanatorily autonomous. Facts about correct belief are explained by the evidence and the operative epistemic rule or rules. In particular, the weight of epistemic reasons bearing on whether p is explanatorily independent of how useful it would be to have a correct belief concerning p in this case (compare Kelly 2003).²⁵ This important structural feature in our view was given clear expression in Rawls's "Two Concepts of Rules." The question of which rules of baseball are operative, for instance, is presumably subject to practical revision. But in the middle of a game, if the batter asked for a fourth strike, "this would be most kindly taken as a joke" (1955: 26).

Once the evidence and epistemic rule are fixed, there's no further practical intervention in explaining which epistemic state is correct. In

24. We are in broad agreement here with Way's (2012) observation that the transmission behavior of "right" and "wrong" kinds of reasons is different. He is arguing for skepticism about the "wrong kinds of reasons" for attitudes, and, hence, skepticism about practical reasons for beliefs. His argument is restricted only to defenders of "wrong kinds of reasons" who accept a fitting attitudes analysis of practical reasons (see 506n30). As he allows, views like ours have no trouble with these data about transmission.

25. As noted, the converse is not always true. It may be that the usefulness of a belief is closely related to how well supported it is by the evidence. Take a case in which a sailor is likely to perish in a storm (thanks to Carlos Núñez for this example). Perhaps the sailor's evidence would need to be extremely weighty in order for them to believe that all hope is lost—and, on that basis, rather than fighting the storm, perhaps using their remaining moments to pray.

particular, the weight of epistemic reasons bearing on whether p is explanatorily independent of how useful it would be to have a correct belief concerning p in this case. There is a discrete and autonomous explanation involving distinctively epistemic reasons bearing on a distinctively epistemic standard of correctness. Practical reasons play no role in this particular explanation. But it does not follow that there are no practical reasons for belief. For there is a larger explanation of which this sub-explanation constitutes just one important part.

Putting these pieces together, we get a model for the overall interaction of different kinds of reasons in normative epistemology, and in normative theory more generally. For example, if you have an incentive-based practical reason not to believe that p , this can be weighed against the practical reason to believe p that's transmitted from a usefulness-based practical reason to have a correct belief concerning p . The result of this interaction, other things equal, determines whether you ought to believe that p . Suppose that you need to decide whether to cut the blue wire or the green wire. Cutting one will save the day. Cutting the other will blow up the building. Your evidence strongly suggests that cutting the green wire will blow up the building. But now suppose that you're offered fifty dollars to not believe that cutting the green wire will blow up the building. That is, plausibly enough, a practical reason not to have that belief. But since your evidence strongly supports the belief that cutting the green wire will blow up the building, in this situation you also have a practical reason to believe precisely that. On any plausible account, the "avoid explosion" reason is a lot weightier than the "fifty dollars" reason.

But now imagine a slightly different case. You have the same evidence about the green wire. But now it is impossible for you to affect the wire-cutting in any way. Furthermore, someone will blow up the building precisely if you have the correct belief that cutting the green wire will blow up the building. Plausibly, in this case, you have more reason to not believe that cutting the green wire will blow up the building.

4. Normative Reasons, Authoritative Reasons, Just Plain Ought

Our discussion so far has appealed to a distinction between being normative (a property that epistemic reasons have) and being *authoritatively* normative (a property that epistemic reasons lack). It's time for us to put a bit of flesh on this distinction, given how much work it does for us.

We don't think there is any neutral way to characterize or define the property of being authoritatively normative. One common refrain is

that the mark of the authoritatively normative is the relation to normative reasons (see, for example, Scanlon 1998 and Parfit 2011 for the song and Kiesewetter 2017 for discussion). This is a mistake. There are plenty of normative reasons that are not authoritative but that aren't merely motivating or explanatory reasons.

We propose instead a functional role account of authoritativeness, according to which the authoritatively normative properties are distinguished by their essential relations to facts about the “just plain ought.”²⁶

The “just plain ought” is itself identifiable by its functional role. It is the second “ought” employed in questions of the form “according to this activity I ought to x , but ought I to x ?”²⁷ Sometimes we indicate the just plain notion by slapping a “really” on it, as in “it is *really* the case that I ought to x .” Sometimes this ought is modified with “all things considered,” as in, “I know I morally ought to x , but ought I to x all things considered?” or as in, “I know that I ought, as a ticket inspector, to give this impoverished person a ticket, but is it really the case that I ought to do so?” and “. . . ought I to do so all things considered?” Philosophers with a wide range of different substantive and metaethical views have taken questions such as these to be meaningful; we follow them.

Plausibly enough, questions such as these arise whenever different kinds of normative standards bear on the question of which state to be in. It seems intelligible and substantively plausible that there are facts about what one just plain ought to do in cases involving actions in activities. Indeed, many of the cases discussed so far constitute further evidence in favor of the intelligibility and substantive plausibility of such claims, for instance, concerning which move you ought to make in a game of chess when you have reasons to win and also reasons to go easy. By extension, it is natural to appeal to the just plain ought in order to assimilate different

26. See references in footnote 14. Which relations to the “just plain ought”? Most plausibly: constitutive explanatory relations (or identity). But there are different options here. One important thing to notice is that this is an account of *being authoritative* in terms of relations to the just plain ought. This is importantly different from Broome's (e.g., 2013) analysis of being a normative reason in terms of explanatory relations to the just plain ought.

27. See Schroeder 2011, Broome 2013, McPherson 2018, and Woods 2018 for related discussion and characterizations of notions analogous to just plain oughts and reasons. Sometimes they're characterized in terms of *guidance*, sometimes in terms of *closing deliberation*, sometimes in terms of *breaking conflicts between normative domains*. We take all these to be different ways of attempting to latch onto the same notion and, in some way, for each characterization to be picking out part of the functional role of just plain normative notions.

standards bearing on anything else there can be reasons for—doxastic states, affective states, institutions, and so on.

On a simple model, facts about the just plain ought, like facts about domain-specific standards of correctness, are immediately explained by reasons for and against the state in question. Here we are appealing, as a first approximation, to something like the “for-*x*-ing” and “against-*x*-ing” roles in a weighing explanation of ought.²⁸ As before, it is crucial to distinguish the reasons that play this role in explaining the just plain ought from the reasons that play roles in explaining which moves or states would be correct according to some more specific standard. The reasons that play the former role in explaining the just plain ought are the *authoritative* reasons. Reasons that do not play this role in explaining the just plain ought, even if they play some role in the complete explanation of what you just plain ought to do (by way of explaining what state would be correct), are not authoritative reasons.

This distinction between authoritative reasons and other normative reasons, even if at odds with a piece of popular *weltanschauung*, seems to us fairly natural and well-motivated by reflection on a range of cases. Of course, one could deny that the cases of chess-based reasons, Nazi party reasons, and so on pattern with epistemic reasons. This is an available position, and there is clearly more to be said here. Our view for now is that the patterning hypothesis is theoretically tidy and explanatorily powerful.²⁹

28. These terms are from Broome 2013. See also Maguire and Lord 2016. It is not trivial to say that what you just plain ought to do is explained by just plain ought-making reasons of the kind that participate in a weighing explanation. It might instead be explained by some kind of lexical ordering between domains. Consider the following view, for instance: that you just plain ought to do whatever you morally ought to do; if there is nothing you morally ought to do, then you just plain ought to do whatever you most want to do. On this view, the just plain ought is proximately explained not by some weighing of just plain ought-making reasons. The just plain ought is explained by the “morality is overriding” principle together with some fact about whether morality bears on the matter; if so, what you morally ought to do, and if not, what you most want to do. A similar view bearing on belief would maintain that you just plain ought to believe whatever it is epistemically correct to believe unless the practical costs of doing so are sufficiently high (where sufficiency is itself a practical standard), in which case you ought to believe whatever is practically best (this is a simplified version of the view presented in Reisner 2008).

29. For instance, those, such as Howard (2019a), who maintain that fittingness differs from correctness in that only the former is authoritatively normative will need to provide an account of what makes the difference.

However, this does raise the following question: in what sense are nonauthoritative reasons still *normative*? Unfortunately, it is not much easier to say what it is to be normative than to say what it is to be authoritatively normative.³⁰ We take activity-specific reasons, for instance, clearly to favor the options they favor (relative to the ends or values of the activity in question). In the context of the relevant activity, they are the sorts of thing that could be offered as advice, or justification, or used as reasons in reasoning. They participate in weighing explanations of strict normative facts.

Our main point is that we see no compelling reason to deny the semblance that, in many cases, the relevant standards are normative. They do not seem to be nonnormative properties. They are not merely motivating reasons—though we are sometimes motivated by them. And neither are they merely nonnormative explanatory reasons.³¹ The fact that this or that consideration is a chess-based reason, or an etiquette-based reason, or an epistemic reason does not seem to be expressible without loss in nonnormative terms. For instance, Kelly (2003) argues that the epistemic support relation cannot be identified with some relation of entailment. Similarly, though the rules of chess ground the chess standard, the fact that this standard guides us when playing chess isn't simply reducible to this grounding fact. These issues are contentious, of course. Our point is just that, *prima facie*, the relation of support is a normative relation. It is just that it is not an *authoritatively* normative relation as such.

Here is one further observation in support of the normativity of operative standards of correctness. Noncompliance with such standards often renders one liable to certain kinds of criticism.³² When you make

30. For related discussion, see Rosen 2001. Rosen assumes for the purposes of his brief discussion that all reasons are authoritative. Otherwise his discussion is congenial to our way of thinking.

31. Thanks to Benjamin Kiesewetter for useful comments here. Kiesewetter n.d. makes use of similar considerations to also argue that epistemic reasons are normative, though he doesn't accept that there are normative reasons that are not authoritative. See also Joyce 2001 on "institutional reasons" for further discussion.

32. The notion of criticism-liability for noncompliance is not identical with the mere fact of noncompliance with an operative standard. It involves the further claim that criticism of some kind would be correct. We assume that criticism for noncompliance is not merely asserting that one has failed to comply. Of course, it doesn't follow from the fact that you are criticism-labile that anyone just plain ought to criticize you, or indeed, that anyone has any authoritative reason to criticize you. Perhaps criticism will only make you worse; or perhaps someone is set to blow up the Eiffel tower just in case anyone

an obviously incorrect move in chess, perhaps because you were quite properly preoccupied with something else, you are liable to a chess-specific kind of criticism. If you reasonably succumb to a threat and give a student an incorrect grade, you are liable to professional criticism. If you steal from the rich to give to the poor, you are liable to legal criticism. Plausibly enough, if you fail to respond correctly to your evidence, you are liable to a kind of epistemic criticism (for detailed discussion of these claims, see Woods 2016, 2018).

It is crucial to this further appeal to criticism liability that the standards in question are operative. What makes a standard operative? It is natural to think there will be some kind of social acceptance condition. It is also possible that there is some further evaluative or normative standard, for instance one that rules out the possibility of normative reasons yielded by the practices of undetectable torture. One of us is more attracted to some such further condition than the other.³³ We set these questions aside for discussion elsewhere.

We have been using Stockfish standards as a convenient example of the standards of good chess. But, in fact, the Stockfish standards are unlikely to be operative when two beginners are playing each other. The correct move for an expert is not necessarily the correct move for a beginner. Likewise, the operative epistemic standards bearing on beliefs about whether it will rain this afternoon are different in my living room and the local meteorological office. On some ways of spelling out criticism-liability, whichever conditions explain why some particular set of standards is operative also explain why criticism for noncompliance with the standards is correct (for further discussion of the relation between normative standards, operativeness, and criticism-liability, see Brennan et al. 2013; Darwall 2006; Joyce 2001; Gilbert 1989; Bicchieri 2005; Woods 2018).

That's enough by way of introducing the properties of authoritative and nonauthoritative normative standards. It remains to ask: *which* reasons are the authoritative reasons?

We conjecture that all and only the *practical* reasons are the authoritative reasons.³⁴ Which ones are those? As noted above, we are neutral

criticizes you. The correctness of criticism doesn't entail authoritative reasons any more than the correctness of anything else. Thanks to Austen McDougal for discussion.

33. Woods (2018: sec. 10.3.1) argues against an evaluative condition. See Nieswandt 2019 for arguments for such a condition.

34. We are neutral concerning the modal status of the thesis that all and only practical reasons are authoritative reasons, whether a definitional claim about what it is to be

concerning the question of which reasons are the practical reasons. The structure of our view is available to those who accept different substantive accounts of practical reasons, whether value-based, desire-based, welfare-based, or reasons-fundamentalist. The cases mentioned at the outset of this article illustrate the kinds of reasons we have in mind. If the fact that some option would make Mildred happy is a practical reason to take that option, then that fact is an authoritative reason to take that option. Correctness-based reasons for actions, like that this chess move would capture material, are not practical reasons in this sense.³⁵

In summary, here is the simplest version of the picture. Start with the property of being what you just plain ought to do. What you just plain ought to do is immediately explained by the interaction of reasons for and against. These are the authoritative reasons. Our conjecture is that these are all practical reasons. Some of these practical reasons bear for or against one's participation in activities with standards of correctness. This is familiar enough from reflection on the normative status of games and so on. When there is most authoritative reason to be in a state that is correct according to some specific activity, then you ought to be in the state that is correct according to that activity.³⁶ This is a claim about meeting some activity-specific condition. These remarks also apply to beliefs. When there is most authoritative reason—hence most practical reason—to be in the correct doxastic state, then you ought to be in whatever doxastic state is correct—and not otherwise.

This picture allows us to absorb the intuitions we started with in a natural way. Even though you've done what you just plain ought to do when believing against the evidence, you've violated the relevant epistemic standards for belief, and, as such, are liable to epistemic criticism. It's this sense in which the anti-pragmatists are onto something. But this doesn't mean that you are failing to do what you ought to do in believing against the evidence, since in this case you lacked sufficient authoritative reason to adhere to the operative epistemic standards.

a just plain ought-making reason, or as a metaphysically substantive conjecture, or something else.

35. Is this approach inconsistent with the conjecture that practical reasons are "internal" reasons in the "game of life"? Discuss.

36. For simplicity, we are assuming that most authoritative reason entails just plain ought. We don't suppose that all instances of the schema—most reason entails ought—hold (for example, it seems to not hold for epistemic reasons).

5. Comparisons with Related Views

At this point, it will be helpful to contrast our view with some important alternatives and to respond to some objections. It's not our intention to refute alternative positions here. Our argument for our account is as advertised above, namely that it constitutes an account of the interaction of practical and epistemic sources of support for beliefs in a way that explains a range of cases, and that promises to generalize to the interaction of "right" and "wrong" reasons elsewhere in practical philosophy. Contrasting our view with others provides further clarification about the details of our own proposal, and further highlights its distinctiveness.

5.1. *Danielsson and Olson and the Direction of Explanation*

Danielsson and Olson (2007) distinguish between "content reasons" and "holding reasons." Content reasons for a state are considerations that bear on the standard of correctness for the state. These are a subset of our correctness-based reasons. In particular, they are correctness-based reasons bearing on *attitudes* with correctness conditions. Holding reasons for a state play a similar role to our practical reasons. Holding reasons are similarly more restrictive than our practical reasons, since they appear to only support attitudes. Our practical reasons are thus, plausibly, a generalization of holding reasons.

Danielsson and Olson do maintain—as a plausible normative hypothesis, not a definitional matter—that content reasons for an attitude imply a defeasible holding reason for the attitude (in our terminology, that correctness-based reasons imply practical reasons). The defeasibility leaves some room for these holding reasons to be defeated and outweighed, for instance in a case in which considerations of clutter avoidance support having no doxastic attitude toward *p*. But it does not, by itself, explain *why* considerations such as clutter avoidance can play this role, nor how this might pattern with other cases of practical influence on what one ought to believe.

Since this is a normative hypothesis rather than a definitional relation, their view is also consistent with denying that all content reasons entail holding reasons. This is an attractive feature of their view. But it would be helpful to have an explanation of why this defeasible thesis is true. Some "fittingness-firsters" have attempted to offer such an explanation in terms of the fact that correct states are finally valuable (Howard 2019a). We do not find this axiological claim plausible. But if it were true, this would only explain (together with some principle connecting value

with practical reasons) why the fact that a state is *correct* entails a practical reason to be in that state. It doesn't follow that every correctness-based reason to be in that state entails a practical reason to be in that state. Suppose a scientist tells you the earth is flat. That might constitute a weak epistemic reason; it doesn't plausibly entail a practical reason to believe that the earth is flat—even if correct beliefs are finally valuable.

But Danielsson and Olson go on to defend the following definitional claim, that for there to be a holding reason for A is for there to be a content reason to favor A. For them, content reasons are authoritatively normative. Holding reasons share this status, but only because to be a holding reason just is to be a content reason (for favor).

In contrast, we deny that content reasons are authoritatively normative. At most, facts about correctness are *conditionally* authoritatively normative—conditional on the bearing of practical reasons in particular situations. Epistemic reasons explain which beliefs are correct. So, insofar as their status is derivative from the status of that which they explain, epistemic reasons are, at most, derivatively conditionally authoritatively normative. On our view, Danielsson and Olson's holding reasons—identified by us as the practical reasons—are the only authoritatively normative reasons. Not, as they have it, the content reasons—our correctness-based reasons. We think that this alternative account of what we ought to believe fits better with our intuitions about cases and their explanation.

5.2. *Andrew Reisner and Selim Berker on Multilevel Weighing*

The closest thing to a discussion of a multilevel view like ours that we are aware of is due to Andrew Reisner. Here is one version of this view (Reisner 2008: 24): “When [practical] reasons for belief are strong enough, [epistemic] reasons for belief are silent, and . . . otherwise, [practical] reasons for belief are silent in determining what one ought to believe, all-things-considered [that is, what one just plain ought to believe].”

Our view shares many important features with Reisner's view. We both accept the autonomous domain of epistemic reasons and a distinctively epistemic standard of correctness. We both accept an independent deontic standard—Reisner's “ought, all things considered” and our just plain ought. We both maintain that in many cases, you ought to have the belief that is correct. We both maintain that in some cases, practical reasons explain the fact that it is not the case that you ought to have

the correct belief. And we both maintain that in such cases, the facts about what you ought to believe are explained by a complex of an autonomous epistemic explanation and some practical reasons bearing on whether to have the correct belief (or perhaps: whether to have the belief that is correct).

Selim Berker (forthcoming) objects that Reisner's view yields the wrong verdicts in a number of cases. His simplest runs as follows. Suppose one faces the question whether to believe, disbelieve, or suspend belief in some proposition p . Assume that there are practical reasons against being in one of these three states—let's say, against believing that p —that are "strong enough" to "silence" the relevant epistemic reasons. Assume that no practical reasons bear positively or negatively on the other two states. Assume that if practical reasons do not distinguish between two options, both are permissible. It would seem to follow from Reisner's view that both suspending belief in p and disbelieving p are permissible. But suppose that the evidence is equally balanced between favoring and disfavoring p . Berker argues that it would not be permissible to disbelieve p in such a case—one must suspend. He concludes that one of the assumptions attributed to Reisner must be false. We agree (though our view comes under fire from some of Berker's objections as well; more below).

Our view avoids this problematic result and offers a satisfying way to explain what you ought to do in many cases with similar structure. Start with the analogy with actions in activities. Suppose that you are incentivized not to play knight to d4, which in this situation is the best move available to you. But suppose you ought to keep playing competitively. How should you move? Assuming everything else about the situation is normal, you should make the *second-best* move.³⁷ If the second-best move is pawn to f6, you should play pawn to f6.

Suppose that (1) you have a practical reason to have a belief concerning whether p that is as accurate as possible, (2) you have a weightier practical reason against believing that p , namely that you promised to not believe that p , and yet (3) the epistemic reasons support believing p . Let's say we are working with the three doxastic states—belief, disbelief, and suspension. The next most accurate attitude is suspension rather than disbelief. So, the practical reason for accuracy will also provide support for suspension as against disbelief. In this case, you ought to suspend judgment concerning p .

37. Our thanks to Chris Howard for discussions here. For an alternative solution that aims to preserve the authoritative normativity of fittingness, see Howard (2019b).

But now suppose that (4) your practical reason against believing that p is that it would be very upsetting. Perhaps p is “My partner is unfaithful.” You have a weighty practical reason against believing this, plausibly enough. But, equally plausibly, this reason also weighs against suspending judgment in p . So now we have the accuracy-based reason (1) that, given the correctness facts in (3), favors believing that p . You have a weighty reason against believing that p . The next best you can do, epistemically speaking, is to suspend judgment in p . But you have a weighty reason against doing that too. This reason is perhaps not quite as weighty as the reason against believing that p (this will depend on how it will impact you to suspend belief in this matter, rather than believe that your partner is faithful). But this belief against suspension may still be weighty enough to outweigh the accuracy-based reason in favor of believing your partner is unfaithful. In this case, you ought to disbelieve p .³⁸

Reisner’s view can also be interpreted in a slightly different way, as Berker explains:

On this second interpretation, overall all-things-considered verdicts can be determined by two successive processes of weighing. First we weigh the epistemic reasons for and against belief and its alternatives in the standard epistemic-reasons-for-belief-like way. . . . Then we weigh the practical reasons for and against belief and its alternatives in a standard practical-reasons-for-action-like way. . . . with one exception: if the first weighing process resulted in the subject having decisive epistemic reason to hold a given doxastic attitude, then we include an additional reason of a fixed, very high weight in favor of that attitude in the second weighing process. (Berker, forthcoming; compare Reisner 2008: 26–27)

There are some crucial differences between this view and our own. Like our own, it views evidential reasons as not authoritative. However, it holds that facts about the balance of epistemic reasons necessarily entail authoritative reasons (though perhaps only if some further epistemic condition is met, such as the weight of evidence strongly supporting the relevant doxastic state). This entailment goes through independently of any facts about nonepistemic reasons.

38. There are some tricky questions about how to characterize “second-best” in epistemology (for some discussion, see Staffel 2015). But our view still offers some guidance. Suppose the evidence supports a credence of exactly 0.7. Appeal to practical standards will explain which doxastic states are to be avoided (see the example about unfaithfulness in the main text). The answer will be, roughly: get close to 0.7 while avoiding the disincentive.

We deny that purely epistemic considerations entail authoritative reasons, whatever their overall strength. A fortiori, we deny that really strong epistemic reasons will win the day over weak practical reasons. This difference is clearest in the cases of mental clutter and evil beliefs; we see these cases as ones where we have no, or just about no, practical reason to have a particular doxastic state.

Perhaps some further terminology will help to clarify this distinction. We defined evidentialism as the thesis that there are distinctively epistemic reasons for belief. We accept this thesis. We defined anti-pragmatism as the thesis that only epistemic reasons are reasons for belief. We reject this. Since we have introduced the further property of authoritativeness, we can define the following two theses in the ballpark of evidentialism. *Authoritative epistemic sufficiency* maintains that all epistemic reasons for belief are authoritative. *Authoritative epistemic necessity* maintains that no nonepistemic reasons for belief are authoritative.³⁹ The first Reisner view, so far as we can tell, accepts authoritative epistemic sufficiency while rejecting authoritative epistemic necessity. The second Reisner view accepts a version of authoritative epistemic sufficiency, restricted to cases involving decisive epistemic reason, while rejecting authoritative epistemic necessity. We reject all of these theses.

We offer a more natural, and perhaps less theoretical, explanation of the motivating intuition: when practical reasons against having a correct belief are weak, only very weak practical reasons in favor of having the correct belief are needed to carry the day; such reasons are incredibly easy to come by, given the important role epistemic considerations play in our workaday lives. As there are often strong practical reasons to have correct beliefs, and only weak practical reasons not to, this view and ours will look rather similar in many cases, but the difference is theoretically and explanatorily important.

It would be disingenuous to stop there, however, since Berker also develops significant worries for Reisner's second view that affect our view as well. Suppose we have a picture like ours where there's a range of practical reasons in play. Suppose, in particular, we can have a strong practical reason *against* believing *p*, a mild practical reason against suspending belief in *p*, and a mild epistemic reason to believe correctly. Then, as Berker notes, by increasing evidence in *p* without changing

39. Notice that so long as we presume that practical reasons to have correct beliefs transmit to practical reasons to have the beliefs that are correct, this thesis also entails that there are no practical reasons to have correct beliefs.

the practical reasons, we can go from being permitted to disbelieve p , to being required to suspend belief in p , to being permitted to disbelieve p again. As Berker points out, this violates the following principle, that increase in evidence *for* p should never make it the case that we go from being forbidden to disbelieve in p to being permitted to disbelieve in p . Though our view isn't quite committed to this result—there could be authoritative dilemmas for all we've said—it's nevertheless intuitive and deserves addressing.

First a caveat: such cases will be rare. Epistemic considerations often affect the balance of practical reasons through our reasons to be as correct as possible and to avoid epistemically bad doxastic states. However, Berker is right that for any reasonable distribution of practical reasons, we can find a case like his in the abstract. So cases that violate Berker's principle need to be addressed.

Our response is to deny the intuitiveness of the principle as a principle governing the *authoritative* ought, while granting fully it as a principle of the *epistemic* ought. We suspect that much of its intuitiveness comes from the difficulty in disentangling epistemic and authoritative standards, especially given how frequently they coincide. As we address below when talking about followability (sec. 5.5), when we engage with the question of whether to believe p , we frequently do so from an epistemic point of view, and from that point of view, we are not permitted to disbelieve p . And, as mentioned above, we're still liable to criticism for violating our epistemic obligations even though we had sufficient authoritative reason to do so.

All these facts combine in undermining the intuitiveness of Berker's principle (see also the discussion of Kelly cases in sec. 5.4). We do admit that our view can't rescue this principle as such; there will potentially be violations of it. But we also hold that once our view is on the table and, in particular, once we focus on the fact that the epistemic standard is still in force, much of the strangeness of these violations goes away.⁴⁰

5.3. *Susanna Rinard and the Normativity of Evidence*

Our view also shares a number of features with the *Robust Pragmatism* of Susanna Rinard. Rinard emphasizes the significance of an overall deontic standard applicable to both beliefs and actions. In Rinard (2015), the

40. Thanks to Selim Berker for helpful clarification of his view and very useful discussion of our response.

“guidance-giving should” plays a similar functional role to our “just plain ought.” Rinard emphasizes the “equal treatment” of belief and action: that what one just plain ought to believe is explained in just the same way as what one just plain ought to do—that there’s no difference in the ultimate explanatory grounds between the two cases (see 2017, 2019). We are happy to go along with this too, with the important caveat that the *mediate explanatory structure* in the two cases is often different (as described above). And like Rinard, our claim that only practical reasons are authoritative is object-neutral: it applies as much to beliefs as to actions.

To draw out one difference, consider the following:

In most ordinary cases, evidence in favor of P constitutes a [practical] reason to believe it. Typically, evidence that the store is closed now is a pragmatic consideration in favour of believing it, as one would (typically) be inconvenienced by having false beliefs about the store’s hours. Evidence that one’s spouse has pneumonia is (typically) a pragmatic reason to believe it, as one will (ordinarily) be better suited to care for them if one has true beliefs about the nature of their illness. (Rinard 2015: 219)

Elsewhere in that article she says: “Robust Pragmatism does not entail that epistemic considerations are never reasons for belief. On the contrary . . . it is *typically* the case, on Robust Pragmatism, that evidence for P is a reason to believe P. This is because it is typically the case that evidence for P is a pragmatic consideration in favor of believing P” (218). But this is not quite right. An epistemic reason for *p* is never, strictly speaking, a practical reason to believe *p*. It is rather a correctness-based reason to believe *p*. Take the fact, F, that Jones told you that *p* in some context where it’s important to have an accurate belief concerning *p*. Is F a practical reason to believe that *p*? No, it is still just an epistemic reason to believe *p*. If, for instance, it’s painful to believe that *p*, this does not weigh against F; it rather weighs against the practical reasons to have an accurate belief concerning *p*.

In contrast, if Smith told you that everything Jones told you is false, that’s an epistemic reason to not believe *p*, which in turns weighs against F. Practical reasons weigh against other practical reasons and epistemic reasons weigh against epistemic reasons, but epistemic and practical reasons never weigh against each other directly. This isn’t mere pedantry. It is important that epistemic considerations do normative work *in an epistemic domain* by the light of distinctively epistemic standards. This is precisely what our account recommends.

In more recent work (2019), Rinard has expressed things slightly differently in response to just this consideration. She says (2019: 1930): “Evidence for P usually does provide one with a reason to believe it. . . . Since we’re generally better off believing the truth, and since evidence for P indicates that P is true, evidence for P generally provides us with a reason to believe P.” Here evidence “provides” a practical reason rather than constituting one. In that article, Rinard also allows for the possibility of an independent epistemic normative standard, so long as it does not relate constitutively to the “guidance-giving should.” This would fit nicely with our insistence upon an autonomous domain of epistemic normativity, yielding independent reasons, weights, and criticism-liability.

There remains a difference of emphasis, of course. In that same work, Rinard “suspects” that epistemic normativity is merely a “philosopher’s invention.”⁴¹ As can be seen from above, we have no such suspicion; quite the contrary. But our considered views are largely consistent. In personal correspondence, Rinard agrees, and suggests one residual disagreement (Susanna Rinard, pers. comm.):

We share some central theses—in many ways I think our views are animated by the same spirit. . . . However, if I were going to look for a difference, it would be in your claims that epistemic normativity is a real normative standard—that if one believes incorrectly, in your sense, then one is, as you say, criticizable, one has erred, etc., even if one believes as they ought.

We are more compelled by the idea that one might be liable to criticism even if one did what one ought to have done. This is a familiar idea in practical ethics, where many agree that it might be the case that one ought to do something that involves failing to adhere to an operative normative standard—perhaps one ought to break a promise, or disappoint a friend, or cheat on an exam to appease the terrorist.

Take a case in which one has most reason to break some professional code of conduct. One remains criticism-liable in such cases. It would be perfectly correct for another participant in the practice—one’s manager, or another employee, for instance—to criticize you for your action. It may even be that compliance was supported by very weighty practical reasons that just happened to be outweighed by even weightier ones. Nothing we have said about practical reasons rules out the possibility of serious practical conflicts. All the same, plausibly enough,

41. For similar skepticism, see Papineau (2013), and Glüer and Wikforss (2013).

you ought not to comply with the outweighed standards in some such cases. Just so, we think that one might be liable to criticism for failing to adhere to an operative epistemic standard, even if one ought not to. We don't insist that one is always so liable to criticism; there may be cases and cases. We just want to make theoretical room for the possibility in at least some cases.

In this sense, our account preserves some normative significance for the standard of correctness for beliefs even in cases in which one should not have the correct belief. We can accommodate the intuition that one errs, in the relevant respect—even if one is nevertheless doing what one ought to do.⁴²

5.4. *Thomas Kelly on Salient Evidence*

Thomas Kelly (2003) introduces some cases that are notoriously difficult for pragmatist views to explain. Our view gives the pragmatist significantly more to say. Consider:

KELLY CASES: You pass me on your way to see the *whodunit*, but I think you've already seen it. I mention how surprised and delighted I was that the butler actually did it in the end. You are epistemically justified in being *very confident* that the relevant proposition is true, even though it undermines your ability to enjoy the movie. Nothing good would come from your believing this.

These cases show that practical reasons for belief can't undermine epistemic reasons for belief. We accept this with some important reservations (recall the epidemiologist case above, for example). They show, furthermore, that the weight of some epistemic reason to believe that *p* is unaffected by the weight of practical considerations bearing on your doxastic attitude toward *p*. Again, with some reservations (recall the "testimonial affirmative action" case above), we accept this claim, on the whole.⁴³ We also accept that one would be criticism-liable by failure to comply with the relevant epistemic standard in such a case.

But, in the Kelly cases, it remains plausible that you have no practical reasons to believe that the butler did it. So, we deny that you just plain ought to do so.

42. Many thanks to Susanna Rinard for helpful discussion of these issues.

43. These two qualifications, incidentally, do not violate the autonomy of the epistemic explanation. The autonomous explanation proceeds once the relevant conditions and modifiers are in place. This is familiar from the practical domain. See Bader (2015).

Our view attempts to split the intuitions in Kelly cases—Kelly is absolutely right that the practical reasons for belief don't influence the bearing of the epistemic reasons on what it would be correct to believe. These epistemic reasons are normative and the standard is operative. You would, plausibly enough, be liable to criticism for noncompliance. But on our view, there is more to the story since epistemic reasons aren't authoritative. So it's not true that we just plain ought to believe in line with someone's slip of evidence in a Kelly case; we'd just be epistemically correct in doing so.

One of the quirks of Kelly's example is that it is hard to see how one might *avoid* having the relevant belief given such salient decisive evidence. And also, the reason against having the belief (spoilage of the movie) is relatively inconsequential. Consider other cases involving salient evidence that might help to clarify the issues here. Suppose you don't want to know the sex of your baby until it is born. Suppose some nurse, not knowing this, gives you some evidence one way or another. It would behoove you to attempt to ignore this evidence, perhaps even to attempt to discredit it in your own mind (if that helps). These are actions, some of them mental actions. Our point is not about whether you ought to perform these. Our point is that if, somehow, you managed to avoid having the relevant belief, you would not thereby fail to have the beliefs you (just plain) ought to have—no matter how compelling the relevant evidence.

Or suppose you know yourself to be overly sensitive to criticism. You are giving a public presentation, live on television. It is very important that you give it your best performance—even if that is not great. Certain important members of the audience right in front of you are yawning, clearly bored. This is good salient evidence that your talk is going badly. Still, if you form the belief that your talk is going badly, you will not adapt, perhaps throwing in a few jokes to liven the crowd. You will fail utterly. You know all this. It seems quite plausible that you ought not believe the talk is going badly.⁴⁴

An analogy with emotions might also help. Suppose you have an interview for your dream job. It is psychologically plausible that having the correct emotions—excitement about the prospect, anxiety about how you will perform, fear of failure—would hamper your performance. It seems plausible enough that there would be no value in fearing failure

44. This is adapted from Kornblith (1983); see also Greaves (2013).

during the interview, that nothing apart from the correctness of such fear speaks in favor of the emotion. So, plausibly, you have no *practical* reason to fear failure during the interview.⁴⁵ And yet, given the circumstances, it would be perfectly understandable if you had that response and in so doing you'd have, as it were, fear-correctly responded to your situation.

5.5. *On Followability*

One of the most common objections to pragmatism, mentioned at the outset, denies that one can take a nonepistemic consideration as a reason to believe something, or that one can reason from a nonepistemic consideration to a belief. It is further maintained that a consideration is a reason for some state only if one can so respond to it, and hence that nonepistemic considerations are not reasons for beliefs. There is an objection to our view here since we maintain that there are nonepistemic reasons for belief—and indeed, worse, that in an important sense *all* reasons for belief are nonepistemic reasons for belief.

A number of theorists have recently put pressure on the followability argument, pointing out that instances of *bad* reasoning suggest that the first premise is better construed as a claim about legitimacy (or something along those lines) rather than possibility (Way 2016), and also pointing out some cases in which one can take nonepistemic considerations as reasons for belief in conjunction with other epistemic reasons for belief (Leary 2017), and that, even in the case of reasons for action, it isn't clear that all reasons are "followable" (Rinard 2015). We find these replies plausible and accept that the "followability" premise needs more work.

But we want to emphasize an alternative line of response. We affirm the existence of epistemic reasons, and we accept their distinctiveness. Epistemic and nonepistemic reasons have different functional roles in our overall account. It is an open question, so far as anything we have said is concerned, whether all and only the reasons playing the functional role that epistemic reasons play are followable in the relevant sense.

It is important to emphasize that we are not engaged in the project of giving an account of *reasoning* here. We are offering an account of what explains what you ought to believe. We are distinguishing between how one does or should reason about what to believe, on the one hand, and

45. This example presumes that the "right kind of reasons" for emotions are not authoritative reasons. For discussion, see Maguire (2018) and Howard (2019a).

the full explanation of what one ought to believe, on the other hand. The fact—if it is a fact—that we should not reason from practical considerations to facts about what to believe does not entail that practical considerations don't play any role in the metaphysical explanation of what you ought to believe. Neither does it follow that one ought to reason from a consideration if and/or only if that consideration is an *authoritative* reason.

Another idea from Rawls is suggestive here. Rawlsian practices, which are presumably one type of activity in our sense, often involve autonomous ways of reasoning about moves within the practice. For example, Rawls (1955: 15) remarked that “a general utilitarian defense is not open to the promisor: it is not one of the defences allowed by the practice of making promises.” The same might be true of believing. We can distinguish the distinctively epistemic activity of responding to one's evidence in conformity to the operative rules in a context, from a more general activity of figuring out what to do or how to be. It might be that while actively engaged in the epistemic activity, it is only correct, by the standards of that activity, to cite evidence as support for beliefs, even if we can cite practical reasons, as it were, from the outside. This is further evidence that these correctness-based reasons are reasons and that they are normative—given the standard “earmarks” of normative reasons. This fits with our antirevisionist methodology.

This distinction is structurally similar to the more familiar distinction between “criteria of rightness” and “decision-making procedures” in practical normativity. Railton's (1984) “sophisticated consequentialist” will in ordinary contexts take as reasons for action facts about one's friend's interests, or one's promises; one should not take as one's reason to see one's friend in the hospital “that doing so will maximize value.” The pressure to resist this plausibly comes from the practice of friendship, which imposes constitutive constraints on one's motives. Nevertheless, it remains the position of the sophisticated consequentialist that one's authoritative reasons (to put things in our terms) are given by facts about what values one's actions promote. This presents the possibility that in responding to the correctness-based reasons within the practice of friendship, one will fail to conform to one's authoritative reasons, for instance in a case in which one really ought to skip one's friend's party to go phone-banking. Just so, in responding to correctness-based reasons within the activity of believing, one might fail to conform to one's authoritative reasons.

The analogy with affective attitudes is helpful again here. The fact that fear would be costly does not affect the correctness conditions for fear (in the face of a marauding bear). And yet, you may respond, in some sense, to both the fact of the marauding bear and the fact that fear would be costly by not being afraid. That seems psychologically possible, and not obviously robustly criticizable, in just the way it would not be robustly criticizable to fail to update in a Kelly case, even though it would be epistemically criticizable.

Suppose this is right. It is not impossible to respond to practical reasons for belief. It is rather that one should not follow practical reasons for belief—where this “should” is given by the activity of believing. But notice that even this thought is compatible with responding to practical reasons on some other level (just as one might respond to the possible consequences of fear, or delight, by remaining stoical—and just as one might, against the norms of the practice, take consequentialist considerations into account in deciding whether to keep one’s promise).

The following stretch of reasoning seems unobjectionable to us: “Evidence E seems to support believing that p ; but I don’t believe that p , it would be awful to believe that p , and I don’t see any authoritative reason to be properly responsive to the evidence in this matter, so I’m just going to ignore that evidence.” This is wishful thinking, to be sure. But not all cases of wishful thinking are contrary to (authoritative) reason.

Suppose you believe that the fact that your child is not home on time is evidence that they are in danger and that it would be really terrible if your child were in danger; in response to these beliefs, you believe that they are not in danger. That’s a paradigmatically bad case of wishful thinking. But now consider an unpleasant variation on the case. Your child was killed by an unknown attacker. You have evidence that your child suffered before they died; this evidence is just as strong as the evidence that your child was in danger in the original version of the case. But it would be really terrible if your child suffered before they died; plausibly, there are no practical reasons to have this belief, and weighty practical reasons against. We think it is plausible that in this situation, it is not the case that you just plain ought to believe that your child suffered.

You may disagree. Some people prefer to know the worst rather than fear the worst. But others might not. And others should not. And that is all we need, so long as the difference turns on nonepistemic considerations, namely how the belief would affect you. Suppose you failed to believe your child was in danger but, if you so believed, you might be able to do something to help. Plausibly, you would be acting

contrary to (authoritative) reason if you were to fail so to believe. In that case, what you would be criticized for is giving your own peace of mind too much weight relative to the chance that you might be able to do something to help your child. This is precisely what our account would predict, since relevant normative considerations are not epistemic but practical.

6. Conclusion

We have been arguing that it is consistent and attractive to hold that there is an autonomous domain of epistemic reasons that provide support for an independent standard of epistemic correctness. We have suggested they play an important role in many cases in determining what you ought to believe, while also maintaining that there are practical reasons for beliefs, and, indeed, that all authoritative reasons for belief are practical reasons.

Our account will not satisfy everyone. We offer a better response to Kelly cases than hardcore pragmatists. But strong evidentialists may still find our account too weak. Our lack of hostility to certain cases of wishful thinking will not appeal to a certain respectable philosophical turn of mind. And some potential implications are downright startling on first glance. For instance, if you are offered enough money to believe that p , our account suggests that you will be contrary to (authoritative) reason if you fail to believe that p , for any p ; or for another example, the grieving parent would seem to be contrary to authoritative reason if they did believe their child suffered.⁴⁶ These results are particularly worrisome given our hermeneutic methodology. The significance of our line on this sort of case is vulnerable to future work on “ought implies can” and doxastic voluntarism, and future work on related issues in other areas in normativity.

Nevertheless, we take our account to have a number of attractive features. Among the significant merits of this view is that we do not merely account for the existence of a range of cases supporting pragmatism. Our view predicts the many different ways in which practical considerations have influence in determining what we ought to believe—sometimes bearing directly on beliefs, sometimes on having correct beliefs, sometimes on the determination of evidence or the appropriate epistemic

46. Only “suggests” for our view is consistent with further constraints on both epistemic reasons and authoritative reasons (for instance, consider the brief remarks about *being operative* above). Thanks to an anonymous referee for pressing us on these points.

rule, sometimes on our storage capacities. All these interactions are straightforward, so long as one carefully keeps track of what reasons are supporting, which reasons are transmitting, and which reasons are weighing against each other. According to our account, both epistemic reasons and practical reasons usually play important, and importantly, different roles in an overall explanation of what you ought to believe.

This proposal offers an initial framework to account for the interaction of both “right” and “wrong” kinds of reasons and “formal” and “substantive” normativity. This framework promises to bear fruit in other areas of practical philosophy, for instance, concerning the interaction of correctness-based and practical reasons for affective attitudes and, more generally, concerning reasons in the context of Rawlsian practices, professional roles, norms-governed relationships and projects, and social institutions.

Our account offers an attractive combination of normative pluralism (given the range of different domain-specific normative properties) and normative unity (since all and only practical reasons are authoritative). Our account is consistent with (but does not assume) the thesis that there is just one “source” of authoritative normativity. This fits nicely with our motivating picture: a single authoritative standard bears on every aspect of our orientation to each other and the world—not just on our actions, or our actions plus our motives, or our actions plus our motives plus our emotions, but also on our beliefs about ourselves and others.

References

- Archer, Avery. 2018. “Wondering about What You Know.” *Analysis* 78, no. 4: 596–604.
- Bader, Ralph. 2015. “Conditions, Modifiers, and Holism.” In *Weighing Reasons*, edited by Errol Lord and Barry Maguire, 27–55. Oxford: Oxford University Press.
- Basu, Rima. 2018. “Can Beliefs Wrong?” *Philosophical Topics* 46, no. 1: 1–17.
- Berker, Selim. 2018. “A Combinatorial Argument against Pragmatic Reasons for Belief.” *Analytic Philosophy* 59, no. 4: 427–70.
- Bicchieri, Cristina. 2005. *The Grammar of Society: The Nature and Dynamics of Social Norms*. Cambridge: Cambridge University Press.
- Brennan, Geoffrey, Lina Eriksson, Robert E. Goodin, and Nicholas Southwood. 2013. *Explaining Norms*. Oxford: Oxford University Press.
- Broome, John. 2013. *Rationality through Reasoning*. Oxford: Wiley-Blackwell.
- Buchak, Lara. 2013. *Risk and Rationality*. Oxford: Oxford University Press.

- Ceyhan, Özlem, Betül Özen, Handan Zincir, Nuray Şimşek, and Meral Başaran. 2018. "How Intensive Care Nurses Perceive Good Death." *Death Studies* 42, no. 10: 667–72.
- Chappell, Richard. 2012. "Fittingness: The Sole Normative Primitive." *Philosophical Quarterly* 62, no. 249: 684–704.
- Chignell, Andrew. 2016. "The Ethics of Belief." Stanford Encyclopedia of Philosophy Archive, October 4. plato.stanford.edu/archives/win2016/entries/ethics-belief/.
- Cohen, G. A. 2008. *Rescuing Justice, Rescuing Equality*. Cambridge, MA: Harvard University Press.
- Crawford, Lindsay. 2019. "Believing the Best: On Doxastic Partiality in Friendship." *Synthese* 196, no. 4: 1575–93.
- Danielsson, Sven, and Jonas Olson. 2007. "Brentano and the Buck-Passers." *Mind* 116, no. 463: 511–22.
- Darwall, Stephen L. 2006. *The Second-Person Standpoint: Morality, Respect, and Accountability*. Cambridge, MA: Harvard University Press.
- Feldman, Richard. 2000. "The Ethics of Belief." *Philosophy and Phenomenological Research* 60, no. 3: 667–95.
- Field, Hartry. 2000. "Apriority as an Evaluative Notion." In *New Essays on the A Priori*, edited by Paul Boghossian and Christopher Peacocke, 117–49. Oxford: Oxford University Press.
- Friedman, Jane. 2018. "Junk Beliefs and Interest-Driven Epistemology." *Philosophy and Phenomenological Research* 97, no. 3: 568–83.
- Gendler, Tamar. 2011. "On the Epistemic Costs of Implicit Bias." *Philosophical Studies* 156, no. 1: 33–63.
- Gertken, Jan, and Benjamin Kiesewetter. 2017. "The Right and the Wrong Kind of Reasons." *Philosophy Compass* 12, no. 5: e12412. doi.org/10.1111/phc3.12412.
- Gilbert, Margaret. 1989. *On Social Facts*. New York: Routledge.
- Glüer, Kathrin, and Åsa Wikforss. 2013. "Against Belief Normativity." In *The Aim of Belief*, edited by Timothy Chan, 80–99. Oxford: Oxford University Press.
- Greaves, Hilary. 2013. "Epistemic Decision Theory." *Mind* 122, no. 488: 915–52.
- Harman, Gilbert. 1986. *Change in View*. Cambridge, MA: MIT Press.
- Hazlett, Allan. 2013. *A Luxury of the Understanding: On the Value of True Belief*. Oxford: Oxford University Press.
- Hedden, Brian, and Mark Colyvan. 2019. "Legal Probabilism: A Qualified Defence." *Journal of Political Philosophy* 27, no. 4: 448–68.
- Howard, Chris. 2018. "Fittingness." *Philosophy Compass* 13, no. 11: e12542. doi.org/10.1111/phc3.12542.
- Howard, Chris. 2019a. "The Fundamentality of Fit." In *Oxford Studies in Metaethics*, vol. 14, edited by Russ Shafer-Landau, 216–36. Oxford: Oxford University Press.

- Howard, Chris. 2019b. "Weighing Epistemic and Practical Reasons for Beliefs." *Philosophical Studies*, June 12. doi.org/10.1007/s1109-8-019-01307-y.
- James, William. 1896. "The Will to Believe." *The New World* 5: 327–47.
- Joyce, Richard. 2001. *The Myth of Morality*. Cambridge: Cambridge University Press.
- Keller, Simon. 2004. "Friendship and Belief." *Philosophical Papers* 33, no. 3: 329–51.
- Kelly, Thomas. 2003. "Epistemic Rationality as Instrumental Rationality: A Critique." *Philosophy and Phenomenological Research* 66, no. 3: 612–40.
- Kiesewetter, Benjamin. n. d. "Are Epistemic Reasons Normative?" Unpublished manuscript.
- Kiesewetter, Benjamin. 2017. *The Normativity of Rationality*. Oxford: Oxford University Press.
- Kornblith, Hilary. 1983. "Justified Belief and Epistemically Responsible Action." *Philosophical Review* 92, no. 1: 33–48.
- Leary, Stephanie. 2017. "In Defense of Practical Reasons for Belief." *Australasian Journal of Philosophy* 95, no. 3: 529–42.
- Maguire, Barry. 2018. "There Are No Reasons for Affective Attitudes." *Mind* 127, no. 507: 779–805.
- Maguire, Barry, and Errol Lord. 2016. "An Opinionated Guide to the Weight of Reasons." In *Weighing Reasons*, edited by Errol Lord and Barry Maguire, 3–24. Oxford: Oxford University Press.
- Marušić, Berislav. 2013. "Promising against the Evidence." *Ethics* 123, no. 2: 292–317.
- McCormick, Miriam Schleifer. 2014. *Believing Against the Evidence: Agency and the Ethics of Belief*. New York: Routledge.
- McLeod, Owen. 2001. "Just Plain Ought." *Journal of Ethics* 5, no. 4: 269–91.
- McPherson, Tristram. 2018. "Authoritatively Normative Concepts." In *Oxford Studies in Metaethics*, vol. 13, edited by Russ Shafer-Landau, 253–77. Oxford: Oxford University Press.
- Nieswandt, Katharina. "What Is Conventionalism about Moral Rights and Duties?" *Australasian Journal of Philosophy* 97, no. 1: 15–28.
- Papineau, David. 2013. "There Are No Norms of Belief." In *The Aim of Belief*, edited by Timothy Chan, 64–79. Oxford: Oxford University Press.
- Parfit, Derek. 2011. *On What Matters*, vol. 1. Oxford: Oxford University Press.
- Pascal, Blaise. (1960) 1995. *Pensées*. Translated and edited by A. Kreilsheimer. New York: Penguin.
- Preston-Roedder, Ryan. 2013. "Faith in Humanity." *Philosophy and Phenomenological Research* 87, no. 3: 664–87.
- Rabinowicz, Wlodek, and Toni Rønnow-Rasmussen. 2004. "The Strike of the Demon: On Fitting Pro-attitudes and Value." *Ethics* 114, no. 3: 391–423.
- Railton, Peter. 1984. "Alienation, Consequentialism, and the Demands of Morality." *Philosophy and Public Affairs* 13, no. 2: 134–71.

- Rawls, John. 1955. "Two Concepts of Rules." *Philosophical Review* 64, no. 1: 3–32.
- Raz, Joseph. 2013. "Value: A Menu of Questions." In *Reason, Morality, and Law: The Philosophy of John Finnis*, edited by John Keown and Robert P. George, 13–23. Oxford: Oxford University Press.
- Reisner, Andrew. 2008. "Weighing Pragmatic and Epistemic Reasons for Belief." *Philosophical Studies* 138, no. 1: 17–27.
- Reisner, Andrew. 2018. "Two Theses about the Distinctness of Practical and Theoretical Normativity." In *Normativity: Epistemic and Practical*, edited by Conor McHugh, Jonathan Way, and Daniel Whiting, 221–40. Oxford: Oxford University Press.
- Rinard, Susanna. 2015. "Against the New Evidentialists." *Philosophical Issues* 25, no. 1: 208–23.
- Rinard, Susanna. 2017. "No Exception for Belief." *Philosophy and Phenomenological Research* 94, no. 1: 121–43.
- Rinard, Susanna. 2019. "Equal Treatment for Belief." *Philosophical Studies* 176, no. 7: 1923–50.
- Rosen, Gideon. 2001. "Brandom on Modality, Normativity, and Intentionality." *Philosophy and Phenomenological Research* 63, no. 3: 611–23.
- Scanlon, T. M. 1998. *What We Owe to Each Other*. Cambridge, MA: Harvard University Press.
- Schroeder, Mark. 2010. "Value and the Right Kind of Reasons." In *Oxford Studies in Metaethics*, vol. 5, edited by Russ Shafer-Landau, 25–55. Oxford: Oxford University Press.
- Schroeder, Mark. 2011. "Ought, Agents, and Actions." *Philosophical Review* 120, no. 1: 1–41.
- Schroeder, Mark. 2015. "What Makes Reasons Sufficient?" *American Philosophical Quarterly* 52, no. 2: 159–70.
- Shah, Nishi. 2006. "A New Argument for Evidentialism." *Philosophical Quarterly* 56, no. 225: 481–98.
- Sharadin, Nathaniel. 2013. "Schroeder on the Wrong Kind of Reason Problem for Attitudes." *Journal of Ethics and Social Philosophy* 7, no. 3. doi.org/10.26556/jesp.v7i3.
- Shirer, William L. 1991. *The Rise and Fall of the Third Reich: A History of Nazi Germany*. New York: Random House.
- Staffel, Julia. 2015. "Measuring the Overall Incoherence of Credence Functions." *Synthese* 192, no. 5: 1467–93.
- Steglich-Petersen, Asbjørn. 2009. "Weighing the Aim of Belief." *Philosophical Studies* 145, no. 3: 395–405.
- Steglich-Petersen, Asbjørn. 2011. "How to Be a Teleologist about Epistemic Reasons." In *Reasons for Belief*, edited by Asbjørn Steglich-Petersen and Andrew Reisner, 13–33. Cambridge: Cambridge University Press.
- Stroud, Sarah. 2006. "Epistemic Partiality in Friendship." *Ethics* 116, no. 3: 498–524.

The Game of Belief

- Talbot, Brian. 2014. "Truth Promoting Non-evidential Reasons for Belief." *Philosophical Studies* 168, no. 3: 599–618.
- Way, Jonathan. 2012. "Transmission and the Wrong Kind of Reason." *Ethics* 122, no. 3: 489–515.
- Way, Jonathan. 2016. "Two Arguments for Evidentialism." *Philosophical Quarterly* 66, no. 265: 805–18.
- Way, Jonathan, and Connor McHugh. 2016. "Fittingness First." *Ethics* 126, no. 2: 575–606.
- Williams, Neil W., and Joe Saunders. 2018. "Practical Grounds for Belief: Kant and James on Religion." *European Journal of Philosophy* 26, no. 4: 1269–82.
- Wodak, Daniel. 2018. "Mere Formalities: Normative Fictions and Normative Authority." *Canadian Journal of Philosophy* 49, no. 6: 828–50.
- Woods, Jack. 2016. "The Normative Force of Promising." *Oxford Studies in Normative Ethics*, vol. 6, edited by Mark Timmons, 77–101. Oxford: Oxford University Press.
- Woods, Jack. 2018. "The Authority of Formality." In *Oxford Studies in Metaethics*, vol. 13, edited by Russ Shafer-Landau, 207–29. Oxford: Oxford University Press.
- Wright, Crispin. 2004. "Warrant for Nothing (and Foundations for Free)?" *Aristotelian Society Supplementary Volume* 78, no. 1: 167–212.