Reasons and Belief

Daniel Fogal

Penultimate version - please cite published version where possible

I. Introduction

Much recent work in epistemology has concerned the relationship between the epistemic and the practical, with a particular focus on the question of how, if at all, practical considerations affect what we ought to believe. Two main positive accounts have been proposed: reasons pragmatism and pragmatic encroachment. Following Quanbeck and Worsnip (forthcoming), we can understand them as follows:

> **Reasons pragmatism**: practical (including moral) considerations can affect what we ought to believe by constituting distinctively practical (i.e., non-epistemic) reasons for or against belief.

> **Pragmatic encroachment:** practical considerations bear on what we ought to believe by affecting epistemic justification (e.g. how much justification is required to justifiably believe).

Although the debates over reasons pragmatism (henceforth 'pragmatism') and pragmatic encroachment (henceforth 'encroachment') have often been pursued separately, they have also been pursued together, including by Quanbeck and Worsnip.

Both debates center around intuitive judgments about cases together with various principles involving reasons (and/or justification) and belief, with many contributors in the pragmatism and —to a lesser extent—encroachment debates helping themselves to talk of belief without saying much about what exactly they have in mind, other than the minimum necessary to distinguish "full" or "outright" belief from its graded cousin, confidence or "partial" belief. There has also been a tendency to overlook potentially relevant distinctions among different kinds of reasons that have been drawn elsewhere. The goal of this paper is to argue that greater clarity and care concerning both reasons and belief is called for. Increased sensitivity to various subtleties will not only minimize the chances of unwittingly engaging in merely verbal disputes but also allow us to better navigate the pragmatism and encroachment debates.

The paper consists of two parts. In the first part (Section II), I explore a range foundational issues concerning belief. In particular, I address the need to distinguish (a) different belief-like states, (b) different kinds of doxastic control, and (c) different ways 'believe' might be used. In the second part (Section III), I apply some of the lessons learned to the debate(s) over pragmatism and encroachment, revisiting some of the standard examples animating them and highlighting what I take to be important grains of truth in the different views on offer. The view I ultimately favor is a constrained form of pragmatism, according to which practical and epistemic considerations both constitute genuine reasons for belief but don't usually combine or interact to determine what we ought to believe.

II. Preliminaries

A. The Nature of Belief

Let's start with the distinction between full and partial belief.[1] Although the distinction can be drawn in different ways, it's standardly assumed that full belief is a fundamentally *categorical* state —you either fully believe p or you don't.[2] In contrast, partial belief is fundamentally *graded*—one can be more or less confident, have a lot or very little confidence, and so on. (I say 'fundamentally' because while full beliefs can vary in confidence or strength and so can be aptly described in *both* categorical *and* graded terms, it's the former that's fundamental to full belief *qua* full belief.) Hence our first point of contrast:

> *Contrast #1*: full belief is categorical, while partial belief is graded.

Beyond this basic difference, controversy abounds. It'll nonetheless be helpful to consider some of the most commonly suggested points of contrast between full and partial belief. It's often claimed, for example, that full belief, along with disbelief and agnosticism, are settled states of commitment: whereas full belief in p involves commitment to the truth of p, disbelief involves commitment to the falsity (or at least non-truth) of p, and agnosticism is form of committed neutrality with respect to p.[3] Partial belief, on the other hand, is *not* a settled state of commitment: one's degree or level of confidence often varies, and no amount of confidence in p (at least short of the maximum) commits one to the truth of p. This is partly why partial belief, unlike full belief, is not naturally assessable as true or false, though it can be assessed for accuracy in a broader sense (i.e. one can be closer or further from the truth). Thus:

> *Contrast #2*: full belief in p is a settled state of commitment to the truth of p, while partial belief is not.

Another common, complementary way of distinguishing full and partial belief states is in terms of functional roles. It's often said, for example, that to fully believe p is (in part) to be disposed to "rely on" p—i.e. to be disposed to use p as a premise in reasoning, to take it for granted in acting and decision-making, to assert it, and so on.[4] To disbelieve p is in turn (in part) a matter of being disposed to rely on not-p, while agnosticism is a matter of being disposed to not rely on either p or not-p. A high degree of confidence in p, by contrast, doesn't bring with it the same sorts of dispositions to (not) rely on p, and hence is compatible with either full belief or agnosticism. Thus:

> *Contrast #3*: full belief in p involves being disposed to rely on p, while partial belief doesn't.

---

[1] For useful surveys of (and references concerning) the relationship between full and partial belief, see Jackson (2020) and Weisberg (2020).

[2] I'm ignoring the possibility of vagueness and indeterminacy.

[3] For recent development of this idea, along with references, see Singh (forthcoming).

[4] Cf. Ross and Schroeder (2014).

A related idea is that to fully believe p is (in part) to be disposed to treat the question of whether p as settled and be resistant to re-opening deliberation.[5] Merely having a high degree of confidence in p, by contrast, is fully compatible with inquiry and deliberation with respect to p. Thus:

> *Contrast #4*: full belief in p disposes one to treat the question of p as settled and be resistant to re-opening deliberation, while partial belief doesn't.

While Contrasts #2-#4 tell us a fair bit about full belief, they don't tell us much about partial belief other than what it is *not*. Perhaps the most prominent idea is that degrees or levels of confidence play an important role in serving as weights in cognitive processes whose parameters are "continuous" (e.g. Bayes Theorem, expected utility calculations), rather than "discrete" (e.g. inference to the best explanation, dominance reasoning), though representing degrees of confidence in numerically precise ways (e.g. as credences) is, at best, an idealization.[6] It has also been suggested that partial belief in p tracks or embodies one's assessment of the likelihood of p, and hence (unlike categorical attitudes) is sensitive to small changes in one's evidence, and that partial belief is useful in more cautious, careful, and effortful inquiry and deliberations, triggered by (e.g.) high stakes, salience of error, or the need to defend one's stance.[7]

What *unites* full and partial belief—that is, what makes them both states of *belief*—is that they play various roles that are characteristic of belief, together forming our internal map or model of the world.[8] There are two main kinds of roles. Following Soter (forthcoming), I'll call the first *appraisal*: belief states, whether partial or full, form and change in response to (what we take to be) evidence, and are assessable in terms of both accuracy and justification.[9] The second is *guidance*: belief states guide a range of psychological and behavioral processes, including reasoning, action, thought, planning, feelings of conviction, and so on.[10] Putting these points together: belief states characteristically form in response to (perceived) evidence and jointly constitute our "default cognitive background" (to use Bratman's phrase), interacting with other mental states to affect "our patterns of attention, memory, thought, motivation, judgment and inference, goal selection, action tendencies, etc., often automatically and without our direct oversight" (5).

Although there is disagreement over the relative importance of the appraisal vs. guidance roles when it comes to the nature of (full or partial) belief, as well as over which particular appraisal and/or guidance roles are most central, for present purposes we needn't settle this question.[11] That's because I'm mostly concerned with what is *characteristic* of (or otherwise closely

---

[5] Cf. Holton (2014) and Friedman (2019).

[6] Cf. Ross and Schroeder (2014).

[7] See Weisberg (2020) for discussion and references.

[8] This gloss is intended to be neutral concerning the correct metaphysics of mind, especially the debate between "psychofunctional representationalists" like Mandelbaum and Quilty-Dunn (2018) and "liberal dispositionalists" like Schwitzgebel (2002).

[9] For an empirically-informed defense of the evidence-responsiveness of beliefs, see Flores (forthcoming).

[10] In Ryle's terminology, beliefs correspond to "multi-track" dispositions, signifying "abilities, tendencies or pronenesses to do, not things of one unique kind, but things of lots of different kinds" (1949: 118).

[11] Soter (forthcoming) surveys the debate and provides extensive references.

associated with) belief rather than what is strictly speaking *constitutive* of it. Soter represents this "two-pronged" model of belief as follows:
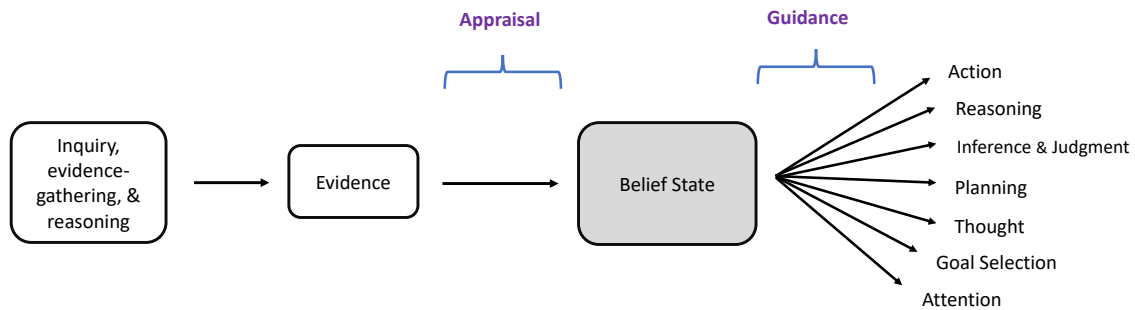

Figure 1 - The two-pronged architecture of belief

Note: "The processes stemming from guidance are meant to be representative but not exhaustive or precisely taxonomized, especially in their relations to each other (e.g., action often comes after, and as a result of, many of the other processes, which themselves can be broken into more precise mechanistic components and processes)."

The **first box** represents inquiry-related activities, understood as "practices of evidence-gathering, reasoning, and thinking through a matter: the 'upstream' practices that lead us to have [a given] body of evidence bearing on *p*" (4). The **second box** represents the body of evidence that results not only from inquiry, but also from whatever other, non-inquiry-related sources of evidence that may be relevant (e.g. unguided perception, memory, etc.). The **third box** represents one's belief state, whether partial or full.[12]

Whereas what unites full and partial belief states is *that* they play various appraisal and/or guidance roles, what distinguishes them (in part) is *how* they do so—i.e. the specific ways in which they are responsive to evidence and impact our patterns of reasoning, attention, thought, motivation, feelings of conviction, and the like. To illustrate, consider the (full) belief that there is beer in the fridge:

> Some of the dispositions associated with this belief include: the disposition to say, in appropriate circumstances, sentences like 'There is beer in my fridge'; the disposition to look in the fridge if one wants a beer; a readiness to offer beer to a thirsty guest; the

---

[12] Why depict a single 'belief state' box rather than two separate boxes, one for full belief and one for partial belief? Because that better reflects the storage of information. As Weisberg (2020) notes:

> Humans do not store information in two, separate belief structures, one for partial belief and the other for full belief. Full and partial belief share one, common store: long term memory. That single, shared store is used to generate both categorical judgments and levels of confidence. (18)

This is compatible with there being several different long-term memory stores, such as for episodic memory and for semantic memory (18). Regardless of the details, having a single, complex storage system functioning as a common "supervenience base" for both full and partial belief helps alleviate worries about double storage and/or conflict between them.

disposition to utter silently to oneself, in appropriate contexts, 'There is beer in my fridge'; an aptness to feel surprise should one go to the fridge and find no beer; the disposition to draw conclusions entailed by the proposition that there is beer in the fridge (e.g., that there is something in the fridge); and so forth. (Schwitzgebel 2002, 251)

In contrast, if one is merely more confident than not that there is beer in the fridge, some of the associated dispositions will include: the disposition to say in appropriate circumstances, things like 'There is probably beer in my fridge, but I'm not sure'; a reluctance to offer beer to a thirsty guest and instead first check the fridge; the disposition to utter silently to oneself, in appropriate contexts, 'There is probably beer in my fridge'; an aptness to feel only mild surprise should one go to the fridge and find no beer; and so on.

We can summarize the foregoing by saying that full and partial belief differ not only in their constitutive natures—e.g. full belief is a categorical form of commitment, while partial belief is not—but also in the specific ways in which they fulfill the various appraisal and guidance role(s) characteristic of belief. Although there is a debate over whether humans have both full and partial belief states, I take the case for *doxastic dualism* to be quite strong: human beings have both categorical doxastic attitudes (e.g. full belief, disbelief, agnosticism) and graded doxastic attitudes (e.g. partial belief), and neither fully reduces to the other.[13] Although there are a multitude of other important distinctions that might be drawn—such as between implicit and explicit beliefs (cf. Schwitzgebel 2023, §2.2), "merit" and "crony" beliefs (Simler 2016; Westra 2023), etc.—it's the distinction between categorical and graded belief states that matters most for my purposes.[14]

B. Doxastic Control

There is a large, sprawling literature on doxastic control, and in particular on whether beliefs are under our *voluntary* control. (By 'control' I'll mean *voluntary* control—unless otherwise noted—in

---

[13] For philosophical defense, see Ross and Schroeder (2014); for empirical defense, see Weisberg (2020).

[14] The introduction of additional, cross-cutting distinctions motivates *doxastic pluralism*, according to which there are a range of related but distinct categorical and graded doxastic states we might be in, none of which are metaphysically privileged over all others. The case for pluralism stems from taking seriously the fact that most if not all of the dispositions associated with the appraisal and guidance roles characteristic of belief are (a) graded and (b) can come apart from each other. In cases where they come apart and/or hold to different degrees—as they do in cases of explicit vs. implicit beliefs, merit vs. crony beliefs, ordinary vs. delusional beliefs, human vs. hon-human animal beliefs, and so on—the question then arises: does the subject "really" believe? At least in many such cases, one is likely to feel conflicted, and it's natural to reply with something like "well, in one sense they do, and in another sense they don't" (cf. Schwitzgebel 2002). What's more, many of the seemingly binary doxastic distinctions (implicit/explicit, merit/crony, etc) really represent a spectrum, with some belief states falling between the extremes (and hence being more or less implicit/explicit, etc.), varying in terms of the fine(r)-grained appraisal and guidance dispositions associated with them. Although there are different ways one might try to regiment or precisify the notion of belief, with some "clusters" of finer-grained, appraisal- and/or guidance-related dispositions being more natural or useful for various practical or theoretical purposes than others, it's doubtful that any particular cluster will have exclusive claim to metaphysical priority. (Cf. Pryor (ms) on there being a "complex n-dimensional" doxastic space, with different theorists picking out distinct, possibly overlapping regions in the space.)

what follows.) To simply matters, I'll use recent work by Soter (forthcoming) as a guide—both because she's particularly clear and careful and because I think she's basically right.

Let's start with the following expanded diagram and glosses by Soter (forthcoming):
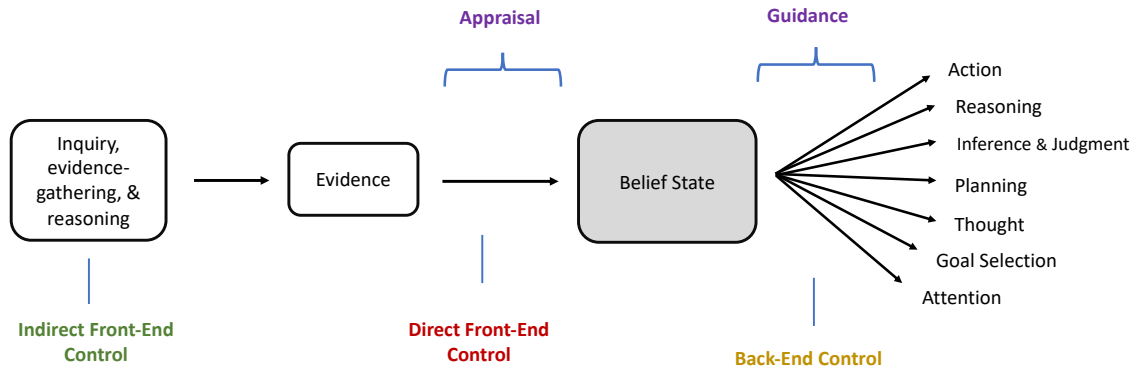

Figure 2 - Different questions of control

The current debate over doxastic (in)voluntarism has mainly focused on whether we have "front-end" control over belief's evidence-responsiveness or *appraisal* function. More specifically, the debate concerns whether we have *direct* front-end control, since everyone agrees we have indirect front-end control (hence the green). Among other things, we can choose "whether to inquire into some topic, how much and what kind of evidence to gather, how carefully and exhaustively to think through a question, etc., and we can do all of this for practical reasons" (4).[15] However, while exercising indirect front-end control is a reliable means of voluntarily *influencing* our belief state, it's not a reliable means of *selecting and securing* a specific belief state. As Alston (1988) observes, there is an important "difference between doing A in order to bring about E, for some definite E, and doing A so that some effect within a certain range will ensue" (271).

Another point of near-unanimity among philosophers is that we generally *lack* direct front-end control (hence the red): we cannot simply choose what we believe, or do so on the basis of just any kind of reason, and this lack of direct control flows from the nature of belief.[16] The basic idea is that beliefs by their nature represent their objects as true, and so—at least in general—can only be directly formed in response to (what are taken to be) epistemic reasons, understood as reasons that are in some way relevant to getting at the truth and avoiding error (cf. Leary 2017).

Details aside, I agree with the standard view—i.e. that we can exercise various forms of indirect front-end control but lack direct front-end control. What Soter adds to the standard view is a compelling defense of the possibility (and reality) of "back-end" control over belief's *guidance* function(s). To motivate the idea, she draws on what we know from cognitive science about the kinds of control we have over other relevantly similar mental states—i.e. ones with a functionally similar "two-pronged" appraisal/guidance architecture (8). She focuses in particular on emotions, summarizing the parallels as follows:

---

[15] Cf. Alston (1988) on "long-range" control and "indirect influence".

[16] Cf. Alston (1988) on the lack of ''basic" and "non-basic immediate" control.

On the appraisal side, both emotions and belief states form and update automatically and nondeliberatively in response to state-appropriate stimuli. Just as threat stimuli automatically elicit fear, evidence (appraised as such) automatically elicits the formation/updating of belief states. In both cases, we are limited in our ability to intervene on the front-end: state-formation is restricted by responsiveness to state-relevant input.

[O]n the guidance side: once elicited, beliefs and emotions both automatically affect a diverse range of psychological mechanisms in state-congruent ways. Indeed, both seem to influence the *same* wide range of cognitive mechanisms and processes. Like emotions, belief states affect our patterns of attention (guiding it towards information that is relevant (Shinoda et al., 2001), supportive of important beliefs (Rajsic et al., 2015), or surprisingly incongruent (Võ & Henderson, 2009)); what we encode into memory and how we recall that information (e.g., Frost et al., 2015; Tuckey & Brewer, 2003; Brewer & Treyens, 1981; Pezdek et al., 1989), how we select actions and set goals (e.g., though shaping assessments about what options are possible (Phillips et al., 2019; Phillips & Cushman, 2017) and what goals we should stick with (Kushnir, forthcoming; Cushman & Morris, 2015)), what thoughts and options spontaneously come to mind (e.g., Bear et al., 2020; Mills & Phillips, 2022), how we evaluate novel information and draw inferences—and so on, across a range of psychological processes. (10-11)

Given that beliefs and emotion-type states share the same type of appraisal-guidance architecture, and given that both exert the same kind of influence over a similar range of psychological processes, Soter draws the conclusion that the guidance processes of beliefs and emotion-type states are mechanistically similar, and that "beliefs also shape cognition and action via the production of state-congruent biases across diverse cognitive mechanisms" (11).

We thus arrive at the possibility of exercising back-end "interventionist" control over our beliefs in much the same way we do over our emotions:

We can prevent belief-appraisals from instantiating their default effects on patterns of attention, thought, memory, goal selection, planning, reasoning, deliberation, action, and so on—systematically blocking the guiding function, and redirecting default responses towards motivation-congruent patterns. Moreover, we [can] deploy this back-end control for any set of moral, practical, or otherwise goal-directed reasons, in cases where our appraisal states are inconsistent with our practical goals or motivations and we don't want them to have their usual effects on our patterns of reasoning, cognition, and action. This reveals an important kind of control: even if our (direct) influence on appraisal side of belief is constrained by evidential reasons, our (direct) influence on guidance-instantiation is not. (11)

Three key features of back-end control are thus that it is (p. 15):

(a) *direct*: the capacity to intervene on guidance is not mediated by other processes;

(b) *non-evidential*: it can be deployed in response to goal-directed, practical, and moral reasons; and

(c) *intentional*: it can be exercised volitionally, overriding default psychological processes.

Soter appeals to these features in explaining the psychological underpinnings of the practical attitude of *acceptance* (as opposed to belief) as well as exploring the possibility of a novel ("back-end") form of doxastic voluntarism. On Soter's (2023) view, acceptance centrally involves a cognitive "gating" function in which we exercise back-end control to prevent a target belief state from having its characteristic downstream effects on reasoning, cognition, and action, and to restructure those downstream processes in ways consistent with what's being accepted. Back-end doxastic voluntarism, by contrast, involves the sustained, comprehensive deployment of back-end control arising out of a self-directed commitment to an alternative. But whereas mundane exercises of back-end control like those involved in gating are common—consider, for example, instances of lying or hypothetical reasoning—full-blown exercises of back-end voluntarism are not (Soter forthcoming, 15). For present purposes, only the former matters.

C. Belief and 'believe'

The first two preliminaries concerned the functional profile of belief and the nature of our control over it. The final preliminary concerns the relationship between belief and the verb 'believe'. For although 'believe' can be used to express the kind of doxastic state of traditional interest to epistemologists and philosophers of action (namely, full belief-like states) it oftentimes doesn't. Instead it's often used in a way that is roughly interchangeably with 'think' to express doxastically weak states, where by '(doxastically) weak' I just mean they lack many of the hallmarks characteristic of full belief.[17] For example, if you ask who Peter thinks is going to win the NBA Championship, I might reply:

(1) Peter believes/thinks that the Bucks will win.

(1) can be true even though Peter merely thinks it's likely—or perhaps merely more likely than the alternatives—that the Bucks will win, and so isn't disposed to take it for granted in reasoning or otherwise treat it as a settled matter. So the truth of (1) doesn't require anything like full belief—it only requires that the Bucks be (something like) Peter's best guess as to who will win.[18] Although guesses are epistemically assessable—they may be more or less justified, given the evidence—they're not the sort of thing that typically amount to knowledge or serve as an appropriate basis for flat-out assertion.

The weakness of (many) ordinary uses of 'believe' is not unique to cases of guessing or prediction—instead, it's rather widespread. For another example, consider the following from Beddor (forthcoming):

Terminal Check Abby hops into a cab, and tells the driver to head towards JFK. "Which terminal?", the driver asks. Abby replies:

(2) I believe my flight leaves from terminal 7, but I'm not sure. Let me check.

---

[17] My use of 'weak' differs from—and is less committal than—that used by Hawthorne, Rothschild, and Spectre (2016).

[18] Cf. Holguín (2022).

Abby pulls out her phone, and, after a minute of navigating the United app, confirms that her flight is indeed leaving from terminal 7. (4)

As Beddor remarks:

> Abby's utterance of [(2)] seems like a perfectly ordinary remark in this context… It also seems like [(2)] could well be true in this situation; it does not seem that Abby is necessarily deluded about her own mental states. [However,] she does not take herself to have settled the question of which terminal her flight leaves from. After all, she immediately follows up her belief ascription with a comment ("Let me check") and an action (confirming the terminal on the United app), both of which reveal that she does not regard the matter as settled. (3)

Many other examples could be (and have been) provided.[19]

Given the doxastic weakness of many ordinary 'believe'-ascriptions, it should come as no surprise the same is true of ordinary 'reason(s)-to-believe'-ascriptions. If you ask Harris who he thinks will win the championship, for example, he might say:

(3) I'm not sure, but the fact that the Celtics have the best record is a reason to believe/think that they'll win.

Here Harris is doing (at least) two things. First, he's denying he has a settled or full belief—in ordinary English talk of being 'sure' often expresses (something like) full belief, rather than (say) Cartesian certainty.[20] Second, he's citing a consideration that clearly increases the likelihood that the Celtics will win, and so counts in favor of increasing one's confidence that they will, but doesn't clearly count in favor of *fully believing* it. This is evidenced by the fact that although (3) is perfectly acceptable, (4) is decidedly less so:

(4) I'm not sure, but the fact that the Celtics have the best record is a reason to fully believe/ be sure/take it for granted that they'll win.

The use of 'believe' in (3) is therefore plausibly a weak one. For other examples, consider:

(5) The dark clouds on the horizon are a reason to believe it's going to rain.
(6) The defendant's lack of an alibi was a reason to believe they committed the crime.
(7) The fact that the odds are so long is a reason to believe your lottery ticket will lose.

In (5)-(7) the reasons cited are all naturally understood as calling for *some* kind of positive doxastic response, but not necessarily full belief in the relevant proposition. *What* kind of positive doxastic response? It's not clear. A lot seems to depend on the strength of the relevant reason.[21] Suppose, for example, that it rains roughly half the time dark clouds appear on the horizon. In such a scenario

---

[19] See Hawthorne, Rothschild, and Spectre (2016) and the ensuing literature on so-called "weak belief".
[20] Cf. Goodman and Holguín (2023).
[21] Here and throughout I set aside complications arising from accepting the "cluster view" of reasons that I defend elsewhere (e.g. Fogal and Worsnip 2023).

(5) is true, even though seeing dark clouds (given the relevant background information) only supports being roughly 50% confident that it'll rain. Alternatively, suppose dark clouds on the horizon *always* portend rain—the latter follows the former without exception. Once again (5) is true, though in this case seeing dark clouds supports being practically certain that it'll rain, and so (presumably) full belief would be justified. But what if it only rains about 20% of the time dark clouds appear on the horizon? In that case (5) remains true, even though seeing dark clouds only supports being roughly 20% confident that it'll rain. This nonetheless counts as a positive doxastic response insofar as it involves increased partial belief—other things being equal, you should be more confident that it'll rain upon seeing dark clouds than you were before.

All that's required, then, for (5) to be true is that the reason cited be a reason to become (as I'll put it) *more positively doxastically disposed* in one way or another towards the relevant proposition. Something similar is true of (6)-(7), along with most (if not all) other epistemic reasons ascriptions.[22] Indeed, this is just what we should expect given the close relationship between epistemic reasons and evidence, together with the familiar "probabilifying" conception of evidence. For as Kelly (2014, §1) notes, it's "natural to think that 'reason to believe' and 'evidence' are more or less synonymous, being distinguished chiefly by the fact that the former functions grammatically as a count noun while the latter functions as a mass term". More carefully, we can say that 'reason to believe' in its *epistemic* (as opposed to, say, practical) sense is more or less synonymous with 'evidence'. Or, to leave open the possibility of non-evidential epistemic reasons, we can say that the *evidential* use of 'reason to believe' is more or less interchangeable with 'evidence'—both are used to pick out sources of evidential support. If we think of evidence as something that bears on the truth or likelihood of a given proposition *p*, with evidence for *p* positively impacting the likelihood of *p* and thereby calling for a positive change in—i.e. an increase in either the level, stability, or specificity of[23]—one's confidence or partial belief, it's a short step to thinking of evidential reasons like (5)-(7) as first and foremost bearing on the truth or likelihood of the relevant proposition, with evidential reasons to believe *p* calling for a positive change in one's partial belief that *p*—one that will often, if not always, fall short of full belief.

Nothing I've said so far is meant to suggest that full belief-like states are *never* expressed by 'believe' nor that reasons for belief are *never* reasons for full belief. Indeed, as Williamson (forthcoming) argues, there are a variety of contexts "in which 'believe' expresses something like the state philosophers have taken it to express", and "[if] 'believe' sometimes stands for full belief, that is vindication enough of philosophers' way of using 'believe'" (20). And for cases where reasons to believe are plausibly reasons for full belief we can look to paradigmatic examples of so-called "wrong kind"—or, more neutrally, "non-standard" (cf. Gertken and Kiesewetter 2017)—reasons to believe, such as those involving bribes, threats, or Pascalian considerations concerning the afterlife. Take, for instance, *Threat* from Quanbeck and Worsnip (forthcoming):

> **Threat**   A powerful evil demon credibly threatens to torture your family for eternity unless you believe that 2+2=5.

---

[22] (7) is particularly noteworthy because philosophers often appeal to lottery-type cases as being ones where even *extremely strong* (but otherwise "bare") statistical evidence fails to support full belief.

[23] Cf. Joyce (2005).

Insofar as the demon's threat is a reason to believe that 2+2=5, it would seem to be a reason first and foremost to *fully believe* that 2+2=5, not merely to become more favorably doxastically disposed (in some way or other) towards the proposition. It's doubtful, for example, that the demon will be satisfied if you merely slightly increase your (initially near-zero) confidence that 2+2=5, or if you merely manage middling confidence that 2+2=5. Something similar would be true if the demon, switching tactics, had instead offered you $1,000,000 to believe that 2+2=5. Insofar as weaker-than-full-belief states are supported at all by the threat or bribe, they are supported *indirectly*—the demon's threat/bribe may be a reason to increase one's confidence that 2+2=5, for example, but if so that's *because* it's a reason to fully believe that 2+2=5, and the latter requires the former.

It's worth noting that 'believe' is not unique in expressing "stronger" and "weaker" mental states. A similar situation arises with respect to 'want' and the desire-like states it can (be used to) express. Condoravdi and Lauer (2014) note, for example, that while '*S* wants *A*' can express that *S* has a mere desire for *A*, it can also express that *S* has an "effective" or action-guiding preference for *A*. To illustrate, they note the two replies below are not contradictory:

(8) Do you want to play tennis?

      a.   I want to, but I have to teach.
      b.   No [= I don't want to], I have to teach.

A natural way to explain the compatibility of (8a) and (8b) is to assume that 'want to play tennis' in (8a) expresses a mere desire to play tennis, while it expresses something stronger—an effective preference or intention—in (8b).

For a similar case involving 'believe', consider (9a) and (9b):

(9) Peter, do you believe the Bucks will win the championship?

      a.   I do—they have the best shot.
      b.   No—though I do think they have the best shot.

The reply in (9a) is plausibly understood as the affirmation of a doxastically "weak", best-guess-like state, while (9b) is plausibly understood as denying a stronger, full-belief-like state.

The overarching moral to draw is that whereas 'pro attitude' does yeoman's service for Davidson in talking about desire-like states relevant to action, 'believe' does yeoman's service for ordinary English speakers in talking about belief-like states: it can express a variety of different states (or "doxastic pro attitudes"), ranging from best guesses to firmly held political and religious convictions, and many things in between. Does this mean 'believe' is semantically ambiguous? No —at least not in the sense of being homonymous. While some take 'believe' to be polysemous (e.g. Westra 2023), others take it to be context-sensitive (e.g. Williamson 2020). Less committally, we can take 'believe'—like most expressions in natural language—to be semantically underspecified: its linguistic meaning constrains but doesn't determine what a competent speaker, speaking literally, can use it to say or communicate (cf. Sperber & Wilson 1995; Carston 2002).

To sum up so far: in addition to recognizing a variety of belief(-like) states as well as different kinds of doxastic control, we should be sensitive to the fact that 'believe' can be used to express "stronger" and "weaker" doxastic states.

III. Application

With the foundation laid, let's now turn to the relationship between the practical and the epistemic. Following Berker (2018) and Quanbeck and Worsnip (forthcoming), we can distinguish three main forms of pragmatism:

> **Austere pragmatism**: only practical considerations constitute genuine (or "authoritative") normative reasons for belief; epistemic reasons are at best "formally" normative (cf. the norms of etiquette or chess).[24]

> **Interactionist pragmatism**: practical and epistemic considerations both constitute genuine normative reasons for belief and both contribute to determining what we all-things-considered ought to believe.[25]

> **Separatist pragmatism**: practical and epistemic considerations both constitute genuine reasons for belief, but cannot be compared or weighed to determine an all-things-considered verdict about what we ought to believe.[26]

All three forms of pragmatism are opposed by:

> **Anti-pragmatism**: practical considerations do not constitute genuine normative reasons for belief; at best only epistemic reasons do.[27]

In what follows, I'll be assuming the falsity of austere pragmatism—as others have argued, epistemic reasons bear many of the hallmarks of genuine normative reasons, and the best explanation of this fact is that they *are* genuine normative reasons.[28] The form of pragmatism I favor is a limited form of interactionist pragmatism:

> **Limited interactionist pragmatism** (**LIP**): practical and epistemic considerations both constitute genuine normative reasons for belief but don't usually combine to determine what we all-things-considered ought to believe.

Though LIP contradicts the letter of separatist pragmatism and anti-pragmatism, there are nonetheless important grains of truth in each, and my defense of LIP will, in different ways, be in the spirit of both.[29] That's because practical and epistemic considerations differ in striking and

---

[24] Cf. Rinard (2019), Mantel (2019), Maguire and Woods (2020).

[25] Cf. Reisner (2008), Leary (2017), Howard (2020).

[26] Cf. Feldman (2000), Kauppinen (forthcoming).

[27] Cf. Berker (2018) and so-called "wrong-kind reason skeptics" like Way (2012).

[28] Cf. Kiesewetter (2021).

[29] For an interesting and interestingly different development of a LIP-like view, see Hirvelä (2023).

systematic ways—ways that typically prevent them from combining to determine what we all-things-considered ought to believe (hence the truth in separatism) and that typically justify denying that practical considerations are normative in the same sense that epistemic ones are (hence the truth in anti-pragmatism).

In what follows I'll highlight three important ways practical and epistemic reasons to believe differ, explain why they matter, and conclude by exploring their collective upshot.

A. First Difference: What is Supported

The first difference between practical and epistemic reasons to believe concerns what they support. Although both bear on belief states (and hence count as "reasons for belief"), they typically do so in importantly different ways. As I'll put it, whereas practical reasons tend to directly support being in a certain doxastic *state*—one that is specifiable independently of the reason and its strength—evidential reasons support doxastically *responding* in a certain way—a way that is *not* specifiable independently of the reason and its strength.

To get an initial handle on the difference, return to the observation made in Section II.C that textbook cases involving so-called "wrong kind" or "non-standard" reasons to believe—think threats, bribes, and Pascalian considerations—typically concern *full belief*, whereas textbook cases involving "right-kind" or "evidential" reasons to believe typically concern a weaker doxastic response, which I'm referring to as "becoming (more) positively doxastically disposed" towards the relevant proposition.[30] Contrast *Threat*, for example, with the following:

> **Expert**   The world's preeminent mathematician holds a major news conference on April 1 announcing the discovery of a proof that 2+2=5.

The news conference, we can suppose, is a joke. To enhance comedic effect, however, the "discovery" is presented as genuine, and it subsequently goes viral on social media, with many people mistakenly thinking the discovery might be real. It's plausible that, for at least some social media users, the fact that there is a genuine-looking announcement of a proof that 2+2=5 is a reason (however weak) to believe that 2+2 does, in fact, equal 5.

The contrast between *Threat* and *Expert* helps illustrate the more general point that practical reasons for belief are (at least typically) *state-oriented* in a way that evidential reasons are not, and that evidential reasons are *content-oriented* in a way that practical reasons are not.[31] The difference can be cashed out, in part, in terms of what is directly supported. In *Threat*, for example, the reason supports *having the full belief that 2+2=5* (i.e. being in that state), while in *Expert* the reason supports

---

[30] Worsnip (2021) also considers cases like *Bribe*:
> **Bribe**. You have strong (but not utterly infallible) evidence that the game starts at 3pm. But I, an eccentric millionaire, offer you a bribe: if you can avoid believing that the game starts at 3pm, I will give you $1,000,000.

Here the bribe concerns the absence (rather than the presence) of belief, but like *Threat* it's full belief that's plausibly at issue.

[31] This is different than the more familiar distinction between "state-given" and "object-given" reasons, which concerns what *provides* reasons.

becoming more favorably doxastically disposed toward the proposition that 2+2=5 (i.e. responding in that way), where what that amounts to isn't settled by the reason itself but instead depends, inter alia, on the strength of the reason as well as one's existing level of justified confidence. (More doxastic adjustment is called for in the case of the benighted social media user than the professional mathematician, for example.) The content of the full belief, unlike that of the graded doxastic response, is in a sense incidental—it's the *state* with that content that the reason supports. This remains true even if we modify *Threat* so that one merely needs to increase one's confidence in 2+2=5 rather than fully believe it—the threat targets (and hence supports being in) an independently and antecedently specifiable state with a certain content, but isn't otherwise sensitive to the content itself. Evidential reasons, by contrast, *are* sensitive to the content itself and do *not* generally target an independently and antecedently specifiable state.[32] All we can say absent further information is that they support becoming more favorably doxastically disposed in some way or other towards the relevant proposition, but what specific way that is depends on a range of contextually-variable factors.

This doesn't automatically mean that epistemological tradition is mistaken: all of the evidential reasons cited so far—including (3) and (5)-(7) above—might also, as a matter of fact, support full belief to some (possibly small) degree, even if that's not what they're ordinarily understood as doing. But given that the reasons cited *also* support becoming more positively doxastically disposed towards the relevant proposition—i.e. positive changes in graded belief—we face the question: what is the relationship between the evidential support provided for full and graded belief? Do evidential reasons support full belief in virtue of supporting (positive changes in) graded belief? Do they support graded belief in virtue of supporting full belief? Or do they provide support for graded belief and *separately* also provide support for full belief?

None of these options are terribly attractive. I myself take evidential reasons to first and foremost support (positive changes in) one's graded doxastic attitude towards *p,* and to "support" full belief only in an extended, derivative sense. To illustrate, consider what Worsnip (2021) calls the "beaker of reason(s)":
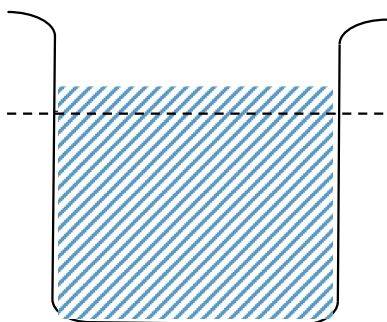
As Worsnip notes, "the shaded area represents how much [overall evidential reason or support] there is to believe *p*, where this is a function of the reasons for and against believing *p*, and their

---

[32] This helps explain why there can be "wrong kind" reasons to be in states lacking an object (e.g. someone might offer you money to feel itchy) but not "right kind" reasons (cf. Howard and Leary 2022).

weight. The higher up in the beaker that the shaded area goes, the more reason there is to believe p… The black dotted line represents [the] threshold for justified belief. When the [amount of reason exceeds] the black dotted line, there is sufficient reason to believe" (534). I like the beaker model, and Worsnip's description of it, but given the foregoing we shouldn't assume the reasons for/against believing p are reasons for/against *fully* believing p—that would need to be argued for. On my view, the relevant reasons only directly bear, whether positively or negatively, on one's graded doxastic state: reasons *for* belief call for positive changes in one's graded doxastic state (i.e. becoming more favorably doxastically disposed) while reasons *against* call for negative changes. What they cumulatively (and directly) support overall is also a graded doxastic state, albeit an "overall" one: how confident one should be overall is a function of how much evidential support there is overall—i.e. how high up the shaded area goes.

So far full belief hasn't entered the picture—evidential reasons, whether individually or collectively, only directly concern graded doxastic states (and changes to them). Insofar as evidential reasons "support" full belief, they do so indirectly and in an importantly different sense. It's *indirect* insofar as what evidential reasons (directly) contribute to is the size of shaded area, and it's the size of the shaded area—not any of the reason(s) contributing to it—that then determines whether full belief justified. By increasing the size of the shaded area, the relevant reason(s) can bring the shaded area *closer to* or *past* the dotted line (if it initially falls short) or *further above* the dotted line (if it already meets or exceeds it), and in *that* sense "support" full belief. But this is importantly different from the primary notion of support obtaining between evidential reasons and (changes in) graded doxastic states. Indeed, "support" for full belief is doubly derivative: it obtains in virtue of facts about the primary notion of support (i.e., the reason's impact on the shaded area) as well as facts about the justificatory status of full belief (i.e., the relationship between the shaded area and the dotted line), with facts concerning the latter in turn depending on facts about overall support in the primary sense. So while it's true that evidential reasons *help determine* whether full belief is justified, they do so indirectly, and they don't "support" full belief in anything like the way they support (changes in) graded belief.

B. Second Difference: Transmission and the Varieties of Reasons

The difference in what practical and evidential reasons support is reflected in how the different kinds of support "transmit". For example, whereas evidential reasons to believe p typically transmit to, e.g., the obvious entailments of p, practical reasons to believe p don't, and whereas practical reasons to believe p typically transmit to the necessary/sufficient/probability-increasing means to believe p—where this may (and often will) include actions—evidential reasons don't. Way (2012) takes this difference in transmission to motivate "skepticism" about non-standard, practical reasons for belief, since he takes the latter to be the best explanation of the former. I think it's a mistake, however, to deny that practical reasons to believe are what they appear to be— namely, reasons to be in a certain doxastic state. Indeed, there are practical reasons to believe in much the same way (and in the same sense) as there are practical reasons to be in various other states, whether mental or not, and to have certain dispositions or character traits. There are good practical reasons to be debt-free, to (not) be a parent, to own a home, to be happy, to have lower

blood-pressure, to be punctual, etc. Practical, non-standard reasons to believe are just an instance of this general phenomenon.

There is therefore no need to try to analyze practical reasons to believe in terms of some *other* kind of state or activity that there can be reasons for. Skeptics have claimed, for example, that practical reasons to believe are "really" reasons to *want* to believe, or *intend* to believe, or to (try to) *bring it about* that you believe.[33] While I think there is an important grain of truth in such proposals, I see no need to deny the datum that practical reasons to believe are, in a straightforward and literal sense, reasons to believe. That's because state(s-of-affairs)-oriented reasons—to be debt-free, to (not) be a parent, and so on—are all, in the first instance, naturally understood as so-called "objective" or "fact-relative" reasons. The example standardly used to motivate the distinction between more "objective" and more "subjective" senses of 'reason(s)' is Bernard Williams' (1981) famous petrol case. (Terminological note: 'reason(s)' is meant to do double-duty for both mass and count noun uses of 'reason'.[34]) In this case, you have a glass of liquid that you believe to be gin and tonic, but that is in fact petrol. In an objective sense, there is good reason for you *not* to drink, since the liquid is petrol, and petrol is bad for you. But in a *subjective* sense, there is good reason for you *to* drink, since you believe the liquid to be gin and tonic, and you want a gin and tonic. A related objective/subjective distinction arises with 'ought' and other deontic terms.

As I argue elsewhere, however, the dichotomy between objective and subjective reason(s)—and between objective and subjective oughts—is insufficiently fine-grained (Fogal and Worsnip 2021).[35] To see why, distinguish two versions of the petrol case:

> **Cleverly Disguised Petrol.** Though it is in fact petrol, the stuff in the glass looks like gin and tonic, smells like gin and tonic, and has been served to you by a barman in response to your request for gin and tonic. On the basis of this evidence, you believe that it is gin and tonic.

> **Obviously Petrol.** As well as in fact being petrol, the stuff in the glass looks (to you and everyone else) like petrol, smells like petrol, and is sitting around in a car mechanic's garage. In defiance of all this evidence, however, you believe that it is gin and tonic.

In *Cleverly Disguised Petrol*, your belief that the glass contains gin and tonic is well-supported by your evidence, and hence justified or (as it's now put) substantively rational. In *Obviously Petrol*, it isn't. This helps us see that there are (at least) *three* bodies of information to which our 'reason(s)'-talk can be relativized.[36] In the objective, "fact-relative" sense, you have good reason to refrain from drinking both in *Cleverly Disguised Petrol* and in *Obviously Petrol*, since the glass in fact contains petrol in both cases, and there is little to no countervailing reason to drink. In the subjective, "attitude-relative" sense, you have good reason to drink both in *Cleverly Disguised Petrol*

---

[33] Cf. Way (2012), Berker (2018).

[34] For more on the distinction, see Fogal (2016).

[35] The following paragraphs are adapted from Fogal and Worsnip (2021).

[36] Reason(s)-talk is much more context-sensitive than this, though I'm trying to suppress additional complications. For a survey of ways epistemic reasons ascriptions are (or might be) context-sensitive, see Fogal and Sylvan (2017).

and in *Obviously Petrol*, since in both cases you *believe* that the glass contains gin and tonic and you *want* a gin and tonic. In the third, more intermediate sense, you have good reason to drink in *Cleverly Disguised Petrol*, but you *lack* good reason to drink in *Obviously Petrol* (and indeed have good reason *not* to drink).

This intermediate notion of reason(s) is sensitive to the salient difference between *Cleverly Disguised Petrol* and *Obviously Petrol*, namely that in *Cleverly Disguised Petrol*, you have good (evidential) reason to believe that the glass contains gin and tonic, whereas in *Obviously Petrol*, you don't. This helps explain why, given your (presumably unobjectionable) desire for a gin and tonic, it's substantively rational for you to (intend) to drink in *Cleverly Disguised Petrol* but not in *Obviously Petrol*.[37] So let's call it the "evidence-relative" notion of reason(s).

For good measure we can introduce a third case:

Cleverly Disguised G&T. Though in fact a gin and tonic, the stuff in the glass looks like petrol, smells like petrol, and is sitting around in a car mechanic's garage. On the basis of this evidence, you believe that it is petrol.

In *Cleverly Disguised G&T*, you have good reason to drink in the fact-relative sense but not the evidence-relative or attitude-relative sense. We can summarize the verdicts about cases as follows (adapted from Fogal and Worsnip 2021):

	What there is good reason to do in… / What you ought to do		
	…*Cleverly Disguised Petrol*	…*Obviously Petrol*	…*Cleverly Disguised G&T*
Fact-relative	Refraining	Refraining	Drinking
Evidence-relative	Drinking	Refraining	Refraining
Attitude-relative	Drinking	Drinking	Refraining

Although many philosophers will want to exclude attitude-relative reasons from the realm of the genuinely normative, that still leaves two not-merely-formal normative notions: a fact-relative one and an evidence-relative one. And while one might seek greater unity by trying to analyze one in terms of the other, doing so is a tall order—and one that, even if successful, doesn't change the fact that we need to be sensitive to the (often overlooked) distinction when theorizing about reasons.

Though there can be fact-relative reason(s) without evidence-relative reason(s) (e.g. *Cleverly Disguised G&T*) and evidence-relative reason(s) without fact-relative reason(s) (e.g. *Cleverly Disguised Petrol*), there are important relationships between them. It's plausible, for example, that being aware of a given fact-relative reason for you to (say) call your mom, be debt-free, or be less anxious will generally give you evidence-relative reason to call your mom, be debt-free, and be less anxious. What's more, as a result, you will also have evidence-relative reason to *do what you can* to call your mom, become debt-free, and become less anxious. As a matter of fact, of course,

---

[37] Substantive rationality is commonly contrasted with structural rationality, or coherence. See Kiesewetter and Worsnip (2023) for discussion and references.

there may be little if anything you can do to perform the relevant action or to help bring about the relevant state(-of-affairs), in which case the reason will be effectively neutered, or normatively idle, when it comes to substantive rationality and agency[38]—subtleties aside, you're not rationally criticizable for failing to be or do (believe, intend, etc.) what you can't be or do.[39] Considerations concerning the ability (i.e. the capacity and opportunity) to respond correctly to reasons—a form of largely *non*-voluntary "control" we have over our attitudes[40]—thus distinguish what I'll call "substantive" reason(s) from the broader category of evidence-relative reason(s): it's only reason(s) in the substantive sense that determine(s) what it's substantively rational to do, believe, etc.

C.   Third Difference: Control Revisited

This brings us to the third (and most familiar) difference between practical and epistemic reasons: in addition to differences in *what* they support—both directly and indirectly via transmission— practical and epistemic reasons differ in *how we can respond* to them. As noted in Section II.B, we typically lack direct "front-end" control over our belief state, and can't directly form beliefs on the basis of (what we take to be) merely practical reasons.[41] Instead, our belief state is generally—and largely automatically—only directly responsive to be (what we take to be) epistemic reasons, understood as reasons that are relevant to getting at the truth and avoiding error.

Assuming a suitable 'ought-implies-can'-style principle, this broadly psychological difference makes an important normative difference. In particular, it means that practical reasons don't generally bear on the question of what one substantively ought to believe *here and now*(-*ish)* —or, when evaluating attitudes located at another point in time, *there and then*(-*ish*). (Note that here 'believe' is used in its broad, inclusive sense—one that's roughly interchangeable with 'think'— and hence includes states of both full and partial belief.) I take this ought to be closely related, if not identical, to what is called the "central" or "deliberative" ought, though given certain controversies concerning the latter I'll refer to it as 'ought$_{SR}$' ('SR' for 'substantive rationality').[42] I'll also leave the temporal qualification implicit, despite it's importance. For although it's rarely remarked upon, non-finite clauses in general—and 'to'-infinitivals in particular—never carry their own primary tense to convey the location in time of the relevant situation (state, event, action, etc.). The relevant time has to instead be inferred from the broader context. Philosophers would do well to be more explicit, similar to how we've learned to be more explicit about the different "flavors" and information-sensitivity of modals like 'ought' (and 'reason(s)').

---

[38] Though it might still make a difference to what it's substantively rational to *want* (in the "weak" sense).

[39] Cf. Schwarz (2020) on the difference between stronger, "agential" uses of 'can' and weaker, "circumstantial" uses. The use of 'can' in the present context plausibly requires both the opportunity to respond to *r* and the general capacity to respond to *r in the right way* (roughly, "as a reason")—the mere possibility of being influenced by *r* in some way or other doesn't suffice for having the relevant ability.

[40] For relevant discussion and references, see Hieronymi (2009) and McHugh (2017).

[41] At least not in a conscious, clear-eyed way. The mechanisms at work in, say, motivated reasoning standardly operate subconsciously and are outside of our direct control (cf. Kunda 1990).

[42] Pittard (2023) argues that there are a multitude of evidentially-constrained "deliberative" oughts concerning the future, but his view is compatible with there being a single "deliberative" ought concerning the here-and-now, or what I'm calling the substantive ought.

The upshot so far: given our lack of direct front-end control, the only considerations typically bearing on what doxastic attitude(s) we ought$_{SR}$ to have are epistemic reasons. However, although practical reasons don't typically bear on the question of what we (here and now) all-things-considered ought$_{SR}$ to believe, they *do* often bear—albeit derivatively—on what we all-things-considered ought$_{SR}$ to *do*. For as noted above, if we have fact-relative reason to perform some action *A* or be in state *S*, and if we're aware of that reason, then we (at least typically) thereby have evidence-relative reason to do what we can to *A* or be in *S*. And if there *is* something we can do (here and now) to help bring it about, or otherwise make it more likely, that we *A* or are in *S*, then we have reason in the substantive sense—'reason$_{SR}$' if you like—to do so, where the strength of the relevant substantive reason is determined, in part, by the strength of the initial reason and modulated by the likelihood of success. So if we're aware that we have practical (fact-relative) reason to believe *p*—as we are in the relevant cases involving threats, bribes, and Pascalian considerations—and if there is something we can do to help bring the belief about, then we have substantive reason to do so.[43] This means that when it comes to practical reasons for belief we will typically have "actionable" (substantive) practical reason to exercise various forms of indirect front-end and back-end doxastic control—control that, depending on the case (e.g. believing 2+2=5), may be rather limited and guaranteed to be ineffective.

I've covered a lot of ground rather quickly. It may help to illustrate the basic view on offer:
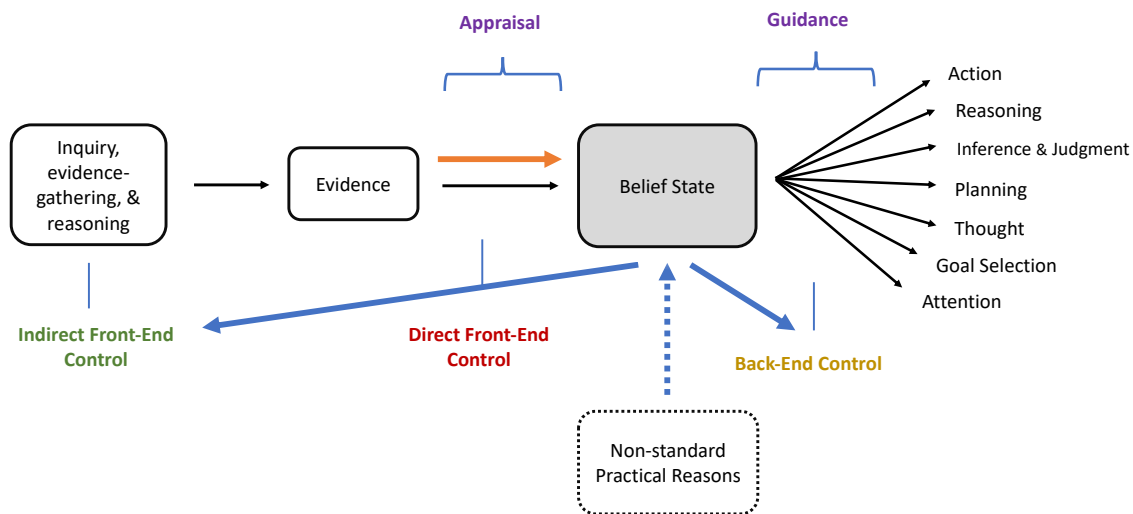

Figure 3 - A Process Model of Non-Standard Reasons for Belief

The **orange line** represents evidential support and the **blue lines** represent non-evidential (including practical) support, with the **solid lines** representing "substantive" forms of support and the **dotted line** representing "non-substantive" (e.g. merely fact-relative) forms support. Not

---

[43] What about cases involving misleading evidence? Suppose, for example, that you justifiably but falsely believe that an eccentric billionaire will pay you a million dollars to believe that 2+2=5. In that case, I think you still have substantive (but not fact-relative) reason to do what you can to get yourself to believe 2+2=5, assuming there is something you can do. (Thanks to Selim Berker for prompting this clarification.)

represented is the fact that the practical support for indirect front-end and back-end control depicted derives from the support for the belief state—the former obtains in virtue of the latter.

We can further modify Figure 3 by drawing attention to an important class of epistemic ("right kind") reasons to believe that are *not* evidential reasons—i.e. reasons that are relevant to getting at the truth and avoiding error concerning some proposition p and that one can be epistemically criticizable for failing to properly respond to, but that don't directly bear on the truth or likelihood of p itself. For example, many reasons relevant to whether we should suspend judgment—or whether we should put off, postpone, or "wait" to form a belief, as McGrath (2021) puts it[44]—are not directly truth-related, despite being epistemic. This includes what McGrath calls "future-comparative factors", such as whether you will later have better (or worse) evidence concerning whether p than you have now or whether you will later be in a better position to assess your current evidence. Many of the reasons bearing on whether we should inquire and/or continue to inquire into a subject matter are also typically non-evidential in nature, despite at least some plausibly counting as epistemic. This includes what McGrath calls "goal-related factors", such as how valuable it would be to know whether p and how likely it is that, if you inquired further and acquired more evidence, you could come to know whether p.

The kinds of reasons just mentioned straddle the traditional practical/epistemic divide. The support they provide is not evidential and what they support (in the first instance) are various kinds of mental actions or processes that we typically have (more or less) direct control over—e.g. ceasing deliberation, deciding to (not) make up one's mind, inquiring further, double-checking, etc —rather than any particular belief state or response. The kinds of things they support nonetheless impact—in varying ways and to varying degrees—one's overall belief state, and are plausibly understood as forms of indirect front-end and/or back-end control. We can thus modify Figure 3 to get:
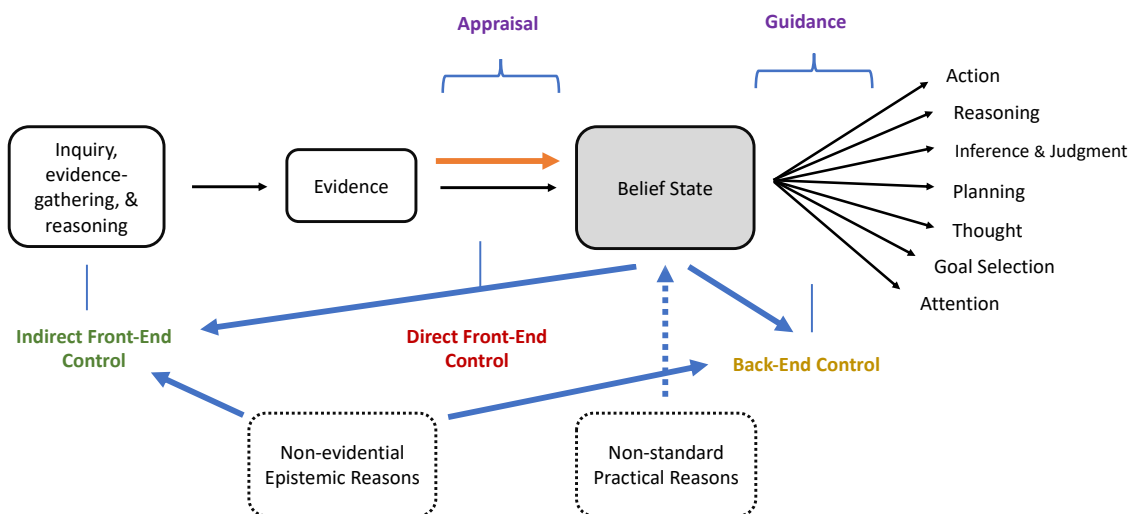

Figure 4 - A Process Model of Reasons for Belief

---

[44] Suspending judgment is thus distinct from agnosticism, or committed neutrality.

D. Limited interactionist pragmatism and encroachment

Summarizing in something like slogan form: practical reasons to believe typically bear on how we ought to *be* doxastically—i.e. on what doxastic state it would be good or best for us to be in—but not on how we ought to *respond* doxastically—i.e. on what doxastic attitude we rationally, here and now, ought to have with respect to *p*. Only epistemic reasons typically bear on the latter. This is the important grain of truth in anti-pragmatism, even though anti-pragmatism is false since both kinds of reasons for belief are what they appear to be—namely, genuine, not-merely-formal normative reasons for belief. Practical and epistemic reasons nonetheless differ in the not-merely-formally-normative oughts they typically contribute to determining: practical reasons typically bear (in the first instance) on the fact-relative ought, while epistemic reasons typically bear on the substantive ought.[45] Although what we ought to do in the fact-relative sense often comes apart from what we ought to do in the substantive sense, such differences in verdicts don't "conflict" in any problematic or puzzling sense. This is the important grain of truth in separatist pragmatism, even though separatist pragmatism is false since it's in principle possible for *both* practical *and* epistemic considerations to bear on what one substantively ought to believe, all-things-considered.

Whether a given practical consideration bears on what one ought$_{SR}$ to believe depends largely on whether one has the ability—i.e., the general capacity plus opportunity—to directly respond to it. Although we usually lack this ability, there are plausible exceptions. Quanbeck and Worsnip (forthcoming), for example, defend what they call "permissivist pragmatism", according to which "when there is more than one epistemically permitted doxastic attitude, practical (including moral) considerations can come in to determine which epistemically permitted doxastic attitude one all-things-considered ought to have" (3). They argue that permissivist pragmatism avoids many of the standard worries about pragmatism, including worries about (lack of) control. For whereas we lack the ability to form or revise our beliefs in response to "incentives to believe something for which one obviously has inadequate evidence", as in cases like *Threat*, it's far less obvious that "practical considerations can't motivate us to believe (or suspend) in cases where we take the evidence to merely *permit*, but not *require*, believing" (19).

Like many others, including Quanbeck and Worsnip, I take epistemic permissivism to be particularly plausible concerning categorical doxastic atittudes (full belief, disbelief, agnosticism). For given the picture provided so far, how we epistemically ought to respond to our evidence (including our higher-order evidence) is first and foremost a matter of how confident we should be —i.e. what graded doxastic state we should be in. Insofar as full belief is incompatible with low (say, below .5) confidence—the appraisal and/or guidance dispositions characteristic of each can't be jointly satisfied[46]—our evidence will also, derivatively, impose constraints on what categorical doxastic states we can or should be in. But even so, as many others have noted, our evidence will often (if not always) leave some leeway concerning which particular categorical attitude (if any) to adopt. One's evidence can support a high degree of confidence in *p*, for example, and therefore be rationally incompatible with disbelief in *p*, without thereby mandating full belief that *p*—one

---

[45] Note that I'm *not* skeptical about there being an all-things-considered ought. I'm only skeptical of there being just one.

[46] At least absent fragmentation and the like.

needn't be making an epistemic mistake in not settling the matter of whether p or not being disposed to take p for granted in reasoning, etc.[47]

Permissivism thus opens up the possibility of (not) believing for practical reasons in the relevant sense—i.e. of deciding whether or not to form a given categorical doxastic attitude at least partly in response to practical considerations. (The recognition of this possibility has a long, illustrious history—including by Augustine, Aquinas, Descartes, and Kierkegaard—and a range of more recent authors have (re-)recognized it as well.[48]) Quanbeck and Worsnip write:

> Suppose [one] judges oneself to have pretty strong but not completely infallible evidence for p, and that one judges that either believing or suspending on p would be epistemically permissible. One nevertheless faces the question of whether to believe p or suspend judgment about whether p. We think it's quite plausible that in such a case, one can choose to believe p or suspend judgment about whether p on the basis of practical considerations.[49]

Indeed, as they note, we *must* have a certain amount of control over whether we (fully) believe if it's possible to respond to pragmatic and moral considerations in the standard sorts of cases motivating encroachment, all of which are plausibly taken to be permissive. Consider, for example:

> **Wine Stain.** Suppose that you have struggled with an alcohol problem for many years, but have been sober for eight months. Tonight you attend a departmental reception for a visiting colloquium speaker, and are proud of withstanding the temptation to have a drink. But when you get home, your spouse smells the wine that the colloquium speaker spilled on your sleeve while gesticulating to make a point, and you can see from her eyes that that she thinks you have fallen off of the wagon. (Basu and Schroeder 2019: 182)

According to Quanbeck and Worsnip, in *Wine Stain* it's "natural to say that one can choose to suspend judgment on whether one's alcoholic spouse has relapsed – as opposed to settling on the belief that they have done so… One can choose to inquire into the matter more *rather than* making up one's mind and believing now" (19, italics in original). Although this doesn't amount to having full front-end doxastic control—one's confidence level, for example, is largely automatically generated—it does involve more-or-less direct control over whether to make up one's mind (or remain settled), and hence whether one forms (or maintains) a full belief. Whether it's best to categorize this as a limited form of direct front-end control or instead a very direct (proximate) form of indirect control is unclear, but not much of substance hangs on the choice.

Quanbeck and Worsnip proceed to argue that permissivist pragmatism allows us to say much of what proponents of encroachment want to say *without* saying that practical considerations

---

[47] Indeed, I'm sympathetic with the thought that one's evidence alone *never* positively requires full belief (cf. Nelson 2010).

[48] For additional historical references, together with critical discussion, see Alston (1988).

[49] There is an important difference between cases that *are* epistemically permissive and ones that we— perhaps wrongly—*take* (if only implicitly) to be permissive. Although Quanbeck and Worsnip tend to blur the difference, it's the latter that arguably matters most when it comes to the ability to respond directly to practical reasons.

encroach on epistemic justification. I'm sympathetic, though my view slices things up slightly differently, splitting the difference between the encroacher and the permissivist pragmatist. To begin with, note that the official gloss of pragmatic encroachment uses the mass noun 'justification' in two different ways:

> **Pragmatic encroachment:** practical considerations bear on what we ought to believe by affecting epistemic justification (e.g. how much justification is required to justifiably believe).

The first use of 'justification' expresses a categorical, threshold-y notion—one that is interchangeable with talk of being justified. The second use (in parentheses) expresses a graded notion, admitting of degrees. The former is naturally understood in terms of the latter: to have justification in the categorical sense—i.e. to be justified—is to have *adequate* or *sufficient* justification in the graded sense—i.e. sufficient to make it epistemically appropriate to believe.[50] The categorical notion is thus fit-related or "aptic" (to use Berker's 2022 phrase): to be epistemically justified is to be epistemically *appropriate* or *fitting*, where this isn't analyzable in standard deontic terms (i.e. in terms of what you *ought*, *may*, or *must* believe).[51]

To see how my view differs from the encroacher's and permissivist pragmatist's, consider the following classic case motivating encroachment:

> **Low-Stakes Train.** You're at the train station in Boston, preparing to travel to Providence for a relaxing vacation. You ask someone standing beside you whether the train stops in Foxboro. They reply, "It does—they told me when I bought the ticket."

> **High-Stakes Train**. You absolutely need to be in Foxboro, the sooner the better. Your career depends on it. You overhear a conversation like that in Low-Stakes Train concerning the train that just arrived and leaves in 15 minutes. (Cf. Fantl and McGrath 2002)

According to the encroacher, you need more evidence to be justified in (fully) believing that the train stops in Foxboro in *High-Stakes Train* than you do in *Low-Stakes Train*, since the costs of relying on your belief—i.e., of taking it for granted and acting accordingly—in the former are much higher than in the latter. The encroacher thus thinks that you'd be making an epistemic mistake in *High-Stakes Train* (but not *Low-Stakes Train*) if you were believe that the train stops in Foxboro merely on the basis of the casual testimony. Quanbeck and Worsnip, by contrast, think you'd merely be making a *prudential* mistake if you were to (fully) believe in *High-Stakes Train*—believing is epistemically permissible, but one all-things-considered ought to suspend judgment instead.

My preferred diagnosis splits the difference. Like Quanbeck and Worsnip, I take *High-Stakes Train* and *Low-Stakes Train* to represent epistemically permissive situations: (full) belief is permitted but not required. However, like the encroacher I think there is a good sense in which you would be *epistemically* (and not just prudentially) criticizable if you were to form that belief in *High-Stakes Train* and act on it without double-checking, etc. This is evidenced by the fact that it would be

---

[50] Bare claims concerning whether *S* has justification to believe *p* (or concerning what provides such justification) are thus ambiguous between graded vs. categorical construals. The difference is roughly the difference between claiming that *S* has *at least some* justification vs. *enough* justification to believe *p*.

[51] See Berker (2022) on the need to distinguish the deontic, the evaluative, and the aptic.

appropriate (among other things) to reduce *epistemic trust* in you in response to your treating the matter as settled in *High-Stakes Train* without seeking confirmation (cf. Kauppinen 2018). What's epistemically problematic, however, is *not* the belief itself but rather the failure to "gate" the belief —i.e. the failure to exercise back-end control to prevent the belief state from having its characteristic downstream effects on reasoning and action. If you trust what you (over)hear but seek confirmation to be extra sure, any sense of epistemic impropriety seems to disappear. In other words, even though it's epistemically okay to believe the train stops in Foxboro in *High-Stakes Train* on the basis of the testimony, and thus be *generally* disposed to take it for granted in acting and reasoning, it's *not* okay—given the stakes—to *occurrently* take it for granted in acting and reasoning.[52] An epistemically responsible agent will either *not* form the relevant belief to begin with (absent further confirmation) or else form the belief but "gate" it (via the exercise of back-end control), thereby reasoning and acting in ways that treat the matter as not-yet-settled.

What's more, whereas Quanbeck and Worsnip treat the practical considerations in encroachment cases as constituting practical reasons for/against (full) belief, I think they are typically more plausibly viewed as reasons to exercise various forms of indirect front-end and/or back-end control—to (not) inquire further, (not) consider alternative hypotheses, (not) deliberate carefully, (not) seek confirmation, and so on. Importantly, however, these practical considerations, when appropriately taken into account, will have an *epistemic bar-raising* effect in high-stakes cases —they will make a difference to the amount and kind of evidence or justification you end up having for believing *p*, resulting in an epistemically improved position. So there is a grain of truth in encroachment: although the amount of justification required for justified—i.e. epistemically appropriate—belief plausibly remains the same in high-stakes and low-stakes cases, the amount of justification required for justified *ungated* belief does not.[53]

References

Alston, William P. (1988). The deontological conception of epistemic justification. Philosophical Perspectives 2:257-299. doi.org/10.2307/2214077

Basu, R. & Schroeder, M. (2019). Doxastic wronging. In B. Kim & M. McGrath (Eds.), *Pragmatic Encroachment in Epistemology* (pp. 181-205). New York: Routledge.

Beddor, B. (forthcoming). Unsettled belief. *The Philosophical Quarterly*. Manuscript available online: https://philarchive.org/rec/BEDUBG

Berker, S. (2018). A combinatorial argument against practical reasons for belief. *Analytic Philosophy, 59*, 427-470. doi.org/10.1111/phib.12140

Berker, S. (2022). The deontic, the evaluative, and the fitting. In C. Howard and R. A. Rowland (Eds.), *Fittingness: Essays in the Philosophy of Normativity* (pp. 23–57). Oxford: Oxford University Press.

Bratman, M. E. (1992). Practical reasoning and acceptance in a context. *Mind, 101*, 1-16. doi.org/10.1093/mind/101.401.1

Carston, R. (2002). *Thoughts and Utterances*. Oxford: Blackwell.

Lauer, S. and Condoravdi, C. (2014). Preference-conditioned necessities: detachment and practical reasoning. *Pacific Philosophical Quarterly*, *95*, 584-621. doi.org/10.1111/papq.12049

Fantl, J. and McGrath, M. (2002). Evidence, pragmatics, and justification. *Philosophical Review*, *111*, 67-94. doi.org/10.1215/00318108-111-1-67

Feldman, R. (2000). The ethics of belief. *Philosophy and Phenomenological Research*, *60*, 667-695. doi.org/10.2307/2653823

Flores, C. (forthcoming). Resistant beliefs, responsive believers. *Journal of Philosophy*.

Fogal, D. (2016). Reasons, reason, and context. In E. Lord and B. Maguire (Eds.), *Weighing Reasons*. Oxford: Oxford University Press.

Fogal, D. and Sylvan, K. (2017). Contextualism about epistemic reasons. In J. J. Ichikawa (Ed.), *The Routledge Handbook of Epistemic Contextualism*. Routledge.

Fogal, D. and Worsnip, A. (2021). Which reasons? Which rationality? *Ergo, 8*. doi.org/10.3998/ergo.1148

Fogal, D. and Worsnip, A. (2024). What the cluster view can do for you. In R. Shafer-Landau (Ed.), *Oxford Studies of Metaethics, 19*. New York: Oxford University Press.

Friedman, J. (2019). Inquiry and belief. *Noûs*, *53*, 296- 315. doi.org/10.1111/nous.12222

Gertken, J. and Kiesewetter, B. (2017). The right and the wrong kind of reasons. *Philosophy Compass, 12*. doi.org/10.1111/phc3.12412

Goodman, J. and Holguín, B. (2022). Thinking and being sure. *Philosophy and Phenomenological Research, 106*, 634-654. doi.org/10.1111/phpr.12876

Hawthorne, J., Rothschild, D., and Spectre, L. (2016). Belief is weak. *Philosophical Studies*, *173*, 1393–1404.

Hieronymi, P. (2009). Believing at will. *Canadian Journal of Philosophy*, *39*, 149–187.

Hirvelä, Jaakko (2023). The structure of moral encroachment. *Philosophical Studies, 180*, 1793-1812. doi.org/10.1007/s11098-023-01949-z

Holguín, B. (2022). Thinking, guessing, and believing. *Philosophers' Imprint, 22,* 1-34. doi.org/10.3998/phimp.2123

Holton, R. (2014). Intention as a model for belief. In M. Vargas and G. Yaffe (Eds.), *Rational and Social Agency*. Oxford University Press.

Howard, C. (2020). Weighing epistemic and practical reasons for belief. *Philosophical Studies, 177*, 2227-2243. doi.org/10.1007/s11098-019-01307-y

Howard, C. and Leary, S. (2022). In defense of the right kind of reason. In C. Howard and R. Cosker-Rowland (Eds.), *Fittingness* (pp. 221–242). Oxford: Oxford University Press.

Jackson, E. (2019). How belief-credence dualism explains away pragmatic encroachment. *Philosophical Quarterly, 69*, 511-533. doi.org/10.1093/pq/pqz006

Jackson, E. (2020). The relationship between belief and credence. *Philosophy Compass, 15*, 1–13. doi.org/10.1111/phc3.12668

Joyce, J. M. (2005). How degrees of belief reflect evidence. *Philosophical Perspectives, 19*, 153-179. doi.org/10.1111/j.1520-8583.2005.00058.x

Kauppinen, A. (2018). Epistemic norms and epistemic accountability. *Philosophers' Imprint*, *18*. hdl.handle.net/2027/spo.3521354.0018.008

Kauppinen, A. (2023). The epistemic vs. the practical. In R. Shafer-Landau (Ed.) *Oxford Studies in Metaethics 18* (pp. 137-162). New York: Oxford University Press.

Kelly, T. (2014). Evidence. *Stanford Encyclopedia of Philosophy.* In E. Zalta (Ed.) *Stanford Encyclopedia of Philosophy* (Winter 2016 Edition). https://plato.stanford.edu/archives/win2016/entries/evidence/.

Kiesewetter, B. (2021). Are epistemic reasons normative? *Noûs, 56,* 670-695. doi.org/10.1111/nous.12379

Kiesewetter, Benjamin & Worsnip, Alex (2023). Structural rationality. In E. Zalta & U. Nodelman (Eds.) *Stanford Encyclopedia of Philosophy* (Fall 2023 Edition). https://plato.stanford.edu/entries/rationality-structural/

Kunda, Z. (1990). The case for motivated reasoning. *Psychological Bulletin, 108,* 480–498. doi.org/10.1037/0033-2909.108.3.480

Leary, S. (2017). In defense of practical reasons for belief. *Australasian Journal of Philosophy, 95*, 529-542. doi.org/10.1080/00048402.2016.1237532

Maguire, B and Woods, J. (2020). The game of belief. *Philosophical Review, 129*, 211-249. doi.org/10.1215/00318108-8012843

Mantel, Susanne (2019). Do epistemic reasons bear on the ought simpliciter? *Philosophical Issues, 29*, 214-227. doi.org/10.1111/phis.12149

McGrath, M. (2021). Being neutral: Agnosticism, inquiry and the suspension of judgment. *Noûs, 55*, 463-484.

McHugh, C. (2017). Attitudinal control. *Synthese, 194*, 2745-2762. doi.org/10.1007/s11229-014-0643-7

Nelson, M. T. (2010). We have no positive epistemic duties. *Mind, 119*(473), 83-102. doi.org/10.1093/mind/fzp148

Pittard, J. (2023). The many 'oughts' of deliberation. *Philosophical Studies*, *180*, 2617-2637. doi.org/10.1007/s11098-023-01973-z

Pryor, J. (Ms). Doxastic skepticism.

Quanbeck, Z & Worsnip, A. (forthcoming). A permissivist alternative to encroachment. *Philosophers' Imprint*.

Quilty-Dunn, J. and Mandelbaum, E. (2018). Against dispositionalism: belief in cognitive science. *Philosophical Studies*, *175*, 2353-2372.

Reisner, A. (2008). Weighing pragmatic and evidential reasons for belief. *Philosophical Studies, 138*, 17-27. doi.org/10.1007/s11098-006-0007-3

Rinard, S. (2019). Believing for practical reasons. *Noûs, 53*, 763-784.

Ross, J. and Schroeder, M. (2014). Belief, credence, and pragmatic encroachment. *Philosophy and Phenomenological Research, 88*, 259-288.

Ryle, G. (1949). *The Concept of Mind.* London: Hutchinson.

Schwarz, W. (2020). Ability and possibility. *Philosophers' Imprint, 20*.

Schwitzgebel, E. (2002). A phenomenal, dispositional account of belief. *Noûs, 36*, 249-275. doi.org/10.1111/1468-0068.00370=

Schwitzgebel, E. (2023). Belief. In E. Zalta & U. Nodelman (Eds.) *Stanford Encyclopedia of Philosophy* (Spring 2024 Edition). URL = <https://plato.stanford.edu/archives/spr2024/entries/belief/>.

Simler, K. (2016). Crony beliefs. *Melting Asphalt.* Available at: https://meltingasphalt.com/crony-beliefs/

Singh, K. (forthcoming). Belief as commitment to the truth. In E. Schwitzgebel and J. Jong (Eds.), *The Nature of Belief*. Oxford: Oxford University Press.

Sperber, D., and Wilson, D. (1995). *Relevance: Communication and cognition* (2nd ed.). Blackwell Publishing.

Soter, L. K. (2023). Acceptance and the ethics of belief. *Philosophical Studies, 180*, 2213-2243. doi.org/10.1007/s11098-023-01963-1

Soter, L. K. (forthcoming). A defense of back-end doxastic voluntarism. *Noûs.* doi.org/10.1111/nous.12501

Way, J. (2012). Transmission and the Wrong Kind of Reason. *Ethics, 122,* 489-515. doi.org/10.1086/664749

Weisberg, J. (2020). Belief in psyontology. *Philosophers' Imprint, 20*.

Westra, E. (2023). Symbolic belief in social cognition. *Philosophical Perspectives, 37*, 388-408. doi.org/10.1111/phpe.12196

Williams, B. (1981). Internal and external reasons. In B. Williams, *Moral Luck* (101-113). Cambridge University Press.

Williamson, T. (forthcoming). Knowledge, credence, and the strength of nelief. In A. Flowerree and B. Reed (Eds.), *Expansive Epistemology: Norms, Action, and the Social World*, London: Routledge.

Worsnip, A (2021). Can pragmatists be moderate? *Philosophy and Phenomenological Research, 102*, 531-558. doi.org/10.1111/phpr.12673