

## Acceptance and the Ethics of Belief

Laura K. Soter  
Duke University

Forthcoming in *Philosophical Studies*; Penultimate draft—please cite published version

### Abstract

Various philosophers authors have argued—on the basis of powerful examples—that we can have compelling moral or practical reasons to believe, even when the evidence suggests otherwise. This paper explores an alternative story, which still aims to respect widely shared intuitions about the motivating examples. Specifically, the paper proposes that what is at stake in these cases is not belief, but rather *acceptance*—an attitude classically characterized as taking a proposition as a premise in practical deliberation and action. I suggest that acceptance’s theoretical usefulness in the ethics of belief has been hidden by its psychological obscurity. I thus aim to develop an empirically adequate and mechanistically specific psychological profile of acceptance. I characterize acceptance as centrally involving a cognitive gating function, in which we prevent a target belief state from having its characteristic downstream effects on reasoning, cognition, and action, and restructure those downstream processes. I then argue that there is substantial empirical support for the existence of the cognitive mechanisms needed to instantiate this view, coming from the science of emotion regulation. I argue that acceptance involves deploying the same mechanisms used in emotional response modulation to belief states: acceptance is *doxastic response modulation*. I then propose that having a better understanding of the psychological profile of acceptance leaves us better positioned to appreciate its potential usefulness for making progress on various puzzles within the ethics of belief.

### 1. Introduction: A Puzzle about the Ethics of Belief

We sometimes seem to have compelling moral or practical reasons to believe something unsupported by our evidence. An athlete might perform better if she believes she will defy the stacked odds and win her race;<sup>1</sup> someone might feel morally compelled to believe a friend’s innocence despite evidence of his guilt;<sup>2</sup> or one might regard belief in a social stereotype as evidentially justified but

---

**Acknowledgements.** I am grateful to Peter Railton, Chandra Sripada, Renée Jorgensen, and Maegan Fairchild for detailed discussion and comments on multiple versions of this paper. Additional thanks are owed to Mica Rapstine, Adam Waggoner, Brian Weatherson, Corey Cusimano, Sarah Buss, Alex Madva, Susan Gelman, Ethan Kross, Zach Barnett, Gabrielle Kerbel, and Matt Stichter for their feedback or discussion, to Aliosha Barranco Lopez, Henry Schiller, Tez Clark, and Caitlin Mace for excellent conference comments, and to two anonymous referees for their constructive feedback. Versions of this paper were presented at the Michigan Graduate Student Working Group, the Michigan Candidacy Seminar, the 2021 Princeton-Michigan Metanormativity Workshop, the 2021 Austin Graduate Ethics and Normativity Talks, the 2022 Southern Society for Philosophy and Psychology, the 2022 Pacific American Philosophical Association Meeting, Athena in Action 2022, and the 2022 Moral Psychology Research Group at Cornell; thanks to those audiences for their engagement and discussion.

<sup>1</sup> Or perhaps a sick person is more likely to recover if they believe against the evidence about their chances (e.g., Rinard 2015; Reisner 2008).

<sup>2</sup> Stroud (2006) and Keller (2004; 2018).

nevertheless morally undesirable.<sup>3</sup> The apparent tension between practical and evidential reasons for belief crops up in a range of philosophical debates.<sup>4</sup> In each, we must grapple with a central theoretical challenge: on the standard picture of belief, we cannot choose to believe something unsupported by our evidence for practical or moral reasons. Rather, belief is “involuntary”—a psychologically immediate reaction to our evidence. The specter of doxastic involuntarism hovers in the background of many puzzles in the ethics of belief: how can we make sense of moral or pragmatic duties, responsibilities, obligations, or motivations of belief (henceforth, I will refer to doxastic “duties” for simplicity), if we lack the requisite ability to believe in response to these non-evidential reasons?<sup>5</sup>

Two overarching strategies are available for trying to preserve a robust ethics and pragmatics of belief.<sup>6</sup> We might find a way to make belief itself do the necessary work—for instance, by allowing the threshold for (justified or rational) belief to vary with the moral stakes (the strategy favored by some encroachment theorists), by explaining away any apparent conflict between morality and rationality,<sup>7</sup> or by adopting a view of belief that involves an active psychological commitment.<sup>8</sup> Alternatively: we could instead find a belief-adjacent attitude that plays the role of allowing us to respond doxastically to non-evidential reasons. Such an attitude must be sufficiently belief-like to earn

---

<sup>3</sup> Basu (2019b) and Begby (2013; 2021).

<sup>4</sup> Including, though certainly not limited to: the debates around moral and pragmatic encroachment (see Jorgensen Bolinger (2020) for a thorough overview), doxastic wrongdoing (Basu and Schroeder 2019; Basu 2018), epistemic partiality (e.g., Stroud 2006; Keller 2004; Arpaly and Brinkerhoff 2018; Kawall 2013), and how we should believe in and about others and ourselves (e.g., Morton and Paul 2019; Paul and Morton 2018).

<sup>5</sup> Various authors have argued that we can make sense of epistemic deontology in the face of involuntarism: e.g., see Hieronymi (2008; 2006), Shah (2002), Chrisman (2008), Weatherson (2008), Steup (2012), Flowerree (2017), among others. However, these defenses of epistemic agency generally do not claim that we have the control to believe for non-epistemic reasons. Though see Jackson (2021) and Roeber (2019; 2020) for discussion of the latter.

<sup>6</sup> Some deny that we ought to pursue an ethics of belief at all, rejecting the idea that we ought not have certain kinds of “inappropriate” beliefs identified by those in the literature; see, for example, Enoch and Spectre (forthcoming) and Sher (2019). Note that even those who reject that we ought to think our beliefs are sensitive to *moral* evaluation, we might still think that there can be cases in which it would be *practically* beneficial for an agent to believe against her evidence (for instance, that she will succeed even when the odds look slim)—thus, even those who want to reject a morality of belief might still have interest in the question of belief against the evidence more broadly.

<sup>7</sup> E.g., Brinkerhoff (2021) does this with some prejudiced beliefs.

<sup>8</sup> Something like this “active endorsement” picture of belief seems to be at work in much of Basu’s discussions of doxastic wrongdoing, for instance (e.g., Basu 2018), though she does not explicitly defend such a view there (though see Basu 2022 for some explicit discussion of this). See also McKaughan (2007) for an overview of this “active endorsement” account.

its proper place within the ethics of belief while still being legitimately responsive to moral and pragmatic reasons, capable of bearing the weight of the ethical and pragmatic puzzles at hand (that is, recasting these issues in terms of the alternative attitude must plausibly capture our doxastic duties in morally significant cases), and it must be *psychologically realistic*—i.e., plausibly part of our actual cognitive economies.

In this paper, I defend a candidate for the second route: *acceptance*. I am specifically interested in acceptance characterized as an alternative to belief: we can accept a proposition *that we do not believe* when we want to think, reason, and act in accordance with it (processes that would normally rely on belief by psychological default). I anchor my account in the existing literature,<sup>9</sup> which historically distinguishes acceptance as more clearly under our voluntary control than belief. But the picture I propose goes substantively beyond existing accounts, by paying special attention to developing a *mechanistically precise and empirically adequate psychological profile of acceptance*.

We can thus identify two overarching tasks that the acceptance account will have to accomplish, if it is to be successful:<sup>10</sup>

- (1) **The Empirical Project:** Provide a psychological profile of acceptance that is mechanistically specific and empirically adequate: that is, that details how the process of acceptance could plausibly be realized in human minds, consistent with functions understood by psychology and cognitive science.
- (2) **The Normative Project:** Show that acceptance can play a prominent role in the ethics of belief: that apparent moral/practical duties of belief can be productively recast as duties of acceptance.

---

<sup>9</sup> Most prominently, this distinction has been developed by Cohen (1989; 1992), Engel (1998), and Bratman (1992) in epistemology (see also Frankish (2007b) and Van Fraassen (1985) for related discussion).

<sup>10</sup> Thanks to Tez Clark and two anonymous reviewers for encouraging me to be more explicit about this division.

The primary focus of this paper will be the empirical project: understanding what goes on operationally in the mind of an accepting agent. The secondary goal will be to suggest that progress on the empirical project is integral to the normative project: that a psychologically rigorous understanding of acceptance will better position us to evaluate acceptance's contributions to the ethics of belief. I will not in this paper address whether my account of acceptance can settle any particular debate within the ethics of belief; instead, I will simply suggest that there are many features of acceptance that make it a promising candidate to help untangle various puzzles in this domain.

This paper proceeds in four stages. I start by highlighting the features of belief that make the idea of moral or pragmatic doxastic duties so tricky, then turn to existing accounts of acceptance as an alternative (§2). I suggest that acceptance is a promising candidate for projects in the ethics of belief, but that we need to better understand the psychological profile of acceptance before we can evaluate its normative promise. I undertake this project, first proposing an operational account of acceptance as centrally involving a *cognitive gating function* (§3). I then offer empirical support for the existence of the mechanisms needed to instantiate this gating function by looking to the science of emotion regulation, ultimately leading to my characterization of acceptance as *doxastic response modulation* (§4). Finally, I return to the normative project, arguing that acceptance holds significant potential for puzzles in the ethics of belief (§5), before concluding (§6).

## **2. Belief and Acceptance**

There are several features that make belief troublesome for practical and moral evaluation. Of course, characterizing belief is not a trivial philosophical task. Here I highlight features that I take to be both reasonably widely accepted by philosophers, and characteristic of the classic evidentialist/involuntarist framework, so as to be a defensible starting point.

### **2.1. A Basic Account of Belief**

## ACCEPTANCE AND THE ETHICS OF BELIEF

Most basically, belief is the propositional attitude of taking some proposition  $p$  to be true, in response to the truth of  $p$  (or a degree of confidence that  $p$  is true).<sup>11</sup> Beliefs operate, functionally, as our *default cognitive background* in deliberation and planning: they provide us with a representation of how the world is and guide us through it.<sup>12</sup> They serve as the basis for inference, reasoning, deliberation, intention, and action—and do so non-inferentially, without requiring our conscious supervision or intervention.<sup>13</sup>

To usefully serve as our default cognitive background, beliefs must be responsive to and reflective of how the world is.<sup>14</sup> Beliefs are thus characteristically evidence-responsive: they change when we encounter new information that bears on the likely truth or falsity of the proposition in question.<sup>15</sup> We need not oversee the formation or updating of our beliefs in response to evidence; rather, beliefs are in general spontaneously evidence- and experience-sensitive—they form more or less automatically in response to the evidence we encounter.<sup>16</sup> Further, we are often rather good at forming accurate beliefs in response to the world's evidence. Under normal circumstances, and given

---

<sup>11</sup> By “in response to the truth of  $p$ ,” I mean something akin to Shah and Velleman (2005)’s discussion: for an agent to believe  $p$ , she must conceive of it as being regulated for truth, and subject to a normative standard according to which the state is correct if and only if it is true. Beliefs can, of course, be false—but an agent cannot take her belief to be false and still believe it (e.g., see Frankish 2007a). This truth-responsiveness is what distinguishes belief from other kinds of attitudes that also involve some version of “taking  $p$  as true,” such as imagining.

<sup>12</sup> I borrow and elaborate on this characterization from Bratman (1992), integrating it in particular with ideas from Railton (2014).

<sup>13</sup> This is crucial for understanding the role of belief in our cognitive lives. The set of things we believe (at least implicitly, though certainly not occurrently) is indefinitely large, and in practical deliberation and reasoning we rely on many beliefs as background premises without consciously entertaining all of them. We could never navigate the world if we had to consider every belief we relied on; we simply lack the time and cognitive resources this would demand. So while we might be able to make many of our beliefs explicit and occurrent, we needn’t do so in order to rely on them.

<sup>14</sup> This is not to say needing to navigate the world will always result in precisely accurate beliefs, especially regarding matters that have little impact on our practical lives. Nonetheless, our need to navigate the world is closely tied to our ability to represent it accurately, and so the connection between belief’s truth-responsiveness and this navigation is important.

<sup>15</sup> E.g., see Shah (2003); Shah and Velleman (2005).

<sup>16</sup> Of course, we do sometimes step in and deliberate about what to believe in response to a complex body of evidence, but such cases actually represent a very small portion of our total belief-formation experiences. Further, this is not to say that rational belief formation never goes wrong; rationalization, perceptual distortion, and delusions might all be examples of cases in which we fail to form beliefs that are spontaneously and accurately responsive to our what our evidence actually justifies. All that is needed for present purposes is that belief-formation is responsive (at least normally/often) to what *we take* the evidence to point to. This is compatible with there being substantive questions about cases in which what we take the evidence to justify is different from what it actually justifies.

## ACCEPTANCE AND THE ETHICS OF BELIEF

a reasonably cooperative and accurate information environment, our belief-forming mechanisms tend to deliver us rather reliable results (especially concerning the local, observable world). This is a deeply important feature of our cognitive systems: our beliefs would be a poor guide to navigating the world if they systematically failed to represent the world accurately.<sup>17</sup>

This framework highlights why conceiving of an ethics of belief is so notoriously tricky. First, belief is thought to be, on the standard evidentialist picture, rationally and psychologically determined by our evidence: information that bears on the likely truth or falsity of  $p$ . Pragmatic and moral considerations do not themselves give us justification to think something is (un)likely to be true. Thus, most philosophers maintain that pragmatic and moral reasons cannot rationally be direct reasons for belief as such. Further, because belief is (either conceptually or psychologically) constrained by evidence and spontaneously evidence-responsive (when all goes well), belief is not under our direct voluntary control: we cannot choose to believe for non-evidential reasons.<sup>18</sup> Herein lies the central problem: insofar as we accept some kind of “ought implies can” principle, it becomes quite difficult to explain how we could have doxastic duties to believe without or against evidence—even in cases where we appear to have compelling practical or urgent moral reason for such a belief—if we lack the doxastic capacity to choose to respond to such reasons.

We could try to escape this tension by rejecting this evidentialist, involuntarist picture—though it is foundational enough that many would find this unattractive. Alternatively, we could focus

---

<sup>17</sup> It is plausible that in often, the domains where people systematically tend to form false beliefs involve epistemic environments that are deficient or distorted in some significant way (perhaps in conjunction with pernicious or disordered motivational factors). When beliefs are false, it often becomes difficult to reason and act on the basis of those beliefs, as the world will continue to push back on the believer. The well-functioning, non-disordered cognitive agent will often find their minds changed by the world, in the end. Of course, this is not true in every case, and there certainly are people who manage to maintain highly unjustified beliefs in the face of significant counterevidence. It may also be easier to cling to such beliefs when they are more abstract and less amenable to everyday evidential support or lack thereof—e.g., the average person’s beliefs about the origins of the universe will receive far less pushback than their beliefs about their singing abilities. Yet we should be wary of focusing too much on these cases at the expense of realizing just how well our belief systems do in general adapt to the evidence they are given.

<sup>18</sup> See Alston (1988) and Williams (1970) for foundational defenses of this claim.

on “indirect” strategies of belief manipulation (e.g., seeking out or avoiding sources of evidence), theorize about amount of evidence needed for a morally significant belief to count as rational or justified, or try to explain away any apparent conflicts between epistemic rationality and morality or prudence. But each of these faces problems: the first of these may be unreliable (if the evidence for the undesired belief is strong, these attempts may be both unsuccessful and epistemically suspect) or self-undermining; the second offers little action-guidance to an agent caught in a moral-doxastic dilemma and wades us into the murky waters of encroachment; and the last seems unlikely to succeed to dispel every possible case, and will be unlikely to satisfy anyone who thinks that it is at least in principle possible for there to be genuine conflicts between rationality and morality or prudence.<sup>19</sup>

Happily, there is another available strategy: we can look for a belief-adjacent attitude that might allow us to be legitimately doxastically responsive to non-evidential reasons. A promising candidate for such an attitude comes from the distinction that has previously been drawn between belief and *acceptance*.<sup>20</sup>

## 2.2 A Promising Alternative: Acceptance

Philosophers have sometimes appealed to acceptance in cases where we want to reason and act on the basis of a proposition we do not believe—when we have reason to prevent our beliefs from playing their default guiding role. Bratman (1992) characterizes acceptance as what we take for granted in a practical context; similarly, Cohen describes accepting as “to have or adopt a policy of deeming, positing, or postulating that p... as a premise in some or all contexts” (1989, p. 368). While we may in many cases believe the propositions we use as premises in practical deliberation, *we need not*—we

---

<sup>19</sup> See Traldi (2022) for an argument about why we should be skeptical of claims that these kinds of norms can never conflict.

<sup>20</sup> My discussion will focus specifically on the *epistemic* conception of acceptance. As Fleisher (2018, p. 2652 fn4) and McKaughan (2007) note, there are a number of other kinds of acceptance discussed in the philosophical literature, including in philosophy of language, philosophy of science, and literature on metacognition. There may be some systematic differences in how acceptance is conceptualized across domains.

can accept something we do not strictly speaking believe. Acceptance thus allows us to intervene on the operations of that default cognitive background of belief.

Acceptance is a propositional attitude: to accept  $p$  is to take  $p$  as a premise in practical reasoning and action. Acceptance differs from belief in several key ways. Perhaps most centrally, acceptance is said to be under our direct voluntary control: we can choose to accept  $p$ , and this is a conscious mental act (Bratman 1992; Cohen 1989, 1992). Accordingly, acceptance lacks belief's distinctive connection to evidence: acceptance does not entail any specific confidence in the likely truth of  $p$ . Further, we can accept directly for practical or moral reasons—and we can accept either in general or in specific contexts (in contrast to belief, which is characteristically context-independent<sup>21</sup>). For example, a lawyer might accept that her client is innocent on professional grounds while trying her case, even if she believes based on the evidence that her client is guilty.

This existing characterization gives us a high-level account of acceptance: a propositional attitude of relying on  $p$  in practical inference and planning—even though  $p$  is not believed—which is under voluntary control, can be responsive to pragmatic and moral considerations, is not limited to evidence-dependent rational norms, and can be selectively deployed in specific contexts or in general. The properties I have highlighted enjoy general consensus from the major authors who have written on acceptance (e.g., Bratman 1992; Cohen 1989, 1992; Engel 1998, Van Fraassen, 1985 among others), despite some important differences in their accounts.

One difference between existing accounts is worth discussing. Some accounts of acceptance, such as Cohen's (1989; 1992), concern *all* cases of taking  $p$  as a premise in practical deliberation,

---

<sup>21</sup> It would be both irrational and highly psychologically odd for one to believe (for example) that Mercury is the closest planet to the sun on Tuesdays, but not to believe this on Thursdays. The same problem does not arise for acceptance: a lawyer might accept her client's innocence in the courtroom, but not at brunch with her friends. Notable dissent to the context-generalness of belief comes from philosophers who argue that belief is “fragmented” (see Elga and Rayo 2021; Egan 2008; Bendaña and Mandelbaum 2021 for recent discussions); though I do not take it that these authors think that the context-dependency of belief is *volitional*. The view of acceptance I develop in this paper is compatible with fragmented belief storage accounts; it will just turn out that there are interesting questions about how precisely the mechanisms I discuss below interact with belief—and this will be different across different accounts of belief architecture.



*regardless* of whether  $p$  is believed by the agent. Though Cohen frequently discusses cases where acceptance and belief come apart to illustrate their different properties, there is not a fundamental difference on his account between cases where acceptance and belief diverge, and where they agree. In all cases, whatever we use in practical reasoning is what we accept; what we believe is an orthogonal question.<sup>22</sup>

Bratman's account, in contrast, diagnoses a substantive difference between cases where acceptance and belief converge and diverge. Because Bratman characterizes belief as our default cognitive background, acceptance is the mental act that of adjusting, or intervening on the operation of, this default background; he thus captures that there is something psychologically unique about relying on something *other* than belief. From a psychological perspective, this is the more useful starting point: we don't need a new account of how reasoning and action are guided when we do not wish to accept against or beyond our evidence, because we already have a perfectly well-functioning cognitive construct to play this role: belief. It is specifically the cases where we choose to adjust or intervene on the operation of that default cognitive background that need elucidation. Accordingly, I follow Bratman's account: we accept when we reason and act on something *other* than what we believe—when we prevent some belief from playing its usual role in deliberation and action.

### 2.3 Is Acceptance up to the Task?

Our hope is that acceptance might offer an alternative, belief-adjacent attitude that can make sense of the strong intuition that there are cases where we can and ought to be doxastically responsive to non-evidential reasons. Is it up to this task?

---

<sup>22</sup> In a similar vein, Stalnaker (1984) characterizes acceptance as treating a proposition as true—but takes this to be a broad category that includes belief as sub-kind, along with presupposition, postulation, assumption, and other nearby attitudes. In virtue of this breadth, Stalnaker's notion of acceptance is sufficiently different from the kind I am concerned with that I will not discuss it further here. Van Fraassen's notion of acceptance may also be in a rather different class than some of the others listed above, insofar as a scientist who accepts in his sense can (and should) still believe her hypothesis is empirically adequate.

Two desiderata are the attitude is sufficiently belief-like to belong within the ethics of belief, but also that it can be responsive to moral and pragmatic reasons. The high-level account meets these criteria: we can choose to accept something we do not believe directly on the basis of moral and practical reasons. This is a promising start.

But there remain two further desired features. First, we want to ensure that the candidate attitude can bear the weight of the puzzles at issue: that acceptance, as opposed to belief, could plausibly fulfill our apparent doxastic duties in morally important circumstances. For instance, some have argued that in a structurally unjust world, agents can sometimes find themselves with bodies of evidence that rationally justify prejudiced beliefs.<sup>23</sup> If our proposal will be that the agent who takes such a belief to be rationally justified but morally undesirable ought to *accept*, acceptance had better be a sufficiently robust attitude to fill these moral shoes. That is, we will have to explain why an agent who accepts some egalitarian claim—even if she doesn't *strictly speaking believe it*—nonetheless fulfills whatever moral duties she has to avoid prejudicial thinking. This is (part of) the normative project identified in the introduction.

Second, the attitude should be psychologically realistic, and our account empirically adequate. We ought to have good reason to suspect that the attitude we propose is really part of the cognitive systems of agents like us, and that we can and do use it in the relevant contexts, along with a clear understanding of what goes on in the mind of an accepting agent. This is the empirical project.

The existing high-level account of acceptance does not, as it stands, clearly satisfy these two goals. Armed only with the descriptions of “departing from our default cognitive background” and “taking a proposition as a premise in reasoning and action,” we lack a rigorous understanding of the psychological profile of acceptance and the cognitive mechanisms used in its deployment. In the absence of such an account, we also lack a precise understanding of what it actually takes for agents

---

<sup>23</sup> See Begby (2013), Basu (2019), among others, for discussion of such cases.

like us to accept, and what differentiates acceptance from other nearby attitudes like supposition. Beyond being of independent interest, lacking this profile leaves us ill-equipped to judge whether acceptance can satisfy the kinds of doxastic duties of concern in the ethics of belief: if we don't understand what it really looks like, psychologically, for an agent to accept, then we cannot hope to adjudicate whether acceptance can play the normative role that philosophers have previously ascribed to belief. In other words: I do not think we can tackle the normative project without first making progress on the empirical project.

Here, I take up the project of developing a profile of acceptance that is mechanistically detailed and empirically plausible. Once we have this in hand, we can return to the question of whether acceptance can indeed fill the role we need it to in the ethics of belief.

### **3. An Operational Profile: The Gating View of Acceptance**

So, what kinds of psychological operations are involved when an agent accepts?

To accept, in the sense developed here, is to intervene on the default cognitive background of belief, preventing a belief (or set of beliefs) from playing its usual (spontaneous, non-inferential) guiding role in reasoning and action. We do this when we have some reason not to rely on this default background: when there is some target underlying belief state (or set of beliefs) which doesn't serve our moral or practical aims, and which we therefore want to override. (Going forward, I will talk of a "target belief" for ease of reference, though it is likely that there is often really a set of related beliefs targeted.)

Acceptance must thus involve identifying where the target belief is activated and operating in deliberation and planning—requiring us to engage in a cognitive monitoring operation. We must then work to block any of the usual inferences, actions, patterns of thought and reasoning, and other default downstream effects caused or licensed by that belief state—and (in some cases, more on this later)

redirect those patterns of thought, reasoning, judgment, and action in accordance with an alternative proposition—updating any inferences, plans, or decisions that depended on that belief accordingly.

In limited contexts, this will not be so difficult. Accepting that someone will keep their promise and meet me for lunch on time—against my evidentially justified belief that they probably will not, given their prior record regarding punctuality—might be fairly straightforward: I might arrive a few minutes early, not make backup lunch plans, and prevent myself from thinking about all the prior times they've been late for our meetings. It will be more difficult to accept in a more context-general way, in part due to the lack of a clearly delineated domain in which the proposition will be relevant and involved in our deliberations. Imagine a parent who believes his child is lying about something significant, but decides to accept what they say—not just in some specific context, but in general. If they are heavily involved in each other's lives, this belief is likely to be activated in a diverse set of inferences, motivations, patterns of reasoning, and decisions. Thus, the parent will have to *monitor* his cognition for all the places where the unwanted belief is likely to be activated in reasoning, decision-making, and other cognitive processes (e.g., attention and motivation). Then, having identified these instances, he must *intervene* to prevent their belief from playing its usual functional role as the default basis for reasoning and action, redirecting their responses appropriately.

We can thus understand the heart of this profile of acceptance as centrally involving a kind of *cognitive gating operation*. This process intervenes to block a particular belief state from having its usual downstream effects and playing its usual role across deliberation, inference, reasoning, and action. Call this the Gating View of acceptance:

GATING VIEW OF ACCEPTANCE: Acceptance, understood as departing from our default cognitive background of belief, centrally involves preventing some target belief state from having its usual downstream effects in deliberation, inference, reasoning, cognition, and action

via a cognitive gating operation—and restructuring those processes in accordance with some accepted proposition.

This gating results in changing the agent’s inferential and deliberative landscape, by blocking certain inferences, actions, or patterns of thinking that would be licensed by their beliefs, and blocking certain inhibitions imposed by those beliefs. Thus, acceptance both prohibits some options that would be available to the agent in virtue of what she believes, but also permits moves that would not be available under belief’s guidance.

### 3.1. Psychological Features of the Gating Operation

The Gating View is consistent with the existing high-level account of acceptance provided by Bratman and others. But fleshing out this operational profile already gives us new insight into key psychological features of acceptance that were not obvious with only the high-level account in hand.

First, acceptance is *cognitively effortful*. One must first be vigilant for all the various places where the target belief(s) may be involved in reasoning, cognition, and action, and intervene to block it from playing its usual role. To do this fully would involve not only monitoring the inferences and decisions that may depend on that belief state, but also anticipating the ways it may affect other cognitive functions—such as motivations, desires, attentional patterns—and updating/redirecting all of those processes accordingly. These operations will involve executive processes which require significant mental engagement and resources (of which we have a notoriously limited supply); this is consistent with well-established evidence about monitoring and intervention processes across cognitive psychology.<sup>24</sup> Thus, we can easily imagine that one might be prone to accidentally relying on their

---

<sup>24</sup> The idea that tasks involving monitoring and intervening on default cognitive processes is effortful appears in many domains in cognitive science. For a few well-known examples: see Evans and Stanovich (2013), Evans (2008), and Evans (2019) for discussion of dual-process theories of cognition (I do not mean to commit to the “intuition vs. reasoning” framing that is sometimes associated with dual-process theories; just the familiar idea that the overriding of default psychological processes is controlled and effortful); see J.D. Cohen (2017) for an overview of the idea of overriding default responses using cognitive control mechanisms and Botvinick et al. (2001) for discussion of the monitoring function in cognitive control; see Devine (1989) and Payne (2005) for discussions of automatic and controlled components of

default belief in situations where she is distracted, or engaged in some other cognitively demanding tasks, and thus lacks the available cognitive resources to successfully monitor and intervene upon her default cognitive background.

Second, the Gating View reveals that acceptance is not a one-off action—we don't simply decide to accept and then it is done, our cognitive landscape instantly shifted, though classic accounts sometimes seem to characterize it this way. Rather, what it takes for an agent to accept is to deploy a series of specific mental control actions. Acceptance is thus more precisely understood as an attitude manifested through a *temporally extended sequence of specific (effortful) mental acts*, than as a single one-off action. Accepting is something we can choose to do, but so choosing involves committing to a pattern of acts over time. This invites us to ask about the extent to which an agent succeeds in accepting  $p$ , evaluating this in terms of how consistently successful she is at identifying contexts where the underlying belief is a premise in her reasoning processes, and how frequently she succeeds at overriding this default belief with the accepted proposition. We also begin to see room for an element of skill or habituation here—ideas that have not frequently been considered in the context of acceptance previously.<sup>25</sup> Importantly, this profile of active regulation is something that distinguishes acceptance from belief. In principle, an agent who believes  $p$  and an agent who is completely successful at accepting  $p$  (despite not believing it) could have identical patterns of reasoning, inference, action, and so on—but even if the two are modally coextensive, the cognitive action profiles of these agents will be radically different: while the believing agent will get that profile by default, the accepting agent will

---

stereotyping in social cognition; see Shenhav et al. (2017) for relevant discussion on mental effort. These are just some examples of an idea ubiquitous in many areas of cognitive psychology.

<sup>25</sup> There seem to be two distinct roles for skill: at the stage of identification of where the belief is involved in influencing our reasoning and behavior, and at the stage of intervening at this points once recognized. We might think that people could be differentially skilled at these two components, and perhaps it will turn out that in the case of belief, the former is particularly difficult (compared to emotion, for example). This idea deserves further exploration; thanks to Matt Stichter for encouraging me to think about it.

have had to put every piece in place intentionally. This is a nuance not usually discussed or captured by traditional accounts of acceptance.

These two features—cognitive effort and a pattern of mental acts—are characteristics of the processes involved in the gating operation. Next, let's develop the mechanistic profile a bit further.

### **3.2. Gating and Response Redirection**

The Gating View emphasizes that the blocking of the default guiding role of belief is a central functional component of acceptance. But there is one more piece to the story: in virtue of gating a belief's default downstream effects, those patterns of thought, reasoning, and action will need to be redirected and restructured. Precisely how this is done will depend on the agent's goals, motivations, and the rest of her epistemic landscape—and different approaches may lend themselves to different functional descriptions.

In some cases, the agent may seek to block a belief and fall back on their ex-ante epistemic background—that is, redirect the relevant downstream processes towards whatever they would look like had the target belief simply never been part of her cognitive landscape. These cases might be naturally described as a kind of “*mere gating*.” To build on an earlier example: consider a parent who accepts that his child didn't lie (or at least, accepts that the case has not been made that his child lied), despite evidence to the contrary. If his ex-ante epistemic background concerning his child generally supports the child being honest, reliable, etc., he might simply block the belief that she lied in this case, redirecting his patterns of thought, reasoning, and action towards whatever they would have been had he never received the unfavorable evidence regarding his child.

In other cases, the agent may not just redirect their cognitive landscape towards what it would have been without the belief in question; rather, they may want to redirect those processes towards some positive alternative (one which does not follow naturally from the ex-ante epistemic background). These cases might naturally be described as a kind of “*gating plus substitution*.” To give a

toy example, imagine you believe (based on your evidence) that Joe stole the money. You might not only want to block your belief that Joe stole the money (i.e., accept that he didn't do it)—but rather, positively accept that Amy stole it. In this case, you'll block your underlying belief in Joe's guilt—and rather than just restructuring the downstream landscape to reflect the absence of that belief, you do so in a way that reflects the commitment to Amy's guilt. Or, there might be an intermediate case of redirecting the downstream responses towards a kind of withholding judgment about  $p$  (you block your belief that the evidence points towards Joe's guilt and restructure your landscape to say: who knows who did it?).<sup>26</sup> This difference between mere gating and gating plus substitution is not meant to be categorical. Instead, the point is that we have latitude in how we restructure the downstream effects of the beliefs that we gate, and how we do that will shape how we naturally want to describe the agent's state of acceptance.

To summarize the basic idea: in any case of departing from the default cognitive background of belief, the agent gates her underlying belief state and restructures her default patterns of reasoning, cognition, and action. But there is a further choice point regarding how she restructures and redirects those responses. If she believes that  $p$  but blocks that belief, she might then redirect her downstream patterns of reasoning and action all the way to not- $p$ , towards some alternative  $q$ , or instead towards a kind of committed agnosticism or withholding judgment about whether  $p$ : not doing things that might normally be licensed by default by her belief that  $p$  (like closing inquiry into  $p$ , relying on  $p$  as an assumption in reasoning, and so on), but neither redirecting her responses all the way to reflect some commitment to reasoning and acting on the basis of not- $p$ .

There is undoubtedly more work to be done in understanding the precise nature of, and relationship between, the different ways an agent can redirect her downstream responses after she has

---

<sup>26</sup> In §5.2 I discuss the relationship between this and suspension of judgment.



gated a belief state; a full treatment of the nuances of this part of the view and the mechanisms involved will have to be saved for future work. The two basic points I want to draw out of this discussion are the following: first, that *gating* of the belief is a central functional component of acceptance, because acceptance involves preventing a belief from having its default psychological effects in guiding reasoning, thinking, and action. Second, that once a belief has been gated, the agent can redirect those processes in various ways—falling back on her ex-ante epistemic background or redirecting them towards some substantive alternative. The same structure underlies these varied cases: there is a wedge to be driven between the underlying belief state and the default downstream effects—but just how big that wedge is can vary (and we might be tempted to call more extreme cases substitution, and less extreme cases mere gating or blocking).

Finally, I'll highlight one more important feature of the Gating View: we can apply this gating operation to *whatever* the underlying belief state is, when we have reason to not want that underlying belief state to guide our reasoning and action. That is, this gating operation isn't limited to cases where we want to block a state of high confidence (what we might specifically call *belief*). Rather, we can also use this gating mechanism—and the associated redirection of responses—against states of low confidence or uncertainty. Belief states of any confidence level give rise to default patterns of reasoning, inference, deliberation, thought, and action—and so we can apply these mechanisms whenever there is a disconnect between the belief state formed in response to our evidence, and the ways that we want to restructure our patterns of thought, reasoning, and action. Thus, the mechanistic profile offered here has a sort of mix-and-match structure: there is a menu of options both in terms of what the nature of the underlying belief state to be gated is, and in terms of how the downstream processes are restructured. Particular combinations may lend themselves to different high-level descriptors (acceptance, blocking/suppression, or suspension [more on this later])—but the key insight is that there is a common operational structure underlying each of these.

#### 4. A Mechanistic Profile: Acceptance as Doxastic Response Modulation

The Gating View provides an initial operational profile of acceptance—one which accords with existing high-level descriptions, but which also provides us with new insight into features of acceptance that were not obvious with only the high-level account in hand. A natural next step is to ask about the view’s empirical plausibility. Do we have good reason to believe that this gating process is something that might actually be realized in human minds? Does the operational process presented above actually accord with any existing phenomena studied by psychologists, that could implement this profile in our cognitive systems?

The answer, I think, is yes. I propose that psychological research on emotion regulation provides support for the kinds of mechanisms at work in the Gating View, thus lending it strong empirical plausibility. Acceptance, I argue, can be understood as the doxastic analogue to emotional response modulation.

##### 4.1. Response Modulation in Emotion Regulation

A thriving body of psychological research explores emotion regulation: how people can and do control the experience and effects of (often maladaptive) emotions. Psychologists have identified a variety of emotion regulation strategies, which are often differentiated and categorized by the point in the emotion generation and experience process that they target. Most simply, we can distinguish between antecedent-focused and response-focused strategies (Gross 1998a).<sup>27</sup>

*Antecedent-focused responses* seek to affect or prevent the formation or elicitation of the target emotion. This includes strategies such as situation avoidance, distraction, and reappraisal techniques. In contrast, *response-focused strategies* seek to regulate the physical, verbal, behavioral, and cognitive

---

<sup>27</sup> Later accounts more precisely divide strategies into five categories: situation selection, situation modification, attentional deployment, cognitive restructuring, and response modulation (Gross 1998b; McRae 2016). Further, because emotions are temporally extended mental processes, the line between categories is in practice somewhat blurry. However, since I will be discussing only response modulation in detail, the coarse two-category distinction is sufficient for present purposes.

consequences characteristic of an emotion, after the emotion has already been elicited, and are thus also referred to as “response modulation” (e.g., Gross and Levenson 1993; see also Gross 1998b; Koole 2009; McRae 2016 for reviews). One especially well-studied kind of response modulation is expressive suppression, in which one tries to inhibit the facial, verbal, or bodily expression of a felt emotion—such as keeping a neutral face upon witnessing something disgusting or getting upsetting news, or resisting an urge to flee in fear or lash out in anger.<sup>28</sup> Successful suppression involves monitoring one’s verbal and behavioral responses for the characteristic expressions of felt emotions, and preventing oneself from enacting those responses (Giuliani et al. 2011).

These strategies can be used in specific one-off instances, such as when trying to suppress a yelp of fear when watching a scary movie, or repeatedly, as a long-term strategy for frequently arising emotions (Gross and John 2003). In both cases, these response modulation strategies are demanding on executive processes (Franchow and Suchy 2015, 2017; Gyurak et al. 2012; Niermeyer, Franchow, and Suchy 2016; Richards 2004; Richards and Gross 1999): it is cognitively effortful for people to deploy them over time, and they interfere with performance on other executively demanding tasks. Crucially (and in partial explanation of their characteristic effortfulness), these response modulation strategies do nothing to directly affect the initial generation of the emotion; their target is instead the downstream action tendencies and cognitive effects that characteristically flow from the emotion’s activation.

---

<sup>28</sup> Although regulation is often discussed in the context of negative emotions, people can regulate positive emotions as well. For instance, someone trying to keep a neutral face and hide excitement upon learning that they were accepted into a prestigious school, or stifling laughter in response to a funny video, are examples of expressive suppression for positive emotions (Gross and Levenson 1993).

#### 4.2. Doxastic Response Modulation

I propose that acceptance involves applying the cognitive mechanisms at work in emotional response modulation to belief states.<sup>29</sup> The response modulation process as developed for emotion regulation accords well with key components of the Gating View.

Most centrally, acceptance–like emotional suppression–does not directly seek to alter the underlying psychological base state (belief) itself. Rather, both seek to prevent the characteristic downstream consequences of the activation of the state in question. The Gating View characterizes acceptance as seeking to prevent the targeted belief from playing its usual role in deliberation, reasoning, and action. To do this is, essentially, to *suppress* the normal downstream effects of an activated belief—just as emotional suppression seeks to suppress the downstream effects of a felt emotion.

There are differences between the characteristic downstream profiles of belief and emotion, but both kinds of states can give rise to the same broad categories of downstream effects. Just as strong emotions often give rise to facial expressions, verbal reactions, and other physical action tendencies, so too do beliefs: for example, encountering evidence that is starkly incongruent with belief-generated expectations may lead to experiences and physical and verbal expressions of surprise. Successful acceptance—e.g., of a friend’s innocence, against a belief in their guilt—may involve preventing oneself from letting out a surprised “oh!”, raising one’s eyebrows, or doing a noticeable double-take when encountering new evidence that *does* seem to speak strongly in favor of their innocence. Just as suppressing a fear response may involve resisting the urge to turn and run, so too does accepting a friend’s innocence involve resisting belief-motivated action tendencies such as not trusting them (behaviorally or psychologically) with sensitive information.

---

<sup>29</sup> Those who endorse accounts on which belief is a form of confidence or trust, such as Railton (2014), McCormick (2022), and perhaps Schwitzgebel (2002), and insofar as confidence and trust are affective states, might simply say that acceptance is a specific kind of emotional response suppression.

Beyond these expression and action tendencies, the characteristic downstream profile of belief also includes a significant cognitive component as one of its primary functional roles: beliefs normally serve as our premises for reasoning and inference (our “default cognitive background”). Theorizing about acceptance thus requires us to focus on the suppression of these cognitive and inferential tendencies. On its face, this may appear to be a point of difference with emotional suppression, where empirical work has focused less on suppressing distinctively cognitive effects. Yet (for better or worse) emotions certainly do affect our reasoning processes. Although we might think of downstream cognitive effects as being more central to the role of belief than to emotion, it would be a mistake (and in conflict with decades of empirical research on how emotions affect reasoning) to suggest that emotions do not in fact have significant cognitive consequences.<sup>30</sup> Thus, fully suppressing an anger response surely involves not only preventing particular facial expressions, but also recognizing when our anger influences our reasoning processes—perhaps driving us to draw more negative conclusions about someone than we would if we were not angry.<sup>31</sup> While this kind of cognitive response may be quite difficult to prevent—and it may well be the case that brute suppression is a particularly ineffective strategy for controlling the downstream cognitive consequences of emotion—this difficulty does not

---

<sup>30</sup> This idea has been established in a variety of domains in psychology. One way it presents itself is that our emotional reactions to stimuli can affect how we reason about them. This focus has been especially central in the study of moral reasoning, where it has been shown that automatic or “intuitive” emotional reactions to stimuli can affect our moral judgments and decision-making (e.g., see Greene 2015; 2007; Greene et al. 2009; 2001; Haidt 2001 for classic discussion). Similar effects have been shown in other domains where emotional reactions can affect reasoning processes, such as framing effects and decisions involving perceived risk (classically, Kahneman and Tversky 1979; see also Keysar, Hayakawa, and An 2012; Costa et al. 2014). More generally, it is well-recognized that emotions, when activated, cause emotion-congruent biases across a range of cognitive mechanisms (e.g.m Brosch et al. 2013; Dolcos and Denkova 2014; Phan and Sripada 2013), including action and goal-selection mechanisms (emotions involve “action tendencies,” e.g., anger biases us towards retaliative goals; see Frijda 1987; Frijda, Kuipers, and ter Schure 1989; Scarantino 2014), attention mechanisms (e.g., fear causes us to be more sensitive to threat-related stimuli) (Domínguez-Borràs and Vuilleumier 2013), how we interpret new information, what and how we remember information, and so on. See Sripada (2021, section 4.2) for a helpful overview of this research framed for a philosophical audience.

<sup>31</sup> An example of empirical research on suppression of specifically cognitive effects of emotion is the suppression of emotion-laden or emotion-activated thoughts (e.g., Matos, Pinto-Gouveia, and Costa 2013; Muris et al. 1992; Roemer and Borkovec 1994; see also Mauss, Bunge, and Gross 2007 for some discussion of automatic suppression techniques in various domains).

itself signify a deep difference between the broad kinds of downstream effects that belief and emotion can have.

We thus have an initial case for a functional similarity between the primary function of emotional response modulation and acceptance: blocking the downstream effects of an activated state. The phenomenological profiles are similarly analogous. In suppressing both emotion and belief, there will be an experience of blocking, gating, working hard to prevent the usual responses from emerging—whether those are facial expressions, action choices, verbal reactions, or reasoning patterns. In the cases of suppressing cognitive and inferential responses, there will be an experience of trying to monitor and identify what parts of the reasoning processes have been influenced by the belief or emotion, and to adjust those accordingly. Further, the phenomenology of effort is prominent across both domains, and this accords well with the research cited above indicating that emotional response modulation is cognitively demanding.<sup>32</sup>

I'll highlight three further dimensions of similarity between emotional and doxastic response modulation. First, we noted that the Gating View allows us to break down the operational steps of acceptance: a basic gating mechanism, and then a question of how one restructures the downstream landscape (mere suppression, or redirection to some positive alternative). We can see that a similar set of options is conceptually available in the emotional domain. Sometimes, an agent merely wants to mask an emotional reaction and suppress the characteristic expressive responses, such as maintaining a neutral face when feeling disgust. Other times, she wants to actively replace the usual disgust response with an alternative response—such as a display of (apparent) enjoyment. Someone feeling disgust towards a friend's poor cooking may not only wish to show a neutral face, but to actively expressed gratitude for the food. Or someone who feels jealousy towards her friend regarding some achievement she herself was pursuing might merely suppress the jealousy response towards neutrality,

---

<sup>32</sup> See Shenhav et al. (2017) for a general discussion of mental effort.

or might redirect those downstream patterns all the way towards being reflective of a happiness for her friend. To (somewhat awkwardly) recast this process into a kind of emotional-acceptance language, we might distinguish between the agent “emotionally accepting” that she is *not jealous of* her friend and that she is *happy for* her friend: both involve the basic suppression mechanism, but they restructure the downstream landscape in different ways. These kinds of cases highlight the operational analogy between the domains of doxastic and emotional response modulation: there is a decision to be made regarding how (and how far) an agent modulates the default patterns of responses from the emotion/belief state.

Second, it is characteristic of both emotional response modulation and acceptance that our deployment of these strategies is not entirely limited by the features that rationally or psychologically constrain the formation of the underlying states. In other words: the very point of these regulation strategies is to provide agents with a way to override the default cognitive and behavioral roles of mental states that arise in response to particular kinds of stimuli. Our affective systems respond to certain kinds of stimuli—danger, injustice, impurity—with specific kinds of emotions—fear, anger, disgust. When these emotional processes are well-attuned, the activated emotions will be *fitting* to their objects; there will be certain perceived features that appropriately give rise to the emotions in question. Emotion regulation strategies allow agents to exert some control over the downstream effects of those emotions, when we have practical or moral reason to do so, even when the elicitation of the emotion was reasonable and fitting. Similarly with acceptance: doxastic response modulation gives agents a way to exert control over the effects of beliefs in our reasoning and deliberation, even when those beliefs are evidentially warranted and well-formed, when we have practical or moral reason to do so. In both domains, these response modulation processes allow agents to be responsive to practical reasons which would not themselves be the right kind of reason (descriptively or normatively) on which to

directly form the underlying base state.<sup>33</sup> Response modulation (emotional and doxastic) thus expands our agential capacities—giving us an additional kind of discretion over whether, when, and how we want these emotional and doxastic states to guide us—without sacrificing the proper functioning of the underlying state-formation system.<sup>34</sup>

Finally, in developing the Gating View we observed that accepting may be less difficult in a limited domain or timeframe, but can become quite challenging over a wide range of contexts or long period of time. This too tracks well with emotional response modulation. While it requires some cognitive effort to mask one’s anger or sadness in a one-off context, it becomes far more difficult to do this repeatedly over time and across a wider range of contexts. In the domain of emotion, relying on expressive suppression in daily life compromises general cognitive performance (Franchow and Suchy 2015), is associated with negative affect (Brans et al. 2013), and has even been thought to contribute to depression and anxiety disorders (Kashdan and Steger 2006; Sperberg and Stabb 1998; Campbell-Sills et al. 2006) and stress-related symptoms (Moore, Zoellner, and Mollenholt 2008). In the context of belief, it is not difficult to imagine that it is far easier to accept a proposition that is only relevant in a specific time or place (such as that it will rain this afternoon) than a chronically relevant one (such as that your friend is innocent, or that your cooking skills are superior)—and that the greater set of contexts in which one seeks to accept, the more difficult it may be. (Though this juxtaposes in interesting ways with the idea that these patterns of suppression could in principle become habitual and skilled with practice; disentangling these possibilities is a question for future work.) Further,

---

<sup>33</sup> For a recent discussion of various ways of understanding the talk of right and wrong kinds of reasons for belief formation, as well as a paper with a helpful overview of relevant literature, see Maguire and Woods (2020).

<sup>34</sup> This raises a question about how we ought to think about the rational or epistemic assessability of acceptance. Though I lack the space for a full treatment here, for now I propose that we should think of the decision to accept as a decision about the tradeoff between on the one hand your evidence, and on the other how you want to be and act in the world given our moral and practical motivations. We often have very good reason to be guided by our evidence and our beliefs—but not always. Decisions about whether to accept are thus cross-domain decisions between the epistemic and the moral/practical; like any cross-domain decision, both sets of norms are going to have some relevance, and neither be decisive. So it’s not the case that acceptance is *not* assessable according to epistemic norms—but it’s not assessable *only* against epistemic norms.



emotion regulation theorists generally agree that suppression is a rather maladaptive regulation strategy (Gross 1998b; Lynch et al. 2001), associated with various negative psychological and social effects (Richards and Gross 1999; Richards 2004; Butler et al. 2003). Regulation strategies that seek to avoid the elicitation of the emotion in the first place (e.g., reappraisal strategies) are more sustainable and beneficial in the long run (Gross 1998b; Richards and Gross 2000).

This too may hold a lesson for the doxastic domain: if you anticipate committing to accepting something in the long term and across contexts, it may be psychologically easier to try to get yourself to believe the proposition in question instead—by searching for new evidence or reinterpreting old, generating further explanations, etc.—or to simply avoid conditions that make the belief occurrent and activated. However, actual belief manipulation strategies face several challenges: they take time, they cannot be guaranteed to work (sometimes the evidence just isn't out there, and of course one of the starting premises of this present project is that we cannot always get ourselves to believe something merely out of prudence), and any such belief-manipulation strategies, even if ultimately successful, may be epistemically suspect. Thus, acceptance is a crucial belief-management strategy available to us in the meantime, or in situations when we do not actually want to interfere with our well-functioning belief-forming mechanisms. It's also worth noting that an agent who successfully accepts something over a long period of time may in fact end up altering their beliefs, without that being their intended goal. Insofar as acceptance will influence the agent's patterns of behavior, intention, inference, and patterns of thought and attention, the accepting agent is likely to alter her evidence-gathering and evaluating practices in a way that the non-accepting agent is not—thus delivering her a different body of evidence than the non-accepter. This means that long-term acceptance may still result in different patterns of beliefs. But importantly, direct belief manipulation is not the target goal of this process.

These similarities across the emotional and doxastic domains gives us compelling reason to think that we can posit response modulation as a plausible mechanism for acceptance. The profile of

response modulation aligns neatly with the key features of both the Gating View and the high-level existing accounts of acceptance, and also has the benefit of being a well-established and well-understood emotion regulation mechanism, thus lending it significant empirical plausibility—a crucial element for the present project of moving beyond high-level epistemological theorizing about acceptance and into the realm of empirically adequate philosophy of mind.

### 5. Acceptance and the Ethics of Belief: Turning to the Normative Project

We now have on the table proposed psychological profile of acceptance, including a clearer picture of what kinds of mental operations and cognitive mechanisms it involves. With this, we have tried to make progress on the empirical project for acceptance. At the outset, I suggested that this was an important prerequisite for progress on the normative project of understanding whether acceptance might indeed be helpful for addressing puzzles in the ethics of belief.

The remainder of the paper will consider the normative project in light of the psychological account developed thus far. Importantly, I will not in these pages seek to settle the normative question; in particular, the role acceptance might play in helping untangle issues in any specific debate (e.g., debates regarding epistemic partiality, profiling, grit, and so on) within the ethics of belief will need to be adjudicated elsewhere. My goal here will be more modest: to suggest that acceptance, in the sense developed above, is a highly promising candidate for making progress on these various issues (and so, worth pursuing further within these debates).<sup>35</sup>

---

<sup>35</sup> There already exist some attempts to integrate acceptance as a solution to some problems in the ethics of belief; one notable recent treatment comes from Renée Jorgensen (see Bolinger 2020b), who appeals to acceptance to make sense of what goes (rationally) wrong in (at least some cases of) racial/social group generalizations. The account of acceptance Jorgensen relies on is a bit different than the one I develop. For one, her discussion is pitched entirely at the level of epistemological theorizing rather than questions of psychological mechanisms (though on that front, I suspect that much of what each of us say is compatible with the other's account). However, Jorgensen specifies that on her account, accepting a proposition involves “taking it for granted” in a sense that it incompatible with thinking *p* is false (2020b, p. 2417 FN 3). But on my account, an agent who believes some proposition to be false can nonetheless prevent that belief from guiding her reasoning and action (though I take no explicit stand here on the *rational* status of so doing). Begby (2021, especially Ch. 9) also discusses acceptance in the context of the ethics of belief (and notes that his discussion is inspired by Jorgensen's, p. 161 FN 10). I agree with much of what Begby has to say, though the psychological profile developed here goes beyond his treatment; I thus think our approaches are complimentary.

### 5.1. Promising Features for the Normative Project

There are several features that make the normative project for acceptance look promising. First, there is the fact that acceptance is under our voluntary control: something we can choose to do in response to the kinds of moral and practical reasons that philosophers frequently take to motivate puzzles in the ethics of belief. Though this feature is not new to my account, it is a very attractive consideration in favor of acceptance, because recasting doxastic duties in terms of acceptance (as opposed to belief) allows us to avoid general worries about the viability of positing duties that can't be directly fulfilled on their own terms. On this topic, however, the psychological profile developed here adds important nuance that was lacking from standard accounts of acceptance: in particular, on the present account the control profile of acceptance is one of effortful regulation over time, rather than a one-off volitional action. This characterization may ultimately have implications for how we ought to think about instances of regulatory failure and blame, a topic that deserves further consideration elsewhere.<sup>36</sup>

The nature of this control profile highlights another normatively significant feature: because accepting involves effortful regulation of default psychological tendencies over time, acceptance turns out to be richly revealing of an agent's values and motivations. In deploying the doxastic response modulation mechanisms described above, the agent is restructuring her default patterns of reasoning, cognition, and action in a way that is expressive of her agency: it takes, and expresses, a serious kind of psychological commitment (and a deployment of mental effort) to continue to regulate the default effects of one's belief over time. Thus, accepting is agentially and morally significant—an attractive feature of the attitude if it is going to be deployed in morally weighty circumstances like those at stake in the ethics of belief.

---

<sup>36</sup> See Sripada (2018) for discussion of responsibility for effortful regulation in the context of addiction, as an example.

Additionally, the psychological profile developed here helps us appreciate that acceptance is a psychologically robust attitude: the commitment discussed above is manifested via the restructuring of diverse behavioral *and mental* processes. That the account of acceptance defended here is decidedly cognitive—involving the agent’s restructuring of her default patterns of reasoning and thinking as well as overt action—is a crucial feature of the view for many philosophers interested in developing an ethics of belief. A key commitment of many theorists in this domain is that our moral obligations can involve not only our outward behavior towards others but also our inner mental lives. Appealing to acceptance appreciates this emphasis and seeks to capture it, while striving to maintain a sensitivity to both lessons from both philosophical theorizing and psychological science about the limits and contours of our mental state control capacities.

## 5.2. Distinguishing Acceptance from Nearby Attitudes

Another benefit of developing the psychological profile of acceptance is that it allows us to more clearly see how acceptance relates to and differs from other nearby kinds of attitudes. This enables us to navigate the normative landscape with more precision and adjudicate between different candidate attitudes. I’ll consider here how the present account of acceptance allows us to better understand its relationship to three alternatives: “acting as if,” suspension of belief, and supposition.

**“Acting as If”:** A common and intuitive worry about appealing to acceptance (instead of belief) in the cases of interest is that acceptance is a thin or unsatisfying alternative to belief: a worry that the accepting agent is just “acting as if” or pretending she believes. Now, there is a grain of truth to this: by stipulation, on the account presented here, an agent accepts when she does not take her evidence to support the desired belief state. But whether the question of “isn’t the agent ‘just’ acting as if she believes” is worrisome for the account of acceptance depends on how thoroughly we flesh out the idea of “acting as if.” If what is meant by this is just acting in the sense of overt or public behavior, speech, and so on—then acceptance as developed here is clearly a different notion, for

precisely the reason noted above: the account of acceptance here is manifested not just through behavioral regulation but also through psychological and cognitive regulation. In other words, accepting isn't just about the way the agent presents herself to others, or the way she intervenes on her impulses right before the moment of action: it's also about how she decides to regulate her own internal patterns of thought, reasoning, planning, and other mental goings-on that are both upstream of and parallel to her actions. If someone who appeals to acting as if means to capture this robust psychological dimension as well, then saying that accepting is a way of "acting as if" no longer holds the same pejorative connotation that seems to be implied by more behavior-focused uses of that term.

**Suspension:** Another attitude of recent interest in discussions of both general epistemology and the ethics of belief is suspension of belief. Some have suggested that if there are moral considerations in play that are in tension with our evidence, rather than forming an affirmative belief we can *suspend judgment*, holding out on belief-formation longer than we would in less morally fraught cases.<sup>37</sup> How does my account of acceptance relate to such discussions of suspension? One immediate difference is that suspension only applies in cases where we want to withhold some belief, while acceptance can also handle cases where we want to redirect towards some positive alternative (e.g., an athlete who has reason to "believe" against the odds that she will succeed). But there is more to say. Though a full treatment of this question deserves more thorough discussion than I can devote here, we can make some initial progress by appreciating that the answer depends in large part on what one means by suspension of judgment. If suspension is understood as something like a first-order attitude of middling confidence, uncertainty, or insufficient evidence to say that one definitively either believes or disbelieves some proposition (what McGrath (2021) distinguishes as "agnoicism"), then this is very different from my account of acceptance, because there is no cognitive regulation at play. Rather,

---

<sup>37</sup> For recent discussions of this kind of view, see Jackson (2021), Roeber (2019), and Ross (2022), among others. For more general theoretical work on suspension, and for work showing very different accounts of the nature of suspension, see e.g., McGrath (2021), and also Friedman (2013; 2017), Masny (2020), Crawford (2022), and Staffel (2019).

this kind of uncertainty is precisely the kind of underlying belief state which has a set of default downstream effects in our reasoning, cognition, and action—the kind of belief state which can be gated if we are so motivated; i.e., if despite feeling uncertain towards a proposition we nonetheless want to block that uncertainty and redirect towards a commitment to a positive or negative attitude towards the proposition.

In contrast, suspension is sometimes discussed as a deliberate attitude an agent takes towards a proposition that can come apart from their evidential support. Here's a natural way to describe this: an agent has an evidentially formed state of somewhat high confidence in some proposition; this confidence would normally lead to a default pattern of effects in the agent's thinking and behavior. But, the agent has reason to want *not* to allow that belief state to guide her as it normally would; instead, she *chooses* to “suspend judgment,” to redirect her patterns of thinking and action as she would if she were more uncertain about the matter in question. If *this* is the kind of thing one means by suspension—wanting to prevent the default cognitive and behavioral effects of a belief state from having their default guiding role—then it turns out that the cognitive mechanisms involved in suspension are *the very same ones* that I have been discussing in this paper under the label of “acceptance.” This is an important insight, because it highlights that there is already a sensitivity in the existing literature to a need for the kinds of mechanisms that I'm aiming to spell out here. It also reveals that, although I am using “acceptance” as a kind of catch-all term here, what I am really interested in is the underlying mechanisms of preventing belief states from playing their default guiding role—and that depending on the precise combination of the underlying belief state and how an agent restructures those downstream processes, different labels (acceptance, suspension, etc.) may be more natural high-level descriptors (even if the underlying cognitive mechanisms are the same).

**Supposition:** I'll consider one final attitude: supposition. Though less frequently discussed in the ethics of belief, supposition is nonetheless worth discussing for the sake of understanding how

acceptance relates to other nearby attitudes.<sup>38</sup> The high-level goals of acceptance and supposition are plausibly quite different: supposition is characteristically a kind of hypothesis exploration, mental simulation, or (potentially counterfactual) investigation of “imagine if  $p$  were true.” We have characterized acceptance as involving a kind of cognitive commitment to reasoning and acting on the basis of  $p$ . Supposition need not involve this commitment, and is unlikely to be deployed over time and across contexts. Instead, its aim is narrower: we generally suppose for the specific purpose of figuring out what might follow if  $p$ . Because of these differing goals, supposition characteristically involves constraining or restructuring our reasoning and inferences, but not our actions, reactions, or any broader set of cognitive mechanisms such as attention and motivation. The fact that supposition does not require the agent to prevent all the various downstream responses of a target belief over an extended period of time, and instead merely consider what might rationally follow from the truth of the proposition in question, suggests that the Gating View of acceptance is likely more demanding than what is needed for supposition. A supposer needn’t monitor her cognition and action nearly as thoroughly as an accepter, and needn’t commit to a long-term policy of doing so; the supposer merely needs to restructure her inferences to explore what follows from  $p$ .<sup>39</sup>

These considerations suggest another dimension of difference, one which also helps clarify an important feature of acceptance: though acceptance lets us act and reason on the basis of propositions that are not supported by our evidence, our ability to successfully accept some  $p$  is not entirely unconstrained by our evidence. The more dramatic the departure from our overall evidence, the more

---

<sup>38</sup> In his (1992) discussion, Bratman also explicitly discusses how supposition differs from acceptance.

<sup>39</sup> These differences in high-level characteristics may also reflect differences in the lower-level psychological profiles of acceptance and supposition. A full treatment of the psychological profile of supposition is beyond the scope of this paper; but as a first pass, one could argue for a gating and response modulation account of supposition with a more limited target: the supposer only needs to gate and suppress the reasoning and inferential processes involving the target belief. Alternatively, perhaps supposition involves somewhat of a different cognitive profile, centrally involving processes of counterfactual reasoning, hypothetical simulation, and cognitive decoupling—and that these processes are more emphasized than the monitoring, gating, and suppression mechanisms that characterize acceptance. Such processes more closely align with the exploratory goals of supposition.

of our default cognitive background of belief we will have to monitor and intervene on, and the more difficult successfully accepting across contexts will be. We thus start to see the limits of our acceptance-regulatory processes, and appreciate that although acceptance is not rationally determined by the features that determine belief, it is *constrained by*, influenced by, and sensitive to them in important ways. Supposition, in contrast, seems not to face these same restrictions. We can suppose things that drastically conflict with our evidence and are wildly implausible—even though it would be extremely difficult to actually succeed at accepting these things. I can suppose that there is a hoard of angry elephants outside my apartment, or that the speed of light is a different constant; I can reason through what might follow if either of these things were true. Yet genuinely accepting these, in the sense that I have argued for in this paper, would be implausibly difficult. The fact that supposition has less practical aims and does not require me to act on its basis, and the fact that we tend to suppose in more limited contexts and timeframes, allows an agent to suppose a whole range of things that they may struggle to accept. This suggests that the more an agent seems to be *regulating*, committing to  $p$ , and using it to guide *action*, reaction, and a range of cognitive capacities such as attention in addition to inference, the more this will look like acceptance<sup>40</sup>. In contrast, to the degree that an agent has merely exploratory aims and is restructuring only their reasoning and inference patterns with the goal of figuring out what would follow if  $p$ , it may be more apt to describe them as supposing.<sup>41</sup>

---

<sup>40</sup> It may be difficult to know precisely where to draw the line between these attitudes in some cases, especially when describing the psychologies of other people—and some may resist making the distinction at all. But, for those who want to distinguish acceptance and supposition, the characteristic aims and psychological profiles seem promisingly different.

<sup>41</sup> One might be tempted to characterize acceptance as having specifically practical (broadly construed) aims, and supposition as having specifically epistemic aims. However, I think acceptance *can* be undertaken for specifically epistemic aims. That is: sometimes, being the most successful epistemic agent in the broad sense will involve responding not merely to the considerations of the evidence directly in front of us. A classic example is a scientist who favors a hypothesis for reasons of theoretical virtue that is less well empirically supported than some alternative: she might accept to forward her epistemic-scientific goals. (The kinds of cognitive regulation mechanisms I've argued for her may actually be stronger than what the scientist needs, though. A more appropriate attitude might be something like Fleisher (2018)'s *rational endorsement*, which focuses on broader norms of inquiry rather than a specific cognitive profile.) For another example, perhaps an agent who knows she'll be entering a deeply unreliable evidential situation thinks that her belief-forming mechanisms might be overwhelmed by the deluge of unreliable evidence. She might seek to regulate her resulting beliefs via the kinds of acceptance mechanisms discussed here, for the clearly epistemic purpose of retaining overall better beliefs. This is a topic that merits further exploration elsewhere.



### 5.3. How Far Can Acceptance Get Us?

Say that you buy my suggestion that acceptance has a significant role to play in the ethics of belief: that we might be able to make progress by recasting some of our doxastic moral duties as duties of acceptance. One might still wonder about an agent who (we stipulate) fully succeeds in accepting, but retains the underlying a belief—might she still be doing something wrong, or somehow failing morally? After all, in recent literature in the ethics of belief, some authors have insisted that their focus is on the *underlying belief itself*, independent of any downstream effects of that belief (e.g. see defenders of doxastic wrongdoing like Basu and Schroeder 2019; Basu 2018; 2019).

I'll say three brief things in response. The first is that my account reveals just how difficult it is to stipulate that an agent entirely succeeds in suppressing her belief. In many real-world cases, there will likely be moments in which the underlying belief state breaks through the agent's suppression efforts—and thus, someone who wants to resist the idea that mere beliefs can (be) morally wrong could locate their complaints about the agent in the places where the undesired belief still influences her cognition and action (though as I suggested above, I think that given the difficulty of the regulatory processes discussed here, we might want to think carefully about how we should think about the evaluative landscape of such instances). This also accords with the fact that the agent herself would likely recognize these breakthroughs as failures, and might feel guilty about them. Thus, we might care about the underlying belief not *independently* of its downstream effects, but precisely *because* those downstream effects—especially when understood to include downstream mental effects—are so ubiquitous and hard to entirely prevent.

Second, we must keep in mind that the cases in question are those in which, by stipulation, the evidential situation is morally or practically undesirable—in which they evidence points to a belief state that is in some sense unwanted by the agent. If, in such a case, an agent nonetheless commits herself to accepting some alternative, and puts in the cognitive work to regulate her cognition and

action in response to the moral or practical reasons she takes to be motivating—why should we think that the underlying belief state, which is primarily reflective of her evidential situation, is the compelling target of moral evaluation, compared to her state of acceptance which is much more richly expressive of her values and commitments?<sup>42</sup> It seems plausible that the latter is far more tied to her overall agency—and thus assessable in moral terms—than the former.<sup>43</sup> Nonetheless, I am ultimately happy to remain conciliatory. For all I’ve said here, it remains an open question—with room for reasonable disagreement—whether the belief state *itself*, independent of any downstream effects on cognition and action (and independent of any upstream failures of irresponsible reasoning or evidence gathering) can (be) morally wrong. My goal is not here to directly argue against the possibility of Basu- and-Schroeder-style doxastic wrongdoing: one could think such wrongdoing happens, and *also* buy the account of acceptance I’ve presented here and think it has a role to play in the ethics of belief.

## 6. Concluding Thoughts

We began this paper by observing that there a recurring tension in puzzles in the ethics and pragmatics of belief: on the one hand, there are many cases in which we appear to have compelling practical or moral reason to believe, but on the other hand, it is classically held that we cannot choose to believe on the basis of such non-evidential reasons. Here, I have I have proposed that perhaps many of these apparently moral/pragmatic doxastic duties can be reframed as duties of *acceptance* rather than belief. The present discussion has had two primary goals. First, I have sought to develop an empirically plausible psychological profile for acceptance (the empirical project); taking the high-level account presented by authors like Bratman as a starting point, I have suggested that a more detailed

---

<sup>42</sup> Rapstine (2021) develops the idea of *epistemic agent regret*, building on Bernard Williams’s conception of agent regret in the moral sphere, I find the heart of Rapstine’s proposal compelling: the idea that we can hold a belief, take that belief to be evidentially justified, but nevertheless regret being a “vehicle” for that belief on moral grounds. Acceptance gives us a resource to do something about our beliefs in such cases, rather than merely resigning ourselves to this regret.

<sup>43</sup> I in fact think there may be less difference between what I call acceptance and what authors like Basu and Schroeder call belief than it initially appears. I suspect that a difficulty in some discussions of the ethics of belief is that people are sometimes trading on importantly different notions of belief, where some are more thick and commitment-like, and others are more thin and merely-evidence-responsive. Untangling this idea is something I am pursuing elsewhere.

## ACCEPTANCE AND THE ETHICS OF BELIEF

discussion of the cognitive mechanisms involved can give us new insight into important characteristics of acceptance that were not apparent with only the high-level account in hand. Second, I have suggested that acceptance thus understood is a promising candidate to help us make progress on various puzzles in the ethics of belief (the normative project). I do not take the present discussion of acceptance to have settled what role acceptance can play in the ethics of belief; I merely hope to have motivated it as an idea worthy of serious further exploration. Acceptance allows us to be genuinely responsive to practical and moral considerations, and to regulate our beliefs in accordance with these reasons—without forcing us to posit an implausibly voluntaristic account on which we can simply choose what to believe directly in response to these reasons. But acceptance also does not force us to require the doxastic gymnastics of trying to trick ourselves into holding a belief, nor does it force us to have the aim of infringing on our well-functioning belief-forming mechanisms in a way that is likely to be epistemically suspect.

I'll conclude by highlighting three key takeaways. First, I hope to have motivated the role of acceptance as an important component of our cognitive economies, and a helpful tool in epistemology, philosophy of mind, and the ethics of belief. Second, I hope this present work invites us to expand how we think about doxastic control and the scope of our doxastic agency. I have tried here to shed light on a way in which we can exert significant control over our doxastic states, through the lens of *regulating* of our underlying beliefs, rather than through the traditional lens of choosing what to believe. Finally, I hope to have shown by example that the topic of doxastic control, which has long been primarily the purview of armchair philosophy, can be significantly elucidated with help from the empirical psychological sciences. Ultimately, this discussion of acceptance should be understood not as an ending point, but rather as a starting point: inviting us to both consider what other doxastic capacities can be understood via empirical investigation, and also to consider what roles belief regulation capacities such as acceptance may play in our epistemological, practical, and moral lives.

## References

- Alston, William P. 1988. "The Deontological Conception of Epistemic Justification." *Philosophical Perspectives* 2: 257–99. <https://doi.org/10.2307/2214077>.
- Arpaly, Nomy, and Anna Brinkerhoff. 2018. "Why Epistemic Partiality Is Overrated." *Philosophical Topics* 46 (1): 37–51. <https://doi.org/10.5840/philtopics20184613>.
- Basu, Rima. 2018. "Can Beliefs Wrong?" *Philosophical Topics* 46 (1): 1–18.
- . 2019. "The Wrongs of Racist Beliefs." *Philosophical Studies* 176 (9): 2497–2515. <https://doi.org/10.1007/s11098-018-1137-0>.
- . 2022. "Belief." *The Philosopher* 110 (2): 7–10.
- Basu, Rima, and Mark Schroeder. 2019. "Doxastic Wronging." In *Pragmatic Encroachment in Epistemology*, edited by Brian Kim and Matthew McGrath, 181–205. Routledge.
- Begby, Endre. 2013. "The Epistemology of Prejudice." *Thought: A Journal of Philosophy* 2 (2): 90–99. <https://doi.org/10.1002/tht3.71>.
- . 2021. *Prejudice: A Study in Non-Ideal Epistemology*. Oxford University Press.
- Bendaña, Joseph, and Eric Mandelbaum. 2021. "The Fragmentation of Belief." In *The Fragmented Mind*, by Joseph Bendaña and Eric Mandelbaum, 78–107. Oxford University Press. <https://doi.org/10.1093/oso/9780198850670.003.0004>.
- Bolinger, Renée Jørgensen. 2020a. "Varieties of Moral Encroachment." *Philosophical Perspectives* 34 (1): 5–26. <https://doi.org/10.1111/phpe.12124>.
- . 2020b. "The Rational Impermissibility of Accepting (Some) Racial Generalizations." *Synthese* 197 (6): 2415–31. <https://doi.org/10.1007/s11229-018-1809-5>.
- Botvinick, Matthew M., Todd S. Braver, Deanna M. Barch, Cameron S. Carter, and Jonathan D. Cohen. 2001. "Conflict Monitoring and Cognitive Control." *Psychological Review* 108: 624–52. <https://doi.org/10.1037/0033-295X.108.3.624>.
- Brans, Karen, Peter Koval, Philippe Verduyn, Yan Lin Lim, and Peter Kuppens. 2013. "The Regulation of Negative and Positive Affect in Daily Life." *Emotion* 13 (5): 926–39. <https://doi.org/10.1037/a0032400>.
- Bratman, Michael E. 1992. "Practical Reasoning and Acceptance in a Context." *Mind* 101 (401): 1–15.
- Brinkerhoff, Anna. 2021. "Prejudiced Beliefs Based on the Evidence: Responding to a Challenge for Evidentialism." *Synthese* 199 (5): 14317–31. <https://doi.org/10.1007/s11229-021-03422-y>.
- Brosch, Tobias, Klaus Scherer, Didier Grandjean, and David Sander. 2013. "The Impact of Emotion on Perception, Attention, Memory, and Decision-Making." *Swiss Medical Weekly* 143 (1920): w13786–w13786. <https://doi.org/10.4414/smw.2013.13786>.
- Butler, Emily A., Boris Egloff, Frank H. Wilhelm, Nancy C. Smith, Elizabeth A. Erickson, and James J. Gross. 2003. "The Social Consequences of Expressive Suppression." *Emotion* 3 (1): 48–67. <https://doi.org/10.1037/1528-3542.3.1.48>.
- Campbell-Sills, Laura, David H. Barlow, Timothy A. Brown, and Stefan G. Hofmann. 2006. "Acceptability and Suppression of Negative Emotion in Anxiety and Mood Disorders." *Emotion* 6 (4): 587–95. <https://doi.org/10.1037/1528-3542.6.4.587>.
- Chrisman, Matthew. 2008. "Ought to Believe." *The Journal of Philosophy* 105 (7): 346–70. <https://doi.org/10.5840/jphil2008105736>.
- Cohen, Jonathan D. 2017. "Cognitive Control." In *The Wiley Handbook of Cognitive Control*, 1–28. John Wiley & Sons, Ltd. <https://doi.org/10.1002/9781118920497.ch1>.
- Cohen, L. Jonathan. 1989. "Belief and Acceptance." *Mind* 98 (391): 367–89.

- Cohen, Laurence Jonathan, and Jonathan L. Cohen. 1992. *An Essay on Belief and Acceptance*. Clarendon Press.
- Costa, Albert, Alice Foucart, Inbal Arnon, Melina Aparici, and Jose Apesteguia. 2014. “Piensa’ Twice: On the Foreign Language Effect in Decision Making.” *Cognition* 130 (2): 236–54. <https://doi.org/10.1016/j.cognition.2013.11.010>.
- Crawford, Lindsay. 2022. “Suspending Judgment Is Something You Do.” *Episteme* 19 (4): 561–77. <https://doi.org/10.1017/epi.2022.40>.
- Devine, Patricia G. 1989. “Stereotypes and Prejudice: Their Automatic and Controlled Components.” *Journal of Personality and Social Psychology* 56: 5–18. <https://doi.org/10.1037/0022-3514.56.1.5>.
- Dolcos, Florin, and Ekaterina Denkova. 2014. “Current Emotion Research in Cognitive Neuroscience: Linking Enhancing and Impairing Effects of Emotion on Cognition.” *Emotion Review* 6 (4): 362–75. <https://doi.org/10.1177/1754073914536449>.
- Domínguez-Borràs, Judith, and Patrik Vuilleumier. 2013. “Affective Biases in Attention and Perception.” In *The Cambridge Handbook of Human Affective Neuroscience*, 331–56. New York, NY, US: Cambridge University Press. <https://doi.org/10.1017/CBO9780511843716.018>.
- Egan, Andy. 2008. “Seeing and Believing: Perception, Belief Formation and the Divided Mind.” *Philosophical Studies* 140 (1): 47–63. <https://doi.org/10.1007/s11098-008-9225-1>.
- Elga, Adam, and Agustín Rayo. 2021. “Fragmentation and Information Access.” In *The Fragmented Mind*, by Adam Elga and Agustín Rayo, 37–53. Oxford University Press. <https://doi.org/10.1093/oso/9780198850670.003.0002>.
- Engel, Pascal. 1998. “Believing, Holding True, and Accepting.” *Philosophical Explorations* 1 (2): 140–51. <https://doi.org/10.1080/10001998058538695>.
- Enoch, David, and Levi Spectre. n.d. “There Is No Such Thing as Doxastic Wrongdoing.” *Philosophical Perspectives*. Accessed December 28, 2022. <https://philarchive.org/rec/ENOTIN>.
- Evans, Jonathan St. B. T. 2008. “Dual-Processing Accounts of Reasoning, Judgment, and Social Cognition.” *Annual Review of Psychology* 59: 255–78.
- . 2019. “Reflections on Reflection: The Nature and Function of Type 2 Processes in Dual-Process Theories of Reasoning.” *Thinking & Reasoning* 25 (4): 383–415. <https://doi.org/10.1080/13546783.2019.1623071>.
- Evans, Jonathan St. B. T., and Keith E. Stanovich. 2013. “Dual-Process Theories of Higher Cognition: Advancing the Debate.” *Perspectives on Psychological Science* 8 (3): 223–41. <https://doi.org/10.1177/1745691612460685>.
- Fleisher, Will. 2018. “Rational Endorsement.” *Philosophical Studies* 175 (10): 2649–75. <https://doi.org/10.1007/s11098-017-0976-4>.
- Flowerree, A. K. 2017. “Agency of Belief and Intention.” *Synthese* 194 (8): 2763–84. <https://doi.org/10.1007/s11229-016-1138-5>.
- Fraassen, Bas C. Van. 1985. *Images of Science: Essays on Realism and Empiricism*. University of Chicago Press.
- Franchow, Emilie I., and Yana Suchy. 2015. “Naturally-Occurring Expressive Suppression in Daily Life Depletes Executive Functioning.” *Emotion* 15 (1): 78–89. <https://doi.org/10.1037/emo0000013>.
- Franchow, Emilie I., and Yana Suchy. 2017. “Expressive Suppression Depletes Executive Functioning in Older Adulthood.” *Journal of the International Neuropsychological Society* 23 (4): 341–51. <https://doi.org/10.1017/S1355617717000054>.
- Frankish, Keith. 2007a. “Deciding to Believe Again.” *Mind* 116 (463): 523–48. <https://doi.org/10.1093/mind/fzm523>.

- . 2007b. *Mind and Supermind*. Cambridge University Press.
- Friedman, Jane. 2013. “Suspended Judgment.” *Philosophical Studies* 162 (2): 165–81. <https://doi.org/10.1007/s11098-011-9753-y>.
- . 2017. “Why Suspend Judging?” *Nous* 51 (2): 302–26. <https://doi.org/10.1111/nous.12137>.
- Frijda, Nico H. 1987. “Emotion, Cognitive Structure, and Action Tendency.” *Cognition and Emotion* 1 (2): 115–43. <https://doi.org/10.1080/02699938708408043>.
- Frijda, Nico H., Peter Kuipers, and Elisabeth ter Schure. 1989. “Relations among Emotion, Appraisal, and Emotional Action Readiness.” *Journal of Personality and Social Psychology* 57: 212–28. <https://doi.org/10.1037/0022-3514.57.2.212>.
- Giuliani, Nicole R., Emily M. Drabant, Roshni Bhatnagar, and James J. Gross. 2011. “Emotion Regulation and Brain Plasticity: Expressive Suppression Use Predicts Anterior Insula Volume.” *NeuroImage* 58 (1): 10–15. <https://doi.org/10.1016/j.neuroimage.2011.06.028>.
- Greene, Joshua D. 2007. “Why Are VMPFC Patients More Utilitarian? A Dual-Process Theory of Moral Judgment Explains.” *Trends in Cognitive Sciences* 11 (8): 322–23. <https://doi.org/10.1016/j.tics.2007.06.004>.
- . 2015. “Beyond Point-and-Shoot Morality: Why Cognitive (Neuro)Science Matters for Ethics.” *Law & Ethics of Human Rights* 9 (2): 141–72. <https://doi.org/10.1515/lehr-2015-0011>.
- Greene, Joshua D., Fiery A. Cushman, Lisa E. Stewart, Kelly Lowenberg, Leigh E. Nystrom, and Jonathan D. Cohen. 2009. “Pushing Moral Buttons: The Interaction between Personal Force and Intention in Moral Judgment.” *Cognition* 111 (3): 364–71. <https://doi.org/10.1016/j.cognition.2009.02.001>.
- Greene, Joshua D., R. Brian Sommerville, Leigh E. Nystrom, John M. Darley, and Jonathan D. Cohen. 2001. “An fMRI Investigation of Emotional Engagement in Moral Judgment.” *Science* 293 (5537): 2105–8. <https://doi.org/10.1126/science.1062872>.
- Gross, James J. 1998a. “Antecedent- and Response-Focused Emotion Regulation: Divergent Consequences for Experience, Expression, and Physiology.” *Journal of Personality and Social Psychology* 74 (1): 224–37. <https://doi.org/10.1037/0022-3514.74.1.224>.
- . 1998b. “The Emerging Field of Emotion Regulation: An Integrative Review.” *Review of General Psychology* 2 (3): 271–99. <https://doi.org/10.1037/1089-2680.2.3.271>.
- Gross, James J., and Oliver P. John. 2003. “Individual Differences in Two Emotion Regulation Processes: Implications for Affect, Relationships, and Well-Being.” *Journal of Personality and Social Psychology* 85 (2): 348–62. <https://doi.org/10.1037/0022-3514.85.2.348>.
- Gross, James J., and Robert W. Levenson. 1993. “Emotional Suppression: Physiology, Self-Report, and Expressive Behavior.” *Journal of Personality and Social Psychology* 64 (6): 970–86. <https://doi.org/10.1037/0022-3514.64.6.970>.
- Gyurak, Anett, Madeleine S. Goodkind, Joel H. Kramer, Bruce L. Miller, and Robert W. Levenson. 2012. “Executive Functions and the Down-Regulation and up-Regulation of Emotion.” *Cognition and Emotion* 26 (1): 103–18. <https://doi.org/10.1080/02699931.2011.557291>.
- Haidt, Jonathan. 2001. “The Emotional Dog and Its Rational Tail: A Social Intuitionist Approach to Moral Judgment.” *Psychological Review* 108 (4): 814–34.
- Hieronymi, Pamela. 2006. “Controlling Attitudes.” *Pacific Philosophical Quarterly* 87 (1): 45–74. <https://doi.org/10.1111/j.1468-0114.2006.00247.x>.
- . 2008. “Responsibility for Believing.” *Synthese* 161 (3): 357–73. <https://doi.org/10.1007/s11229-006-9089-x>.
- Jackson, Elizabeth Grace. 2021. “A Permissivist Defense of Pascal’s Wager.” *Erkenntnis*, September. <https://doi.org/10.1007/s10670-021-00454-1>.

- Kahneman, Daniel, and Amos Tversky. 1979. "Prospect Theory: An Analysis of Decision Under Risk." *Econometrica (Pre-1986)* 47 (2): 263–91.
- Kashdan, Todd B., and Michael F. Steger. 2006. "Expanding the Topography of Social Anxiety: An Experience-Sampling Assessment of Positive Emotions, Positive Events, and Emotion Suppression." *Psychological Science* 17 (2): 120–28. <https://doi.org/10.1111/j.1467-9280.2006.01674.x>.
- Kawall, Jason. 2013. "Friendship and Epistemic Norms." *Philosophical Studies* 165 (2): 349–70. <https://doi.org/10.1007/s11098-012-9953-0>.
- Keller, Simon. 2004. "Friendship and Belief." *Philosophical Papers* 33 (3): 329–51. <https://doi.org/10.1080/05568640409485146>.
- . 2018. "Belief for Someone Else's Sake." *Philosophical Topics* 46 (1): 19–36.
- Keysar, Boaz, Sayuri L. Hayakawa, and Sun Gyu An. 2012. "The Foreign-Language Effect: Thinking in a Foreign Tongue Reduces Decision Biases." *Psychological Science* 23 (6): 661–68. <https://doi.org/10.1177/0956797611432178>.
- Koole, Sander L. 2009. "The Psychology of Emotion Regulation: An Integrative Review." *Cognition and Emotion* 23 (1): 4–41. <https://doi.org/10.1080/02699930802619031>.
- Lynch, Thomas R., Clive J. Robins, Jennifer Q. Morse, and Elizabeth D. Krause. 2001. "A Mediation Model Relating Affect Intensity, Emotion Inhibition, and Psychological Distress." *Behavior Therapy* 32 (3): 519–36. [https://doi.org/10.1016/S0005-7894\(01\)80034-4](https://doi.org/10.1016/S0005-7894(01)80034-4).
- Maguire, Barry, and Jack Woods. 2020. "The Game of Belief." *The Philosophical Review* 129 (2): 211–49. <https://doi.org/10.1215/00318108-8012843>.
- Masny, Michal. 2020. "Friedman on Suspended Judgment." *Synthese* 197 (11): 5009–26. <https://doi.org/10.1007/s11229-018-01957-1>.
- Matos, Marcela, José Pinto-Gouveia, and Vânia Costa. 2013. "Understanding the Importance of Attachment in Shame Traumatic Memory Relation to Depression: The Impact of Emotion Regulation Processes." *Clinical Psychology & Psychotherapy* 20 (2): 149–65. <https://doi.org/10.1002/cpp.786>.
- McCormick, Miriam Schleifer. 2022. "Belief as Emotion." *Philosophical Issues* 32 (1): 104–19. <https://doi.org/10.1111/phis.12232>.
- McGrath, Matthew. 2021. "Being Neutral: Agnosticism, Inquiry and the Suspension of Judgment." *Nous* 55 (2): 463–84. <https://doi.org/10.1111/nous.12323>.
- McKaughan, Daniel J. 2007. "Toward a Richer Vocabulary for Epistemic Attitudes: Mapping the Cognitive Landscape." Ph.D., United States -- Indiana: University of Notre Dame. <http://search.proquest.com/docview/304819154/abstract/D48CC791876E46A5PQ/1>.
- McRae, Kateri. 2016. "Cognitive Emotion Regulation: A Review of Theory and Scientific Findings." *Current Opinion in Behavioral Sciences, Neuroscience of education*, 10 (August): 119–24. <https://doi.org/10.1016/j.cobeha.2016.06.004>.
- Moore, Sally A., Lori A. Zoellner, and Niklas Mollenholt. 2008. "Are Expressive Suppression and Cognitive Reappraisal Associated with Stress-Related Symptoms?" *Behaviour Research and Therapy* 46 (9): 993–1000. <https://doi.org/10.1016/j.brat.2008.05.001>.
- Morton, Jennifer M., and Sarah K. Paul. 2019. "Grit." *Ethics* 129 (2): 175–203. <https://doi.org/10.1086/700029>.
- Muris, Peter, Harald Merckelbach, Marcel van den Hout, and Peter de Jong. 1992. "Suppression of Emotional and Neutral Material." *Behaviour Research and Therapy* 30 (6): 639–42. [https://doi.org/10.1016/0005-7967\(92\)90009-6](https://doi.org/10.1016/0005-7967(92)90009-6).
- Niermeyer, Madison A., Emilie I. Franchow, and Yana Suchy. 2016. "Reported Expressive Suppression in Daily Life Is Associated with Slower Action Planning." *Journal of the*

- International Neuropsychological Society* 22 (6): 671–81.  
<https://doi.org/10.1017/S1355617716000473>.
- Paul, Sarah K., and Jennifer M. Morton. 2018. “Believing in Others.” *Philosophical Topics*. August 1, 2018. <https://doi.org/10.5840/philtopics20184615>.
- Payne, B. Keith. 2005. “Conceptualizing Control in Social Cognition: How Executive Functioning Modulates the Expression of Automatic Stereotyping.” *Journal of Personality and Social Psychology* 89 (4): 488–503. <https://doi.org/10.1037/0022-3514.89.4.488>.
- Phan, K. Luan, and Chandra Sekhar Sripada. 2013. “Emotion Regulation.” In *The Cambridge Handbook of Human Affective Neuroscience*, edited by Jorge Armony and Patrik Vuilleumier, 375–400.
- Rapstine, Mica. 2021. “Regrettable Beliefs.” *Philosophical Studies* 178 (7): 2169–90.  
<https://doi.org/10.1007/s11098-020-01535-7>.
- Reisner, Andrew. 2008. “Weighing Pragmatic and Evidential Reasons for Belief.” *Philosophical Studies* 138 (1): 17–27. <https://doi.org/10.1007/s11098-006-0007-3>.
- Richards, Jane M. 2004. “The Cognitive Consequences of Concealing Feelings.” *Current Directions in Psychological Science* 13 (4): 131–34. <https://doi.org/10.1111/j.0963-7214.2004.00291.x>.
- Richards, Jane M., and James J. Gross. 1999. “Composure at Any Cost? The Cognitive Consequences of Emotion Suppression.” *Personality and Social Psychology Bulletin* 25 (8): 1033–44. <https://doi.org/10.1177/01461672992511010>.
- . 2000. “Emotion Regulation and Memory: The Cognitive Costs of Keeping One’s Cool.” *Journal of Personality and Social Psychology* 79 (3): 410–24. <https://doi.org/10.1037/0022-3514.79.3.410>.
- Rinard, Susanna. 2015. “Against the New Evidentialists.” *Philosophical Issues* 25 (1): 208–23.  
<https://doi.org/10.1111/phis.12061>.
- Roeber, Blake. 2019. “Evidence, Judgment, and Belief at Will.” *Mind* 128 (511): 837–59.  
<https://doi.org/10.1093/mind/fzy065>.
- . 2020. “Permissive Situations and Direct Doxastic Control.” *Philosophy and Phenomenological Research* 101 (2): 415–31. <https://doi.org/10.1111/phpr.12594>.
- Roemer, Lizabeth, and Thomas D. Borkovec. 1994. “Effects of Suppressing Thoughts about Emotional Material.” *Journal of Abnormal Psychology* 103: 467–74.  
<https://doi.org/10.1037/0021-843X.103.3.467>.
- Ross, Lewis. 2022. “Profiling, Neutrality, and Social Equality.” *Australasian Journal of Philosophy* 100 (4): 808–24. <https://doi.org/10.1080/00048402.2021.1926522>.
- Scarantino, Andrea. 2014. “The Motivational Theory of Emotions.” In *Moral Psychology and Human Agency: Philosophical Essays on the Science of Ethics*, 156–85. New York, NY, US: Oxford University Press. <https://doi.org/10.1093/acprof:oso/9780198717812.003.0008>.
- Shah, Nishi. 2002. “Clearing Space For Doxastic Voluntarism.” *The Monist*. August 1, 2002.  
<https://doi.org/10.5840/monist200285326>.
- . 2003. “How Truth Governs Belief.” *The Philosophical Review* 112 (4): 447–82.
- Shah, Nishi, and J. David Velleman. 2005. “Doxastic Deliberation.” *The Philosophical Review* 114 (4): 497–534.
- Shenhav, Amitai, Sebastian Musslick, Falk Lieder, Wouter Kool, Thomas L. Griffiths, Jonathan D. Cohen, and Matthew M. Botvinick. 2017. “Toward a Rational and Mechanistic Account of Mental Effort.” *Annual Review of Neuroscience* 40: 99–124.
- Sher, George. 2019. “A Wild West of the Mind.” *Australasian Journal of Philosophy* 97 (3): 483–96.  
<https://doi.org/10.1080/00048402.2018.1490326>.



- Sperberg, Elizabeth D., and Sally D. Stabb. 1998. "Depression in Women as Related to Anger and Mutuality in Relationships." *Psychology of Women Quarterly* 22 (2): 223–38.  
<https://doi.org/10.1111/j.1471-6402.1998.tb00152.x>.
- Sripada, Chandra. 2018. "Addiction and Fallibility." *The Journal of Philosophy* 115 (11): 569–87.  
<https://doi.org/10.5840/jphil20181151133>.
- . 2021. "The Atoms of Self-Control." *Nous* 55 (4): 800–824.  
<https://doi.org/10.1111/nous.12332>.
- Staffel, Julia. 2019. "Credences and Suspended Judgments as Transitional Attitudes." *Philosophical Issues* 29 (1): 281–94. <https://doi.org/10.1111/phis.12154>.
- Stalnaker, Robert, and Robert Stalnaker. 1984. *Inquiry*. Cambridge, Mass.: MIT Press.
- Steup, Matthias. 2012. "Belief Control and Intentionality." *Synthese* 188 (2): 145–63.  
<https://doi.org/10.1007/s11229-011-9919-3>.
- Stroud, Sarah. 2006. "Epistemic Partiality in Friendship." *Ethics* 116 (3): 498–524.  
<https://doi.org/10.1086/500337>.
- Traldi, Oliver. 2022. "Uncoordinated Norms of Belief." *Australasian Journal of Philosophy* 0 (0): 1–13.  
<https://doi.org/10.1080/00048402.2022.2030378>.
- Weatherson, Brian. 2008. "Deontology and Descartes's Demon." *The Journal of Philosophy* 105 (9): 540–69.
- Williams, Bernard. 1973. "Deciding to Believe." In *Problems of the Self*, edited by Bernard Williams, 136–51. Cambridge University Press.