

Copyright
by
Amelia Kahn
2021

**The Dissertation Committee for Amelia Kahn Certifies that this is the approved
version of the following Dissertation:**

Essays on Ideal Norms and Non-Ideal Agents

Committee:

Sinan Dogramaci, Supervisor

Miriam Schoenfield

Daniel Drucker

David Christensen

Essays on Ideal Norms and Non-Ideal Agents

by

Amelia Kahn

Dissertation

Presented to the Faculty of the Graduate School of

The University of Texas at Austin

in Partial Fulfillment

of the Requirements

for the Degree of

Doctor of Philosophy

The University of Texas at Austin

August 2021

Acknowledgements

I'd like to express my sincerest thanks to my advisor, Sinan Dogramaci for his many years of excellent (if not often heeded) advice, encouragement, and patience. I would also like to extend my thanks the other members of my dissertation committee – David Christensen, Daniel Drucker, and Miriam Schoenfield, each one of whom has been incredibly generous with their time, kind, and thoughtful in their engagement with my ideas. I am also deeply grateful to Matty Silverstein, who introduced me to philosophy and without whose encouragement this dissertation would not have been started, let alone finished.

I would like to further thank my family and friends whose support has been invaluable at every stage and indispensable in the last months and years: my mother, Kathryn Kahn, Abby Attia, Brian Pollex, Briana Toole, Florencia Schlamp, Sachith Joseph Cheruvatur, Bob Porter, Mary Scott, Nicole Lee, Henry Schiller, Anthony Spalvieri-Kruse, Nishant Mohanchandra, Ben Leb, Connie Trang Trinh, Hannah Trees, Kieran Lerch, Erin Bannen, Adam Dolan, Cassie Flores, C. Warren Dale, Stephen Underwood, Alyazia Al Shaibani, Alex Grossman, Gayle Mangan, Robin Fisher, Jerry and Margie Stein, Ornella Halawi, and the folks at the Crown & Anchor Pub.

Abstract

Essays on Ideal Norms and Non-Ideal Agents

Amelia Kahn, Ph.D.

The University of Texas at Austin, 2021

Supervisor: Sinan Dogramaci

My dissertation focuses on idealized epistemology and challenges that arise when we apply it to actual human agents. With the rise of formal epistemology, worries about the role that idealization plays in normative epistemology and the relevance of highly idealized norms to us non-ideal agents have become more pressing. I focus on two types of challenges to idealized norms. The first, which is the subject of the first two chapters, concerns constraints that our limited cognitive resources and time place on our epistemic lives. The second is a metaphysical challenge to theorizing about credences, on the basis that they poorly represent our actual graded belief states.

The first chapter focuses on clutter avoidance worries. Gilbert Harman argues that a strong logical closure norm which would require us to believe all the consequences of our beliefs must be rejected because following such a norm would “clutter up our minds with trivialities.” Call this the Clutter Argument. The Clutter Argument has attracted broad sympathy in the literature, yet there has been surprisingly little attention paid to how the Clutter Argument is meant to work. In this paper I argue that the Clutter Argument fails. First, I argue for a novel account of clutter, the waste view, which identifies clutter with beliefs that waste cognitive resources. I motivate the acceptance of this new account in part by presenting several counterexamples to Jane Friedman’s

(2018) account, which identifies clutter with beliefs about subject matters in which the agent has no interest. I argue that the only plausible clutter principle is one that rules out dynamic norms that would require we believe too much clutter. Finally, I show that even with such a clutter principle, the Clutter Argument fails because actual human agents typically believe consequences of their beliefs implicitly which does require them to clutter their minds.

My second chapter takes as a starting point the recent claim from Jane Friedman that the norms of inquiry conflict with many norms of contemporary epistemology. In response to this conflict, we face a number of options. We can resign ourselves to the incoherence in the epistemic domain. We can radically revise the epistemic to eliminate the conflict. Finally, we can cleave the norms of inquiry from the epistemic, most naturally by insisting that they are, at bottom, practical. Call this final approach the separationist view. I offer a puzzle for the separationist gambit and canvas the possible responses to it. I argue that the best response is to reject a key premise: the claim that when you ought to come to believe P , and you cannot come to believe P and Q at the same time, then you are not permitted to come to believe that Q . However, I show that any satisfactory response the puzzle dissolves the original tension between the epistemic and the norms of inquiry, leaving us with no need to pursue the separationist strategy.

The final chapter tackles the issue of whether comparativism about credences squares with probabilism. I offer a dilemma for comparativism about credences, the view that comparative beliefs, not real-valued credences, are psychologically real and primitive. If we require comparative beliefs to be probabilistically representable in a very strict way, then we should think that many rational agents' comparative beliefs are not representable by a probability function. On the other hand, if we require only a weaker form of probabilistic representability, then credences will fail to represent agents sufficiently well to epistemically evaluate agents.

Table of Contents

Introduction	1
Chapter 1: Clutter, Junk, and Waste	4
1: <i>What is Clutter?</i>	5
2: <i>The Normative Significance of Clutter</i>	14
3: <i>The Clutter Argument</i>	20
4: <i>Conclusion</i>	32
Chapter 2: Dissolving the Tension between the Epistemic and The Zetetic	34
1: <i>Tension Between the Epistemic and the Zetetic</i> 35 2: <i>A Puzzle for the Separationist</i>	40
3: <i>Responses to the Puzzle</i>	42
4: <i>Dissolving the Tension</i>	51
5: <i>Conclusion</i>	53
Chapter 3: Can comparativism about credence deliver probabilism?	55
1: <i>Comparativism</i>	56
2: <i>The Comparativist Argument for Probabilism</i>	64
3: <i>The Totality Condition</i>	70
4: <i>Representability without Totality</i>	77
5: <i>Conclusion</i>	81
Works Cited	85

INTRODUCTION

Epistemology issues demands beyond what we limited human agents can hope to accomplish. Formal epistemology, which models rational belief and knowledge with logic and probability, is perhaps most associated with norms that idealize away from human agents. However, norms that only idealized agents can meet are found in every corner of epistemology. Traditional epistemology gives us evidential norms which normal agents will inevitably fall short of meeting. Recent trends in epistemology offer us still more ideal norms. Knowledge-first epistemology commands us to believe P only if we know it. Given our imperfect access to what we know, we're bound to miss the mark. Even attempts to narrow the focus to humans and our projects have yielded norms beyond human reach. For instance, Jane Friedman's clutter avoidance norm requires that we avoid believing propositions in which we have no interest. This, of course, is wholly impossible for normal agents. Ideal norms abound in epistemology. So too do concerns about what ideal norms mean for human agents.

The central theme of these essays is whether, and how, ideal norms can apply to non-ideal agents. I do not aspire to give a unified approach idealization in epistemology. Rather, my aim is to think through a number of challenges to particular idealized norms—logical closure norms, evidential and knowledge norms, and the norms of probabilism— and address each on its merits. Answering each of these challenges, it turns out, leads to broader questions about how we delineate the epistemic domain, how norms on forming beliefs relate to norms about the beliefs we should have at a given time, among others.

In the first essay, I examine Gilbert Harman's Clutter Argument against a logical closure norm that requires that we believe all of the logical consequences of what we believe. Such a norm, he says, could not apply to normal human agents, because if we were to do as it bids (or try to, for that matter), we would clutter our minds with useless, boring, trivialities. While this looks, at first, like a familiar ought-implies-can concern, there is another, more compelling way to understand clutter avoidance. The idea, roughly, is that the mandate to avoid cluttering our minds falls out of the type of agents we are: ones with limited cognitive resources. Logical closure, on the other hand, is a norm only highly idealized agent can meet. Surely, the thought goes, if one of these has to go, it's the extremely demanding ideal norm.

I attempt to give an account of what this cognitive clutter *is* and why it is worth avoiding that accords with our intuitive judgments and with the relevant psychological findings. Using this account, which I call the *waste view*, I reconstruct what I think is the best version of the Clutter Argument. In the process, some interesting features of the argument come into view. For one, it relies on a surprisingly strong principle about how norms on belief revision relate to norms on belief at a time. What's more, the argument's success depends on our answer to an empirical question: what (typically) is the actual cost in terms of cognitive resources when normal human agents form boring, useless, clutter-y beliefs?

In the second essay, I look into a tension that Jane Friedman points to between the norms that govern inquiry and more traditional epistemic norms. In some cases, she argues, we cannot follow these familiar epistemic norms and inquire as we ought to. Inquiry occupies a somewhat odd place within epistemology. On the one hand, it seems clear that inquiring is central to our epistemic lives. On the other, carrying out inquiry happens in the world and therefore is subject to practical constraints that are typically ignored when we ask how agents ought to believe. While Friedman argues that this motivates a radical revision of our familiar epistemic norms, many others see the tension as reason to relegate norms of inquiry to the practical domain.

However, I argue that each of the epistemic norms that produces tension with the norms of inquiry produces the very same sort of tension with itself. To put it glibly, the tension was coming from *inside* the epistemic all along. I offer a way to dissolve this tension that also dissolves the initial tension between the epistemic and the norms of inquiry. The upshot is that we do not need to radically revise the epistemic to accommodate inquiry, nor do we need to banish inquiry to the practical domain. Undoubtedly, many theorists still will. But if they do, they do so based on their conception of what the epistemic is like, not because norms of inquiry and other epistemic norms put agents in a bind.

The final essay deals with a different sort of concern about idealized norms: the *attitudes* that they govern. Probabilism is a widely discussed and hugely productive view about rational credence. Yet, as many a hushed conversation in the hallway of philosophy departments would attest, some philosophers are somewhat resistant to credence talk. One way that you might respond to worries that precise real-valued credences are psychologically unrealistic is to underwrite them with something less fine-grained, more easily introspectable: comparative belief. This is, in broad outline, comparativism: the view that comparative belief, not credence, is psychologically real and primitive. Now, the credence skeptic is expressing doubt about whether

she, a human agent, has credences. Neither I nor the standard-bearers of comparativism directly take up the question of whether irrational comparative beliefs could underwrite (perhaps improbableistic) credences. However, I ask a related question in the third chapter. I argue that comparativism cannot, in the end, vindicate probabilism. The reason is that probabilities are not a particularly good representation of even a rational agent's comparative beliefs. If a rational agent's comparative beliefs are not well represented by probability functions, then it seems comparativism has little to offer the credence skeptic.

It is probably clear by now that these essays have little in common as far as their content goes. The thread they share is the abiding interest that led me to each of them: How can ideal norms apply to regular agents like us?

CHAPTER 1: CLUTTER, JUNK, AND WASTE

A common but controversial proposal for how logic might constrain rational belief is that our beliefs ought to be closed under logical consequence. That is, we ought to believe all of the logical consequences of our beliefs. Gilbert Harman argues, however, that such a principle would require us to clutter our minds with all sorts of trivial consequences of our beliefs. We are agents with limited time, attention, and memory, and using these limited cognitive resources to believe trivialities would be, as Harman (1986) aptly puts it, “worse than pointless”. I’ll refer to this argument as the Clutter Argument against logical closure norms.

The Clutter Argument has attracted a great deal of support. Wallace (2001), Sainsbury (2002), DePaul (2004), MacFarlane (2004), White (2005), Douven (2010), Wedgewood (2012), Friedman (2018) and Dogramaci (2018) all express some support for the thought that clutter considerations constrain what we should believe. Despite this widespread sympathy, there has been relatively little discussion outside of Friedman (2018) about what clutter is and exactly how the Clutter Argument works.

I argue here that it’s a mistake to reject logical closure norms on the basis of the Clutter Argument (though they might certainly be rejected on other grounds). In Section 1, I unpack the idea of clutter, argue that Jane Friedman’s (2018) account of clutter misses the mark, and offer my own account. In Section 2, I discuss the normative significance of clutter and articulate what I take to be the most plausible clutter avoidance principle. In Section 3, I formulate the Clutter Argument and argue that it fails. Finally, I discuss the relevance of the Clutter Argument to weaker, more plausible logical closure

principles and sketch an alternative picture of an epistemic landscape where clutter avoidance and strong logical closure principles coexist.

1: What is Clutter?

In this section, I aim to clarify the notion of clutter. In 1.1, I introduce Harman's notion of "cluttering the mind" and discuss what it means for a mind to be cluttered. In 1.2, I discuss Friedman's view, which identifies clutter with "junk beliefs," and present several counterexamples to her view. Finally, in 1.3, I develop my own account of clutter, the waste view, which avoids the problems that face Friedman's view.

1.1 HOW CAN MINDS BE CLUTTERED?

Harman suggests that certain trivial beliefs can "clutter" our minds:

Many trivial things are implied by one's view which it would be worse than pointless to add to what one believes. For example, if one believes P, one's view trivially implies "either P or Q," "either P or P," "P and either P or R," and so on. There is no point in cluttering one's mind with all these propositions. (1986, p.12)

To dramatize Harman's point, we can imagine a mirthless logician who attempts to do as logical closure norms bid him, spending hours upon hours believing ever more trivial consequences of his beliefs. This activity is pointless at best and outright irrational at worst. But what does it mean to say that the mirthless logician's mind will be "cluttered" by all of these beliefs? Harman has this to say:

To suppose one's mind could become cluttered with beliefs is to suppose such things as (1) that it takes time to add to one's beliefs further propositions that are trivially implied by them, time that might be better spent on other things, and/or (2) that one has "limited storage capacity" for beliefs, so there is a limit on the

number of things one can believe, and/or (3) that there are limits on "information retrieval," so the more one believes the more difficult it is to recall relevant beliefs when one needs them.(1986, p.12)

The general picture is that clutter uses up limited cognitive resources which might have been used on more interesting or useful matters. Harman points to three such cognitive resources: time, storage capacity, and retrieval from memory.

It's easy to see that deliberating about and forming beliefs often requires some of our limited time.¹ The case of storage capacity, however, is not so straightforward. Attempts to estimate of the storage capacity of the human brain are on shaky methodological ground and none give us reason to suppose that any actual humans run a risk of running out of space (Dudai 1997) As Michaelian (2011) notes, the capacity of human memory is "unlimited in practical terms," whatever theoretical limits it may have. Limitations on storage capacity, then, do not constrain what we should believe.

The ability to retrieve information from memory looks more promising. "Retrieval-induced forgetting," in which recalling certain items from a list causes impaired recall of other items from the list, is well-supported in the psychology literature.² Most research on retrieval-induced forgetting uses category-exemplar pairs (i.e., *fruit-lemon*, *animal-dog*) as cues but the phenomenon appears to generalize to other sorts of cues, including arithmetic facts (Campbell and Thompson 2012; Campbell, Dufour, and Chen 2015), autobiographical memory (Barnier, Hung, and Conway 2004), and location information (Gómez-Ariza, Fernandez, and Bajo 2012).³ However there are important caveats to this general result. First, there's research that suggests that recalling some items from a list sometimes *improves* recall of the non-practiced items in the long term, a phenomenon called retrieval-induced facilitation (Chan, McDermott, and

¹ Some may worry that this applies only *explicit and deliberative* belief formation rather belief formation in general. I'll discuss the cognitive burden of implicit beliefs in Section 3.4.

² Murayama, et al. (2014) provides an overview and meta-analysis of research on retrieval-induced forgetting.

³ These are particularly relevant because, unlike category-exemplar pairs, they are propositions and therefore could be the content of a belief.

Roediger 2006; Chan 2009). Second, the research suggests that it's *retrieving* beliefs already stored in memory, rather than forming new beliefs that induces forgetting. It's not clear, then, that retrieval-induced forgetting gives us reason to avoid forming trivial beliefs in the first place. The research bears out the idea that there are limitations on information retrieval from memory, but fails to support Harman's suggestion that forming new beliefs taxes retrieval.

Research on directed attention fatigue suggests that attention is a resource that can be depleted, as well (Kaplan and Berman 2010). This, no doubt, is anecdotally obvious to those who have tried to cram for an exam or write careful philosophy for ten hours straight.

To sum up, the cognitive resources that clutter squanders are time and attention. While appeals to storage capacity might seem natural given the clutter metaphor, there's little reason to think actual human agents might run up against limits on storage capacity. Similarly, retrieval-induced forgetting does not constrain belief *formation*, since belief formation does not trigger forgetting.

We've seen *how* trivial beliefs can "clutter the mind" by using up limited cognitive resources. Now we can turn to the question of *which* beliefs clutter the mind. Harman focuses on trivial logical consequences because of their relevance to logical closure norms. However, Friedman rightly notes that trivial logical beliefs are not the only beliefs that squander our limited cognitive resources. Boring, useless beliefs about all sorts of topics are worth avoiding for precisely the same reasons as trivial logical consequences. An adequate account of clutter should follow Friedman in considering this wider class of "worse than useless" beliefs.

1.2 THE JUNK VIEW

Friedman defines “junk” as a subject matter in which an agent has no interest.⁴ To “have an interest” in this sense, is to either find a subject matter interesting, or to have “desires served by having beliefs (or knowing) about a subject matter” (2018). For instance, I have an interest in whether prickly pears are native to Texas because I find it interesting, but I have an interest in whether my rent is due on the 1st of each month because I’d like to avoid late fees, not because I find it particularly interesting. Friedman’s view, which I’ll call *the junk view*, takes junk beliefs to be the source of clutter.

Friedman’s “clutter as junk” approach is straightforward and appealing. To avoid clutter, avoid junk beliefs, since these are boring and useless beliefs that waste our cognitive resources. However, the junk view fails to capture our intuitions about which beliefs should be avoided on clutter grounds.⁵ Consider the following example:

Reluctant Belief: A window pops up on Paul’s computer screen. As Paul reads the text in the window, he sees that it presents an excellent argument for P from premises he believes. Paul refuses to give up his beliefs in the premises of the argument. Believing P would not cause Paul distress or distraction. He is conflicted about whether to believe that P on the basis of this argument that is excellent by his own lights because he has no interest whatsoever in whether P. Reluctantly, he comes to believe that P.⁶

P is junk, for Paul, because he has no interest in whether P. Nonetheless, intuitively Paul is at least permitted (if not required) to believe that P, after having

⁴ A subject matter, in Friedman’s terminology, is the answer to a “whether” question -the set of a proposition and its negation.

⁵ I’m supposing here that we have intuitions about which beliefs are clutter. The mirthless logician, inferring away, has beliefs that clearly tax his cognitive resources and gain him very little. Forming beliefs about how many grains of sand are in a handful picked up at random is similarly worse than useless. Our intuition is that these beliefs are to be avoided, because they squander our limited cognitive resources.

⁶ Thanks to Matt Vermaire, Derek Haderlie, and Ian Proops for helpful comments on reworking this example.

considered excellent argument for it from premises that he believes. Once he has already invested the time and attention necessary to entertain P, appreciate the argument, and conclude that he wouldn't give up his belief in any of the premises, believing that P would not further tax his cognitive resources. It may be that he should have avoided reading the pop-up at all, but this has little bearing on what he should do, having already read it. So, Paul's belief is junk, but not clutter.

This points to a general problem: the junk view can't distinguish junk beliefs that tax cognitive resources from those, like Paul's, that don't. This blind spot looks even more troubling when we consider cases where junk beliefs have a net positive effect on cognitive resources:

Study Break: Evelyn is studying for an exam. After hitting the books for several hours, she finds herself distracted and fatigued. To clear her head, she walks to her window and spends five minutes slowly, meditatively counting the people in the courtyard. She forms a series of beliefs of the form "there are at least X people in the courtyard". She is not interested in how many people are in the courtyard, nor does she have any desire served by having this information. When she resumes her studying, she has an easier time focusing and retains the material better than before her break.⁷

The beliefs that Evelyn forms about how many people are in the courtyard are junk, for her. Yet, while she did spend some cognitive resources to form these beliefs, she finishes her study break more able to attend to the material she's trying to learn and will likely learn it faster as a result. That is, she ends up better off in terms of cognitive resources than she was before her break. It seems, then, that Evelyn's beliefs are not clutter, so the junk view gets this case wrong.

⁷ Brian Pollex pointed out in conversation that as soon as Evelyn sees that there are a bunch of people in the courtyard, she already implicitly believes "there is at least one person in the courtyard," "there are at least two people in the courtyard." Her later beliefs of this sort, however, will be genuinely formed as she counts.

One might contend that Evelyn *does* have desires served by having beliefs about how many people are in the courtyard: the desire to clear her head or the desire to study more effectively. This response requires a slight tweak to Friedman's account of having an interest in a subject matter. An agent has an interest in a subject matter if the agent's desires are served by *forming*, not by merely having, a belief about that subject matter. This captures the thought that a belief can serve a subject's desires because of its content *or* because of the psychological effect of forming it. On this broader notion of having an interest, the matter of how many people are in the courtyard is not junk, for Evelyn.⁸ Even though this information is not useful or interesting to her, the act of forming the beliefs can serve her desire to clear her head. The problem with this way of understanding the junk view is that it gets a modified Study Break Case wrong:

No Break: Evelyn is studying for an exam. After hitting the books for several hours, she finds herself distracted and fatigued. Henry interrupts her studying to slowly count aloud the people in the courtyard. Evelyn waits impatiently as Henry informs her that there is at least one person in the courtyard, that there are at least two people in the courtyard, and so on. Evelyn returns to studying no less distracted and fatigued than before.

Intuitively, the beliefs Evelyn forms on Henry's testimony are clutter. They squander time and attention that she could have put to better use. Yet, these beliefs aren't junk, because she has a desire that is served by forming a belief about the subject matter (in the more calming, meditative way). The issue here is that junk is too coarse-grained, since it is defined at the level of subject matters. The junk view, in either form, will give the same verdict in Study Break and No Break, since they involve the same agent with

⁸ Strictly, each proposition of the form 'there are X people in the courtyard' and its negation comprise a distinct subject matter, on Friedman's usage.

the same belief contents and interests. Our intuitions about clutter, however, differ the between the two cases.

While Study Break and Reluctant Belief show that not all junk beliefs are clutter, there's also reason to think that junk beliefs aren't the *only* beliefs that are clutter. Consider the following case:

Plant Fan: Florencia spends several hours memorizing, for various plant species, whether they grow in New York. She finds each of these propositions slightly interesting. In order to learn all of these facts, she forgoes checking how her sister's job search is going, which she's deeply interested in, and learning a software she needs to know how to use for her job.

This is a massively unwise use of cognitive resources, all without forming a single junk belief. Clearly, Florencia's beliefs are clutter, but because of her slight interest, they are not junk. If what I've said here is right, the junk view can't capture our intuitions about which beliefs are clutter. Junk beliefs aren't always clutter, as in Study Break and Reluctant Belief, and some clutter isn't junk, like in Plant Fan. That is, while junk beliefs often contribute to clutter, clutter isn't just junk.⁹

1.3 THE WASTE VIEW

The reason the junk view misses the mark is that it does not take the cost of beliefs into account. As we saw in Study Break and Reluctant Belief, not every boring, useless belief taxes our cognitive resources. Further, Plant Fan shows that even slightly interesting beliefs can be clutter if they're not interesting or important enough to justify the cognitive resources they require. An adequate account of clutter, then, will take the cost of belief into account.

⁹ While the junk view comes apart from our clutter intuitions, this is no obstacle to Friedman's more general point that a truly interest-driven epistemology will look drastically different than traditional epistemology.

The waste view offers us a way to capture the value *and* the cost of a belief. The waste view is that beliefs that waste cognitive resources, or *wasteful beliefs*, are clutter. A belief is wasteful just in case its cost outweighs its value.

I think of the value of a belief as a graded analogue of interests, in Friedman's sense. That is, the value of a belief that P, for an agent, depends on the degree to which forming a belief about whether P serves the agent's desires and the degree to which the agent is interested in whether P. However, the waste view is compatible with other views about which beliefs have value (and how much), including veritism, the view that all and only true beliefs have value.¹⁰

The cost of a belief is most straightforwardly thought of in terms of the cognitive resources required to form it. In light of this, it may seem odd that it could "outweigh" the value of a belief. How are we to make sense of comparing a certain amount of time and degradation of attentional capacities to the value of a belief? The answer, I think, is that cognitive resources are only as valuable as what we can do with them. The cost of forming a given belief is equal to the value of the beliefs that the agent might have formed with the same cognitive resources, had they not formed the belief in question.¹¹ To put the point in a more forward-looking way, the cost of forming a belief is the value of the beliefs an agent would have to forgo to form it. In this way, we can directly compare the cost and the value of a belief, because cost is understood in terms of the value of beliefs.¹²

¹⁰ See Berker (2013) for illuminating discussion of forms of veritism, which he attributes to Alvin Goldman (Goldman 1999). I opt for a notion more in line with Friedman's here because I think it better captures the thought behind Harman's Clutter Avoidance Principle and is more friendly to the Clutter Argument.

¹¹ I take "the beliefs an agent might have formed" to be restricted to close possible worlds. I take the cost of a belief to be determined by the highest value belief the agent might have formed (in a close possible world), rather than the sum of such beliefs.

¹² Why stop at the value of *beliefs* we forgo when determining the cost of a belief? After all, our time and attention are necessary for most of what is valuable, from close relationships to learning to personal hygiene. Similarly, Hannah Trees pointed out to me that everyday activities, like washing the dishes, could detract from forming valuable beliefs. Both of these issues arise because the mandate to use cognitive resources wisely is also important in the practical domain. I restrict the cost of a belief to the value of beliefs forgone, rather than other activities, because we need clutter and the clutter avoidance principle based on it to be plausibly epistemic in character. This is because if clutter avoidance turns out to be a

The waste view is radically agent- and context-sensitive. The value of a belief depends on the agent's interests. The cost of a belief depends on two things. First, it depends on how the belief was formed (or would be formed). Consider some proposition Q, which follows from my beliefs by a complicated proof and which has some small amount of value. Were I to form a belief that Q on the basis of expert testimony, the cost would be minimal, since forming beliefs by expert testimony requires very little in the way of cognitive resources (on my part). However, were I to spend the cognitive resources required work through the complicated proof myself, the cost would be much higher. So, for the very same agent and belief content, a belief could be wasteful or not wasteful depending on how it was formed. Second, the cost of a belief depends on the value of the beliefs an agent might have formed with the same resources. This, in turn, depends on the agent's abilities, which determine which beliefs they might have formed, and the agent's interests, which determine the value of the beliefs they might have formed. Both of these vary from agent to agent, and from time to time for the same agent, since an agent's interests, desires, and which information best serves their desires can change.

The waste view is incredibly sensitive to features of agents and their situations. This is because our clutter intuitions track the cost of belief, which depends on how the belief was formed and the agent's abilities and practical situation. As a result, the waste view can accommodate our intuitions where the junk view struggles. In *Reluctant Belief*, Paul has already done all of the deliberating about P necessary to believe it (though perhaps he should have avoided doing so). Having done all that, there is no further cost to believing P, and so his belief is not wasteful. In *Study Break*, the cost of Evelyn's beliefs is negative—her break replenishes her cognitive resources and allows her to learn more quickly. So, her beliefs are not wasteful. Evelyn and Paul's beliefs are not wasteful because their cost does not outweigh their value, even though their beliefs are not particularly valuable.

practical norm, not an epistemic one, it won't be possible to make a compelling Clutter Argument, for reasons I discuss in Section 2.1.

Similarly, the waste view, unlike the junk view, can make sense of how beliefs that have some value can still be worth avoiding on clutter grounds. In Plant Fan, Florencia's beliefs about which plants grow in New York have some value. However, by spending her time and attention on those beliefs, she forgoes beliefs that she finds more interesting about her sister's job search and beliefs about the software she needs to learn for work that better serve her practical aims. Her plant beliefs are valuable, but the beliefs she misses out on are much moreso. Her plant beliefs, therefore, are wasteful.

In the following section, I discuss how we should make sense of the Clutter Avoidance norm that Harman proposes. In doing so, I rely on the waste view.

2: The Normative Significance of Clutter

What demands do clutter considerations make on our epistemic practices?
Harman suggests the following principle:

Clutter Avoidance: One ought not clutter one's mind with trivialities. (1989, p.12)

This principle leaves several important issues open. Can I clutter my mind by believing a single triviality or do I clutter my mind only when I believe too many trivialities? Is Clutter Avoidance a norm itself, or does it merely constrain norms? The Clutter Argument purports to leverage inconsistency between Clutter Avoidance and logical closure norms to show that we must abandon the latter. This requires figuring out exactly which beliefs or norms conflict with Clutter Avoidance. For that, we need a more precise clutter principle.

In this section, I explore the normative significance of clutter, understood according to the waste view. I discuss a variety of clutter principles and attempt to illuminate the choice points available to clutter theorists. These considerations will lead to the clutter principle I use to formulate the Clutter Argument in Section 3.

2.1: ARE CLUTTER PRINCIPLES *REALLY* EPISTEMIC?

A common reaction to Clutter Avoidance is that it isn't *really* epistemic. Instead, the thought goes, it's a practical norm that happens to concern the epistemic end of (justified) belief formation.¹³ Whether something is clutter doesn't have anything to do with whether it's true or likely to be true. Similarly, the formal and evidential relationships between beliefs have no bearing on whether a belief is clutter or not. What's worse, the view that best captures our intuitions about clutter, the waste view, looks suspiciously like an expected value calculation.

If clutter were, at bottom, practical, this would pose a serious problem for the Clutter Argument. The Clutter Argument, roughly, says that clutter norms conflict with logical closure norms, and so one of them—the logical closure norm—has to go. If clutter and closure norms are from different normative domains, however, then this conflict doesn't mean much. Norms from different domains conflict regularly, and this is not typically damning for the conflicting norms. For instance, the action recommended by prudential norms often differs from the action recommended by moral norms. Yet this fact is not, on its own, reason to reconsider any of the relevant prudential or moral norms. Similarly, a conflict between a practical and an epistemic norm gives us little reason to abandon either.

Luckily for fans of the Clutter Argument and enemies of closure norms alike, the case for clutter norms being non-epistemic is not as straightforward as it may seem. There are plenty of norms that are central to our epistemic lives that have some practical character: the norms of inquiry or zetetic norms.¹⁴ Making wise use of the resources we require to deliberate, investigate, and form beliefs is part of being a good inquirer. Inquiry is a central activity (if not *the* central) activity of our epistemic lives. How to characterize norms of inquiry and how to distinguish them from traditional epistemic

¹³ Roger White (2005) refers to clutter avoidance as a “practical reason” not to take certain attitudes (p.457).

norms (if indeed we can) is beyond the scope of this paper.¹⁵ For now, however, it suffices to say the practical flavor of clutter principles does not necessarily mean that they are not epistemic. In the next section, however, I explore a strategy to avoid the issue of whether clutter principles are *really* epistemic altogether.

2.2 THE WIDE WORLD OF CLUTTER PRINCIPLES

Even with a clear picture of which beliefs contribute to clutter, exactly what a plausible clutter principle looks like remains hazy. The first choice we face is whether to think of a clutter principle as a norm on rational belief or, as Harman suggests, a “metaprinciple that constrains norms on belief revision” Harman, *Change in View*. The second is whether a clutter principle should be strict or lenient. A strict clutter principle would prohibit *all* wasteful beliefs. A lenient clutter principle, on the other hand, would prohibit forming *too many* wasteful beliefs or wasting too many cognitive resources. Finally, if we choose a clutter principle that constrains norms of revision, as Harman suggests, we must decide whether our principle will rule out only norms that *require* clutter, or norms that merely *permit* clutter, as well. I’ll address each of these choice points in turn.

Harman suggests that Clutter Avoidance is not a norm itself, but a principle that rules out norms of revision, or dynamic norms. Such a constraint on dynamic norms is not a norm itself, but a theorist’s tool that guides us in determining what the dynamic norms are (or could be).¹⁶ It tells us not what to believe, but which candidate norms are not correct norms.¹⁷ Harman offers the following case in support of understanding Clutter Avoidance as a constraint on dynamic norms:

¹⁵ Jane Friedman’s “The Epistemic and the Zetetic” (forthcoming) is an illuminating discussion of norms of inquiry and how they relate to traditional epistemic norms.

¹⁶ Sinan Dogramaci helpfully noted that “metanormative” and “metaprinciple” are inapt (since a constraint on dynamic norms is neither a norm about norms or a principle about principles) and unnecessary. I’ll refer to a principle that rules out candidate dynamic norms as a “constraint on dynamic norms” or simply a “constraint.”

¹⁷ I’m using “correct norm” for a norm that *actually* governs us or our beliefs, distinct from something which has the form of a norm, which I’ll call a “candidate norm.”

Suppose George is trying to convince Bob that P. George shows Bob that P is a deductive consequence of things he believes. Bob accepts the validity of George's argument and refuses to change his belief in any of the premises, but he also refuses to accept the conclusion P, citing clutter avoidance as his reason for refusing. (Harman 1986)

Bob cites clutter considerations as a reason not to believe something which, by his own lights, his beliefs strongly support, a defense which Harman rightly calls "absurd." The thought is that norms are the sort of thing we appeal to in defense of our attitudes and make use of in explicit reasoning. Constraints on dynamic norms on the other hand, are theorist's posits, and therefore not the kind of think we should appeal to in defense of our beliefs. Bob's appeal to clutter avoidance is inappropriate because clutter avoidance is not a norm, but a constraint on norms.¹⁸

Opting for a constraint rather than a norm has a further advantage: it allows us to sidestep the question of whether a clutter principle is properly epistemic. First, if we think of constraints as principles that guide our theorizing, it's not clear that a constraint belongs to a normative domain at all. Second, even supposing that a clutter constraint is practical in some sense, this is no obstacle to its constraining our theorizing about epistemic norms. Practical concerns constrain our theorizing in a number of normative domains such as the moral and the legal. For instance, an "ought-implies-can" principle constrains moral theorizing, yet it is at bottom a practical constraint. So, if a clutter principle is a constraint on norms, not a norm itself, it can constrain the epistemic, whatever practical character it may have.

¹⁸ It's not entirely obvious that a constraint on dynamic norms does fare better in this respect. Such constraints might just as easily figure in explicit reasoning. Bob could have defended his refusal to believe P on the grounds that it's trivial and he has no interest in it, and therefore the correct dynamic norms don't require that he believe it. Whether this would be appropriate or not depends on how our answers to theoretical questions in epistemology bear on what we should believe, an issue I won't attempt to resolve here.

The second choice point we face is whether our clutter constraint is strict or lenient. A strict clutter constraint would rule out any candidate norm that requires agents to form any wasteful beliefs. A lenient constraint would allow candidate norms that require agents to waste *some* cognitive resources, but rule out norms that require agents to waste too much.¹⁹ I think the answer here must be lenient for two reasons. First, consider a norm that requires that we believe only the non-wasteful consequences of our beliefs except: we must believe ‘P or not-P’ (for some P). This is an unnatural, gerrymandered norm, to be sure, but the question here is whether we want to rule out this norm on clutter grounds. For some agents, this norm will not require any wasteful beliefs, while for others it will require only one. It seems odd to rule out this norm on clutter grounds because a single, simple wasteful belief doesn’t doom an agent to suffer the bad consequences of clutter. A single wasteful belief does not waste much time, and it won’t, on its own, exhaust our attentional capacity. A norm can require a single wasteful belief without requiring that agents suffer any of the bad consequences that come with forming too many wasteful beliefs. If it’s these bad consequences that explain why we should avoid clutter, then a single wasteful belief here or there is not enough to rule out a candidate norm on grounds of clutter.

There is also dialectical reason to take the clutter constraint to be lenient. As Friedman notes in her discussion of Clutter Avoidance, if we rule out every norm that requires even a single clutter-y belief, then Clutter Avoidance rules out most of our familiar epistemic norms, like evidentialist norms and more modest closure norms. For instance, consider an evidentialist norm that requires us to believe whatever is well-supported by our evidence. For most of us, there are at least some propositions that are well-supported by our evidence but would be wasteful to believe. If we want a convincing Clutter Argument, then, we should avoid a clutter constraint so strong it rules out most of the norms epistemologists are interested in. So, we should take our clutter constraint to be lenient.

¹⁹ I’ve switched from talk of “wasteful beliefs” to “wasted cognitive resources” because the relevant question is not how many wasteful beliefs, but how much is wasted in terms of cognitive resources.

The final choice we must make is whether our clutter constraint should rule out norms that merely *permit* wasting cognitive resources, or only norms that *require* wasting cognitive resources. John MacFarlane (2004) and Sinan Dogramaci (2018) favor the idea that well-supported but trivial or boring beliefs are permissible but not required. The “requirement only” approach captures the idea that following epistemic norms should be consistent with avoiding clutter. We can always follow norms that issue permissions and avoid clutter by simply avoiding the wasteful permitted beliefs. On the other hand, if it’s bad to waste cognitive resources, why should the correct dynamic norms permit us to do so? Friedman argues that clutter avoidance should be understood as ruling out norms that permit clutter (or junk, on Friedman’s view) as well as those that require clutter.

Harman’s compelling thought isn’t that you don’t have to believe all that junk, but that it’s perfectly rationally permissible for you to, rather his thought is that you should avoid believing junk. That’s to say that it’s not rational to clutter your mind, and you shouldn’t do it. But the thought that believing junk is permitted, but not required, leaves it perfectly rationally permissible to clutter your mind. (2018, p.572)

The contrast between a “requirement only” and a “requirement and permission” clutter principle is an important issue. However, the larger goal of this paper is to formulate and evaluate the Clutter Argument against a logical closure norm that issues requirements. If such a norm requires us to clutter our minds, the Clutter Argument will go through whether or not the clutter principle rules out norms that permit clutter. While I’m most sympathetic to Friedman’s “requirement and permission” approach here, I’ll use the more modest clutter principle that only rules out norms that issue requirements in hopes of formulating the Clutter Argument in a way that appeals to the broadest audience.

Our clutter principle, then, is a lenient constraint that rules out candidate dynamic norms that require agents to waste too many cognitive resources. That is, our clutter principle will look like the following:

Clutter Constraint: If a dynamic norm requires one to waste significant cognitive resources, then it is not a correct dynamic norm for one.

Exactly which norms run afoul of Clutter Constraint is somewhat thorny. The dynamic norms we’re concerned with require forming beliefs with certain *content*. But as noted in Section 1.3, whether a belief is wasteful depends on how it is formed, in addition to its content. That is, for the same agent at the same time, a belief with a given content would be wasteful if formed in some ways, and not wasteful if formed in others. This feature is an advantage of the waste view—our intuitions about which beliefs are clutter are sensitive to how beliefs are formed. Unfortunately, though, it does make it difficult to say whether forming a belief with a given content will waste an agent’s cognitive resources. This, in turn, makes it difficult to determine which norms Clutter Constraint rules out.

I offer the following heuristic to test whether a norm will be ruled out by Clutter Constraint. For a candidate dynamic norm, consider how an agent typically follows it, when they do. If an agent would waste significant cognitive resources consistently following a candidate norm in the way they typically follow it (when they do), then Clutter Constraint rules out that norm.

Determining how significant a waste is significant will inevitably require recourse to our intuitions about which uses of cognitive resources are good and which are not. However, as a rule of thumb, we can say that a dynamic norm should be ruled out on clutter grounds when conforming with it would require an agent to waste significant cognitive resources, where a “significant” waste is understood to be enough cognitive resources that it hampers or disturbs of some normal or desirable epistemic practice. That is, when clutter gets in the way, it’s too much clutter.

3: The Clutter Argument

In this section, I formulate what I take to be the most compelling version of the Clutter Argument and argue that it fails. In 3.1, I address the question of which logical closure norm is the target of the Clutter Argument. In 3.2, I formulate the Clutter Argument. In 3.3 I consider various principles that might motivate the crucial premise of the Clutter Argument. Finally, in 3.4 I argue that we have reason to think that the Clutter Argument is unsound.

3.1 THE TARGET OF THE CLUTTER ARGUMENT

Harman offers the following principle as the target of the Clutter Argument:

Logical Closure Principle: One's beliefs should be "closed under logical implication." In other words there is something wrong with one's beliefs if there is a proposition logically implied by them that one does not already believe. In that case one should either add the implied proposition to one's beliefs or give up one of the implying beliefs. (Harman 1986)

Logical Closure principle is, in fact, two distinct norms. One is a static norm which tells us how our beliefs should be, at a given time:

Static Closure: One's beliefs should be closed under logical implication.

The other norm is a dynamic norm that tells us how we should revise our beliefs, having found that we fail to meet Static Closure:

Dynamic Closure: If there is a proposition implied by one's beliefs that one does not already believe, then one should either add the implied proposition to one's beliefs or give up one of the implying beliefs.

Dynamic Closure and Static Closure differ beyond being dynamic and static, respectively. Static Closure is an *ideal* norm, in that its requirement does not take human limitations into account. Dynamic Closure, on the other hand, is a *non-ideal* norm. It kicks in, so to speak, only when an agent has failed to meet an ideal norm, Static Closure. As Declan Smithies puts it, non-ideal norms are “approximations towards ideal standards that take our contingent human limitations into account” (2015). Non-ideal norms capture our “ordinary standards of rationality” on the basis of which we praise or blame human agents (all of whom fail to meet the ideal standards). While the non-ideal norms are, of course, guided by the ideal norms, as we’ll see in Section 3.3, the relationship between them is difficult to articulate.

The most important question for our current purposes is which closure principle will make for the most compelling, dialectically effective argument. The answer is Static Closure.

First, those who are interested in logical closure are typically interested in Static Closure, not dynamic closure. Theorists of ideal rationality have typically conceived of ideal rationality in terms of how the ideally rational agent would believe (at a time) and supposed this will constrain how ideal and non-ideal agents should reason, but not the other way around. David Christensen suggests that this is because the fact that something is wrong with one’s beliefs at a given time is what motivates revising one’s beliefs at all (2004). Further, the formal systems that often motivate strong closure norms rely on static, not dynamic, closure principles. Formal models of belief change, such as dynamic epistemic logic or AGM models of belief revision, offer *ideal* dynamic norms for agents who already meet Static Closure, rather than non-ideal norms like Dynamic Closure.

Second, a Clutter Argument against Dynamic Closure is unlikely to work. Dynamic Closure requires that, if there is some proposition implied by an agent’s beliefs, that they either come to believe it, or give up the implying beliefs. This means, however, that an agent can perfectly well accord with Dynamic Closure by relinquishing

ever more of their beliefs.²⁰ Continual pruning of one's beliefs will not "clutter one's mind with trivialities," as Harman feared. While I'll argue shortly that the Clutter Argument against Static Closure fails as well, it does at least seem fairly promising. After all, Static Closure requires that we believe infinitely many things, many of them boring and useless.

The upshot of these considerations is that an airtight argument against Dynamic Closure would be of little consequence to friends of closure norms, who tend to be committed to Static Closure or similar static norms. Instead, an effective Clutter Argument must take Static Closure as its target.

3.2 THE CLUTTER ARGUMENT

As I discussed in Section 2, most plausible clutter principle, Clutter Constraint, rules out dynamic norms, not static norms. The target of the Clutter Argument, however, is Static Closure, a static norm. How then, is the Clutter Argument meant to go? Harman does not address this issue, possibly because his own Logical Closure Principle included a dynamic norm. However, I'll offer a strategy for overcoming this issue in formulating the Clutter Argument: employing a principle about how the non-ideal dynamic norms relate to the ideal norms.

Static Closure, on its own, does not imply anything about how we should revise our beliefs. However, it seems clear that committing to a norm like Static Closure commits us to *something* about what the non-ideal dynamic norms are like. I'll refer the dynamic norms that must be true if Static Closure is, whatever they may be, as 'the Non-

²⁰ Logical truths somewhat complicate the matter here. While logical truths are implied by anything, even the empty proposition, actual human could continue conforming with Dynamic Closure for the whole of their lifetime without running out of beliefs to jettison. Second, it's not clear that one's *beliefs* imply anything if one has no beliefs. The empty proposition implies all logical truths, but the agent need not believe the empty proposition.

Ideal Dynamic Norms’. By doing this, we can talk generally about features these norms must have while remaining neutral on exactly which norms they are.

With this bit of shorthand in place, we can formulate the argument from Clutter Constraint against Static Closure.

The Clutter Argument

1. If a dynamic norm requires wasting significant cognitive resources, then it is not a correct dynamic norm. (Clutter Constraint)
2. If Static Closure is a correct norm, then the Non-Ideal Dynamic Norms are correct dynamic norms.
3. The Non-Ideal Dynamic Norms require wasting significant cognitive resources.
4. The Non-Ideal Dynamic Norms are not correct norms. (3, 1)
- C. Static Closure is not a correct norm. (2, 4)

In other words, if the Non-Ideal Dynamic Norms would require that one waste significant cognitive resources, then they are ruled out by Clutter Constraint. If the Non-Ideal Dynamic Norms aren’t correct norms, then neither is Static Closure.

It’s worth noting that what’s wasteful for one person may not be wasteful for an agent with different capabilities and interests. I’ll assume in what follows that a norm “requires wasting significant resources” if a human with typical capacities would waste significant cognitive resources by following it in the way they typically do (to the extent that they do).

The crucial question in evaluating the Clutter Argument is whether we have reason to accept Premise 3. Premise 1 is the clutter principle that I argue for in Section 2. Premise 2 is trivial given how I’ve stipulated the meaning of “Non-Ideal Dynamic Norms.” Premise 3 is the claim that believing as we ought to, if Static Closure is true, will waste significant cognitive resources. This the key claim of the Clutter Argument: that Static Closure requires agents to clutter their minds.

3.3: THE DYNAMIC AND THE STATIC NORMS

Premise 3 is a claim about the Non-Ideal Dynamic Norms—that is, about the non-ideal dynamic norms that *must* be the case if Static Closure is true. However, there are plenty of non-ideal dynamic norms that are obviously related to Static Closure that clearly need not waste significant resources. For instance, “form a belief in a logical consequence of your beliefs just in case it would not be wasteful.” Or consider, “if you are inquiring about whether P and P follows from your beliefs, believe that P,” which leaves open that you shouldn’t be inquiring about matters that are not sufficiently interesting or practically important.

This suggests a response to the Clutter Argument for advocates the Static Closure: simply insist that the Non-ideal Dynamic Norms are among the many norms that, in some sense, aim at Static Closure, but do not waste significant cognitive resources. Smithies, an advocate of Static Closure as an ideal norm, characterizes non-ideal norms as “approximations towards ideal standards that take our contingent human limitations into account” (2015). Our limited attention, time, and ability to retrieve information from memory are among our contingent human limitations. On this picture, then, it seems perfectly natural to say that the Non-Ideal Dynamic Norms will not require us to waste significant cognitive resources, for if they did, they would not successfully take our human limitations into account.

To block this sort of response we could, of course, stipulate that taking our limitations into account is only a matter of taking our abilities into account. The non-ideal norms are norms that, under normal conditions, we can abide by, but doing so may well be very costly in terms of cognitive resources. Without some reason to consider our limitations in this narrow way, however, this response seems ad hoc. In order to motivate Premise 3 in a more principled way, we need a “bridge principle” that characterizes the connection between static ideal norms and non-ideal dynamic norms in general.

One straightforward bridge principle says that the non-ideal dynamic norms should get agents to meet the ideal static norms:

Success: Consistently following the non-ideal dynamic norms will lead a human agent to meet the ideal static norms.

Success relies on an ought-implies-can principle about ideal static norms. There are no dynamic norms that will get a human agent to meet Static Closure, because no human agent can meet Static Closure. This poses two problems. First, defenders of Static Closure (and other strong closure and consistency norms) are well aware that human agents can't meet Static Closure, yet nonetheless think Static Closure is a correct norm.²¹ That is, no defender of Static Closure would accept Success. If we want to convince people who think Static Closure might be right, we can't rely on Success. Second and more importantly, Success isn't sufficient to make Premise 3 true. Since there are no dynamic norms that will allow a human agent to meet Static Closure, 'the Non-Ideal Dynamic Norms' has no referent. This will make Premise 3 false (or indeterminate, depending on one's view of non-referring terms).

We can capture the general spirit of Success with a weaker principle:

Long-Run Success: The non-ideal dynamic norms lead to the ideal static norms in the long-run.²²

Long-Run Success does not imply that any human agent can meet the static norms, since we will all stop following the dynamic norms, at some point. Long-Run Success, applied

²¹ For an illuminating discussion of why many reject epistemic ought-implies-can principles (and alternatives to doing so), see Smithies (2005).

²² The construction "lead to" may seem a bit odd. I'm supposing here that the dynamic norms of some sort of form (logical or otherwise) that allows us to abstract somewhat from specific followings by specific agents. For instance, if the dynamic norms are things like modus ponens and disjunctive syllogism, it seems perfectly coherent to talk about where these norms "lead" beyond human capabilities.

to Static Closure, says that the Non-Ideal Dynamic Norms would, if followed forever, require that an agent believe every logical consequence of their beliefs. For a mortal human agent, Long-Run Success means that the Non-Ideal Dynamic Norms will require that they believe ever more logical consequences, but not all of them.

A weaker bridge principle than Long-Run Success will not be sufficient to motivate Premise 3. Take, for instance, the following (still quite strong) bridge principle:

Improvement: Following the non-ideal dynamic norms will leave one always better off with respect to the ideal static norms.

Consider the dynamic norm “believe all the consequences of your beliefs except for the wasteful ones.” Following this norm would leave you always better off with respect to Static Logical Closure. However, forming beliefs in only non-wasteful consequences would, obviously, not waste significant cognitive resources. So, Improvement fails to support Premise 3. A stronger bridge principle than Long-Run Success, however, risks having “the Non-Ideal Dynamic Norms” fail to refer, as is the case with Success.

I argue in the following section that, even if Long-Run Success is true, Premise 3 is false. However, it’s worth considering a preliminary reason to think that Long-Run Success is not sufficient to make Premise 3 true. If Long-Run Success is right, the Non-Ideal Dynamic Norms may well have familiar logical forms, like “if you believe P and believe ‘if P then Q’ and Q is not absurd, you ought to form a belief that Q” or “if you believe P, believe ‘P or Q’.” It seems natural that that norms modeled on natural deduction rules would lead to Static Closure in the way that Long-Run Success requires. These norms require that I believe certain logical consequences of my beliefs. They don’t, however, require that I form these beliefs *fast*, that I form them *soon*, that I give up other pursuits to form them, or that devote lots of time to forming beliefs by this modus ponens norm. Because Static Closure itself is so demanding, there is a tendency to think that the Non-Ideal Dynamic Norms must require us to chug away endlessly like the mirthless logician. However, Long-Run Success leaves open that the Non-Ideal Dynamic

Norms are familiar norms that look very much like modus ponens and disjunction introduction. These norms do require quite a lot, but it's not clear that we couldn't follow them consistently for the whole of our lifetimes without running out of interesting, useful, non-wasteful beliefs to form.

3.4 IMPLICIT BELIEF AND THE CLUTTER ARGUMENT

Premise 3 says that the Non-Ideal Dynamic Norms will require that one waste significant cognitive resources. I suggested in the previous section that if Long-Run Success is correct, the Non-Ideal Dynamic Norms will require that we believe ever more consequences of our beliefs, but not *all* of the consequences of our beliefs. That is, the Non-Ideal Dynamic Norms lead to Static Closure *in the long-run*, but actual human agents can only follow these norms for some finite amount of time. In order to evaluate Premise 3 for human agents then, we need to look at how actual human agents go about following norms that require believing logical consequences of their beliefs, when they do, and ask whether these norms will require wasting significant cognitive resources. I use a familiar norm based on disjunction introduction as examples, but nothing I say depends on this particular norm being among the Non-Ideal Dynamic Norms.

We typically believe logical consequences of our beliefs implicitly, not explicitly. Explicit beliefs are beliefs the contents of which are represented in the mind, whether in conscious thought or in memory. Implicit beliefs are beliefs that are not explicitly represented in the mind.²³ You believe (and believed before reading this) that the table nearest to you does not have a gooey nougat center, despite the fact that until a moment ago, you likely had never so much as entertained it. Since you likely had never explicitly represented in your mind the proposition that that the table nearest to you does not have a gooey nougat center, this belief was implicit (though it's explicit now). For the purposes

²³ The implicit/explicit distinction is often confused with the dispositional/occurrent distinction, but they are importantly different. Roughly, beliefs are occurrent iff they are (currently) the content of a conscious thought, and dispositional otherwise. All occurrent beliefs, therefore, are necessarily explicit, since being the content of an occurrent thought is one way to be explicitly represented in the mind. Dispositional beliefs, however, may be explicit (in the case of beliefs stored in memory) or implicit.

of the present discussion, I'll assume the view, due to Hartry Field (1978), that we implicitly believe the obvious consequences (deductive or otherwise) of our explicit beliefs, though nothing I say hangs on this account of implicit belief.²⁴ Whether an agent implicitly believes P depends on whether it is sufficiently obvious *to her*, not on whether it follows from her other beliefs by an obvious step.

To see that our beliefs about logical consequences of our beliefs are by-and-large implicit, consider the subset of consequences that follow by disjunction introduction. When I come to explicitly believe some proposition P, I come to implicitly believe obvious consequences of P. Since many (but not all) instances of disjunction introduction are obvious to me, I will come to implicitly believe a massive number of disjunctions with P as a disjunct. In fact, I'll come to implicitly believe infinite disjunctions with P as a disjunct, because there are infinite disjunctions of P and other sentences that I can easily understand and any of these will be obvious to me, given my general understanding of disjunction introduction (though this won't be true for agents without this general understanding).²⁵ The general picture here is that for each of our explicit beliefs, we have a massive, in some cases infinite, number of implicit belief about their obvious consequences. No matter how many consequences of our beliefs we explicitly believe, then, we'll implicitly believe many, many more.

The question at hand is whether following the Non-Ideal Dynamic Norms involves wasting significant cognitive resources, as Premise 3 says. Since we overwhelmingly believe logical consequences of our beliefs implicitly, the question of whether the Non-Ideal Dynamic Norms require us to waste significant cognitive resources amounts to the question of whether implicitly believing logical consequences of our beliefs involves wasting significant cognitive resources. We form implicit beliefs

²⁴ For other accounts of implicit belief, see Daniel Dennett (1975) and Mark Crimmins (2002).

²⁵ In some cases it will be indeterminate which norm an agent is following by implicitly believing. To see this, consider consequence, Q, that follows from my explicit beliefs by two distinct patterns of reasoning, each of which would be sufficient to make Q obvious to me. We can imagine that in some contexts in which the matter of whether Q came up, I might use one pattern of reasoning to get to Q, while in other contexts, I might use the other.

in two main ways: by forming new explicit beliefs and by improving at reasoning.²⁶ In arguing that implicitly believing ever more logical consequences doesn't require significant cognitive resources, I'll consider in turn the ways that we form new implicit beliefs.

The first way we form implicit beliefs is by improving at reasoning. As we improve at employing certain patterns of reasoning, more and more consequences that follow by that pattern become sufficiently obvious as to be implicit beliefs. For instance, consider a beginning logic student who sometimes uses disjunction introduction in reasoning. He learns in logic class that disjunction introduction is a good rule in general, and practices recognizing consequences that follow by disjunction introduction. The more adept he becomes at recognizing these consequences, the more obvious such consequences become. So, by getting better at applying this rule, he'll come to believe huge swaths of consequences that follow by disjunction introduction. Improving at reasoning requires some time and attention. However, we can improve at reasoning by doing things like balancing our budget, playing chess, or reasoning carefully about things that are of great interest to us. Cognitive resources put to these uses are not wasted, because these activities and the beliefs formed while doing them are valuable to the agents that pursue them. If we can form implicit beliefs about logical consequences of our beliefs, as the Non-Ideal Dynamic Norms require, by doing whichever reasoning-improving activity we like, then it does not appear we'll waste significant cognitive resources by following the Non-Ideal Dynamic norms.

The second way we form new implicit beliefs by forming new explicit beliefs. When I form an explicit belief that R, I come to implicitly believe obvious consequences of R (that weren't obvious consequences of my beliefs before believing R). Forming new explicit beliefs typically requires at least a little time and attention. The question then is whether we need to form many wasteful explicit beliefs in order to follow the Non-Ideal

²⁶ We also form new implicit beliefs by conceptual enrichment. If I lack the "maple" concept, and then gain it, I'll thereby gain lots of new implicit beliefs involving that concept (like "nothing made of glass is a maple" and "the oak in my backyard is not a maple"). I've left implicit belief formation by conceptual enrichment out of my discussion, because it doesn't apply to our new beliefs in logical consequences.

Dynamic Norms. The answer is no, because we can follow the Non-Ideal Dynamic Norms by improving at reasoning and by forming non-wasteful explicit beliefs about matters that interest us. This won't, in the end, lead us to believe *all* of the consequences of our belief, but all the Non-Ideal Dynamic Norms (according to Long-Run Success) is that we continue believing more consequences of our beliefs. This, it seems, we can do without wasting significant cognitive resources.

One might object here that we are not able to get ever closer to meeting Static Closure without running into a proposition that requires us to waste some cognitive resources. At some point, the consequences that can be believed implicitly will run out and we will be left inferring away like the mirthless logician. However, it's important to note that our beliefs are not fixed over time. We will continue to form more beliefs about our surroundings, daily lives, the news, and so on. Each of these new beliefs will have a vast number of consequences, which we can come to believe by following the Non-Ideal Dynamic Norms (implicitly or explicitly). So, there is little reason to think that we might reach a point where, in order to follow the Non-Ideal Dynamic Norms, we would have to start believing ever more difficult or complicated consequences of our beliefs.²⁷

Premise 3 of the Clutter Argument says that the Non-Ideal Dynamic norms will require agents to waste significant cognitive resources. However, when we look at how we in fact come to believe most of the logical consequences of our beliefs that we do, it's not at all clear that this is the case. We believe new logical consequences of our beliefs mostly implicitly, and we can form these beliefs without significant wasted time or attention. Therefore, Premise 3 is false, and the Clutter Argument is unsound.

For those who would defend the Clutter Argument despite this objection, a clear strategy presents itself. Find a bridge principle between the ideal static and non-ideal dynamic norms that is plausible and delivers the result the Non-Ideal Dynamic Norms will clutter the mind of humans like us, and you have a successful Clutter Argument. Nonetheless, I'm not optimistic that a bridge principle strong enough to make the Clutter

²⁷ Snow Zhang suggested this line of reply in conversation.

Argument work would be plausible. My own attempts to find one have not been fruitful and it's difficult to see how such a principle would remedy the issue short of outright requiring explicit belief.

4: Conclusion

Most philosophers who've addressed clutter considerations have expressed support for or explicitly endorsed the idea that clutter provides a reason to reject or rethink logical closure norms.²⁸ I've argued that this is a mistake. When we look at how we in fact believe logical consequence of our beliefs, it's clear that making consistent progress towards Static Closure does not put us in danger of cluttering our minds. This is because we typically believe these things implicitly, which requires relatively little in the way of cognitive resources. What's more, one way we form these beliefs is by improving at reasoning, and it's difficult to see how this could be bad, from the point of view of rationality.

Static Closure is just about the strongest closure norm on offer. Those who reject it might nonetheless want to accept a weaker closure norm that requires that we believe all obvious or known consequences.²⁹ These weaker norms are subject to versions of the Clutter Argument, since not all obvious or known consequences will be interesting or of practical importance. However, there's reason to think that these versions of the Clutter Argument will fail for the same reason as the Clutter Argument against Static Closure. So, the arguments presented here may vindicate weaker closure norms, as well.

The Clutter Argument, were it successful, would leverage clutter considerations to show that Static Logical Closure could not apply to us human agents with limited time and attention. However, this is not the only way we might make sense of what our cognitive limitations impose on our epistemic practice. Advocates of Clutter Avoidance might make the same move that advocates of closure norms do: treating it as an ideal.

²⁸ In fact, to my knowledge, all of the philosophers who have discussed the Clutter Argument are fairly sympathetic towards it.

²⁹ "Known consequence" here means a proposition that an agent knows to follow from her beliefs, not a consequence of her beliefs that is already known (since it would be trivial to require belief for these).

That is, we might consider it ideal in some sense to *never* have a belief about something in which we have no interest or that wastes cognitive resources. This however, would not in general mean that an agent who believes uninteresting propositions is thereby irrational in every case or that, all things considered, she ought not believe an uninteresting proposition. Instead, Clutter Avoidance (here understood as a norm rather than a theoretical constraint) would be a metric by which an agent could be evaluated. Clutter Avoidance and Logical Closure, on this picture, would be two ideals which need to be balanced against each other, rather than mutually inconsistent mandates.

Reasons-first epistemology provides another way to accommodate the spirit of clutter avoidance and of logical closure norms.³⁰ On such a picture, that forming a belief would be taxing on our cognitive resources (or that the subject matter is junk) is a reason that militates against forming it and perhaps which favors not opening the question at all. That a proposition follows from your beliefs, on the other hand, is reason to believe it. Both of these reasons are defeasible and can be outweighed by other reasons.

What I discuss here does not exhaust the interesting and underappreciated issues around clutter. For one, there is the issue of whether clutter is at bottom a practical matter, not an epistemic one, and if it is, to what degree clutter and other practical considerations can constrain rational belief. Further, as Friedman rightly notes, there is little reason to think that beliefs are the only attitude that can clutter the mind. Entertainings, wonderings, suspicions, doubts, and perhaps even hopes or worries can tax cognitive resources in just the same way beliefs do.

³⁰ I understand reasons-first epistemology here as presented by Jonathan Dancy (2000; 2018). T. M. Scanlon (1998) and John Skorupski (2010) take a similar approach.

CHAPTER 2: DISSOLVING THE TENSION BETWEEN THE EPISTEMIC AND THE ZETETIC

There has been a great deal of recent discussion of the relationship between familiar epistemic norms and the norms that govern inquiry, or zetetic norms. Jane Friedman argues that many norms of contemporary epistemology, like consistency norms or evidential norms, will conflict with the norms that govern inquiry.³¹ This gives rise to what David Thorstad calls the *reaction question*: as theorists of the epistemic, how should we respond to this tension?³² We could accept that there is such a tension within the epistemic and explain it should not worry us and indeed, is an unavoidable feature of the normative landscape, as Thorstad does. We could radically revise the epistemic norms to accommodate the zetetic norms and the prominent role of inquiry in our epistemic lives, as Friedman suggests. Finally, we could eliminate the tension by separating the zetetic and the epistemic, say, by denying that zetetic norms are properly epistemic. Call this strategy *separationism*. Separationism, at first blush, is an appealing option. Zetetic norms, after all, often recommend actions rather than doxastic attitudes and are sensitive to our desires and goals in a way that traditional epistemic norms are not.

Separationism, however, is not as appealing as it may seem. I present a puzzle for the separationist strategy. Any way of solving this puzzle, I argue, will eliminate the tension between zetetic and epistemic norms that motivated separationism in the first place, leaving separationism unmotivated. In Section 1, I explicate the tension between zetetic norms and a number of familiar dynamic epistemic norms. In Section 2, I present a puzzle for separationism. In Section 3, I survey a number of potential responses to the puzzle and explain why my favored solution is preferable to the others. In Section 4, I argue that any of the plausible solutions to the puzzle will eliminate the tension between the zetetic and the epistemic, leaving separationism unmotivated.

³¹ Jane Friedman, "The Epistemic and the Zetetic," *Philosophical Review*, forthcoming.

³² David Thorstad, "Inquiry and the Epistemic," *Philosophical Studies*, 2021.

1: Tension Between the Epistemic and the Zetetic

Friedman argues that the zetetic norms are in tension with a number of familiar epistemic norms.³³ She focuses on the Zetetic Instrumental Principle:

Zetetic Instrumental Principle (ZIP): If one wants to figure out a question Q, then one ought to take the necessary means to figuring out Q.

ZIP tells us how to proceed when we want, or need, to know the answer to some question – that is, when we inquire. A few clarifications are in order. First, we can restrict our attention to cases in which the aim of inquiry is not morally reprehensible, like investigating someone for purposes of blackmail, or inane, like trying to figure out how many blades of grass are growing from a certain stretch of sidewalk. Second, we'll suppose for now that the inquiries under discussion are important and relatively urgent. A more comprehensive account of the norms of inquiry may offer some insight into appropriate zetetic goals and inquiries undertaken on a whim or out of mild curiosity. However, for the purposes of this paper, we should understand ZIP as a norm concerning morally permissible, important, urgent inquiries.

Many inquiries require sustained concentration. If I want to know whether I need to submit a 198S form on my taxes, I'll need to focus intently on comparing the byzantine rules concerning who must submit a 198S with my financial information for at least some stretch of time. Friedman offers the example of trying to figure out how many windows the Chrysler Building has. If I want to find out, I'll have to avoid distraction long enough to finish my count. Focusing intently on the rules about who submit a 198S or counting the windows of the Chrysler building precludes attending to other things or engaging in

³³ Friedman, "The Epistemic and the Zetetic."

other activities. If ZIP is right that I should focus on my inquiry, then I ought not engage in activities that would derail my inquiry. We might express this though with the following principle:

Incompatible Alternatives (IA): If one ought to Φ at t , and one cannot both ϕ and ψ at t , then one is not permitted to ψ at t .

ZIP, together with Incompatible Alternatives, issues a number of prohibitions. If my aim is to figure out how many windows the Chrysler building has, then according to ZIP, I ought to focus on counting, at least enough to avoid losing count. Anything I cannot do while focusing on my counting to the required degree is not permissible, by Incompatible Alternatives. It's this feature of ZIP that generates conflict with familiar epistemic norms. I'll first discuss the tension between ZIP and permission-granting norms, then the tension with requirement norms. Consider the following norms:

Dynamic Evidential Permission (EP): If one has excellent evidence for P at t , then one is permitted to come to believe that P at t .

Dynamic Knowledge Permission (KP): If one is in a position to know that P at t , then one is permitted to come to know that P at t .

Each of these norms permits me to believe a wide swath of propositions which I have excellent evidence for or am in a position to know. However, while I'm counting the windows on the Chrysler building or working on my taxes, ZIP will rule many of these beliefs impermissible. Perhaps I can come to believe *some* of these without losing focus, but there will inevitably be some belief that EP or KP rule permissible that I cannot form without breaking my concentration.

This does not mean that it's impossible to conform to both ZIP and the dynamic permission norms. When they conflict, an agent can simply decline to believe the

proposition that ZIP forbids and EP or KP permits. We can describe the kind of tension we see here as *weakly incoherent*. Two norms are weakly incoherent iff they jointly imply a sentence of the following form:

(W) One is permitted to ϕ at t and one is not permitted to ϕ at t .

In other words, weakly incoherent norms disagree about the permissibility of some action or belief, but they do *not* (on their own) leave the agent with no permissible option. However, the tension is stronger when we turn to epistemic requirements:

Dynamic Evidential Requirement (ER): If one has excellent evidence for P at t , then one ought to come to believe that P at t .

Dynamic Knowledge Requirement (KR): If one is in a position to know that P at t , then one ought to come to know that P at t .

In cases where successful inquiry requires maintaining focus, and therefore not attending to one's evidence or surrounding to form new beliefs, ER and KR will require that one form beliefs that ZIP forbids. We can call this tension *strong incoherence*. A set of norms is strongly incoherent iff they imply a sentence of the following form:

(S) One ought to ϕ at t and one is not permitted to ϕ at t .

Strongly incoherent norms, unlike weakly incoherent norms, do leave agents with no permissible options. If I ought to believe that I've taken more than 30 plane rides in my lifetime, because I have excellent evidence for it, but I ought not believe the same proposition, because reviewing my evidence would distract me from determining whether I need to submit a 198S form. I could, of course, conclude my inquiry about taxes and

then come to believe that I've taken more than 30 plane rides, but according to EP, I ought to come to believe it *right then*.

ZIP, then, generates weak incoherence with dynamic permission norms and strong (and weak) incoherence with the dynamic requirement norms.

1. 1 THE SCOPE OF THE TENSION

Friedman's exposition of the tension between the epistemic and the zetetic focuses on ZIP, but ZIP is not the only zetetic norm that generates incoherence when combined with familiar epistemic norms. For instance, another class of paradigmatic zetetic norm, evidence gathering norms, will face a similar issue.³⁴ A norm that permits an agent to gather evidence at a time, supposing evidence gathering sometimes requires uninterrupted attention, will generate weak incoherence with dynamic requirement norms like ER and DKR. If an evidence-gathering norm generates requirements, as ZIP does, then we will see weak incoherence with the dynamic permission norms, and strong incoherence with the dynamic requirement norms.

Dynamic norms that tell us how to change our beliefs are incoherent with ZIP and other zetetic norms. There remains a question about how the static norms, which permit or require one to be in a certain state, will interact with the zetetic norms. The tension in this case, if there is one, goes by way of the dynamic norms, since it is *forming* beliefs, not merely having them, that is incompatible with maintaining focus. As Friedman notes, whether the tension extends to static norms will depend on the relationship between the static and dynamic norms. She suggests that we should expect dynamic and static norms to "move in lockstep". The thought here is that requirements to *have* a belief drive

³⁴ Richard J. Hall and Charles R. Johnson, "The Epistemic Duty to Seek More Evidence," *American Philosophical Quarterly* 35, no. 2 (1998): 129–39; Elise Woodward and Carolina Flores, "Epistemic Vigilance: In Defense of Epistemic Norms on Evidence-Gathering," [*Manuscript*], n.d.; Kelly Annesley, "[*Manuscript*]," n.d.

requirements to form that belief and the same goes for permissions. I argue in Section 3 that, despite appearances, there is no incoherence between ZIP and familiar dynamic norms. If this is right, then the static norms will not produce incoherence with ZIP either.

1.2 RESPONDING TO THE TENSION

There are three ways to respond to the weak and strong incoherence between ZIP and more commonplace epistemic norms. The first is to simply accept the tension. One way to do this is to accept that there are genuine epistemic dilemmas, that is, cases in which there is no permissible option.³⁵ Another is to endorse the focalist view that Thorstad advocates. According to focalism, the tension between the epistemic and the zetetic comes from inevitable level tension in our epistemic evaluation. The second, which Friedman endorses, is to radically revise the epistemic to eliminate the tension. The thought is that inquiry is central to our epistemic lives and the epistemic norms we posit should accommodate this core feature. The final option is to cleave the zetetic from the epistemic. The most natural way to do this is by insisting that the zetetic norms are *not* really epistemic, perhaps because they are practical norms that happen to do with ends that bear epistemic value. Call this option *separationism* about the epistemic and the zetetic domains.

Separationism is, I suspect, an attractive response for many. It allows us to retain our favorite epistemic norms without accepting incoherence. The cost is relegating the

³⁵ Nick Hughes argues at length that epistemic dilemmas not only exist, but are entirely commonplace. Nick Hughes, “Who’s Afraid Of Epistemic Dilemmas?,” in *Epistemic Dilemmas: New Arguments, New Angles*, ed. Scott Stapleford, Mathias Steup, and Kevin McCain, forthcoming; Nick Hughes, “Epistemic Dilemmas Defended,” in *Epistemic Dilemmas*, ed. Nick Hughes, forthcoming; Nick Hughes, “Dilemmic Epistemology,” *Synthese* 196, no. 10 (2019): 4059–90, <https://doi.org/10.1007/s11229-017-1639-x>.

norms governing evidence gathering, double checking, expert consultation, and other zetetic activities to the practical domain. Given the sensitivity to our desires, interests, and practical constraints that ZIP and norms like it display, perhaps this cost is not so high after all. If we are to pay this cost, however, separationism had better successfully eliminate the incoherence within the epistemic domain.

2: A Puzzle for the Separationist

There is a further obstacle for the separationist strategy: the epistemic is strongly incoherent on its own, without any zetetic norms gumming up the works. This is not simply to say that the set of norms advocated by various contemporary epistemologists are strongly incoherent, though this is no doubt true. Rather, many dynamic requirement norms are strongly incoherent with themselves. We don't need to have a full accounting of the correct epistemic norms to see that a number of them will produce strong incoherence.

Consider the following case. At time t , I have excellent evidence for P and excellent evidence for Q , though I don't yet believe either proposition. It would take me a couple minutes of reflecting on my evidence to come to believe P (in the typical way, rather than by taking a pill or sustaining a blow to the head) and likewise for Q . Because of this, I'm not able to come to believe P and come to believe Q at the same time. That, is, coming to believe P and coming to believe Q are incompatible alternatives. Because I have excellent evidence for P at t , ER will say that I ought to come to believe that P . By stipulation, I cannot both come to believe P and come to believe Q at t , so by the Incompatible Alternatives principle, I am not permitted to come to believe Q at t . But

I have excellent evidence for Q as well, so according to ER I ought to come to believe Q at t, as well. A bit more formally:

1. **The Case:** One has excellent evidence for P and excellent evidence for Q. One cannot both form a belief that P at t and form a belief that Q at t.
2. **Dynamic Evidential Requirement (ER):** If one has excellent evidence for P at t, then one ought to come to believe that P at t.
3. **Incompatible Alternatives:** If one ought to ϕ at t, and one cannot both ϕ and ψ at t, then one is not permitted to ψ at t.
4. One ought to come to believe that P at t (1, 2).
5. One is not permitted to come to believe that Q at t. (3, 4)
6. One ought to come to believe that Q at t. (1, 2).
- C. One ought to come to believe that Q at t and one is not permitted to come to believe Q at t. (5, 6)

This argument shows that ER gives rise to strong incoherence. Recall that a set of norms is strongly incoherent if, and only, if they imply a sentence of the form "one ought to ϕ at t and one is not permitted to ϕ at t. The conclusion of this argument is such a sentence.

While I've used ER as an example, the phenomenon extends to a number of other epistemic requirements. KR, for instance will also generate strong incoherence alone. Inevitably, there will be a number of things that one is in a position to know at any given time, but which are not possible to come to know at the same time. Any dynamic requirement norm whose triggering condition can be met with respect to several different propositions at the same time will require that we come to believe a proposition while, by Incompatible Alternatives, not permitting that we come to believe that proposition. Even epistemic requirements that are less demanding than ER and KR will generate strong incoherence:

Resolve Inconsistency: If one has inconsistent beliefs at t , then one ought to cease to believe one of them at t .

It's perfectly possible (indeed, likely) that at any given time, we have a number of inconsistencies among the things we believe. Figuring out how to resolve these inconsistencies often takes time, so it's plausible that we're not often able to resolve more than one inconsistency at the same time. Resolve Inconsistency, then will require that we resolve a given inconsistency while forbidding, by way of Incompatible Alternatives, that we resolve that inconsistency. I'll discuss which norms face this puzzle in 3.3 and 3.4, but it suffices for now to say that a wide range of dynamic requirements will produce this issue.

This argument presents a puzzle for those who would evict the zetetic from the epistemic. If our rationale for exiling inquiry, an activity central to our epistemic lives, from the epistemic was that it would eliminate incoherence from the epistemic domain, it should give us pause that we have not successfully done so. To restore coherence to the epistemic, we must reject one of the premises.

3: Responses to the Puzzle

In order to resolve the puzzle presented in the previous section, the separationist has three main options: 1) reject the case, 2) reject Incompatible Alternatives, or 3) reject ER (and other strongly incoherent norms) as formulated. I discuss rejecting the case in 3.1 and rejecting Incompatible Alternatives in 3.2. The third approach can take a few forms. First, in 3.3, I discuss rejecting ER and the other incoherent norms wholesale. Second, I discuss replacing ER and other incoherent norms with more modest norms that avoid incoherence in 3.4. Finally, I consider in 3.5 a way to reformulate ER and affected norms in a way that captures the motivation of the norm without running into incoherence.

3.1 REJECT THE CASE

The case, as described, is that one has excellent evidence for two propositions, P and Q, but one cannot come to believe both of them at the same time. This may be because each would require some reflection and deliberation to arrive at, though no particular explanation is essential to the case. Someone who would deny that our case is genuinely possible, then, would have to claim that we can come to believe any two propositions for which we have excellent evidence at the same time. This claim is not particularly plausible. This is not to deny that there are a number of propositions for which we have excellent evidence that we *can* come to believe at the same time. A belief that there are exactly 12 cupcakes in the box and a belief that there aren't 13 cupcakes in the box, for instance, can be formed simultaneously, and typically would be. However, there certainly are some excellently supported propositions which one would only come to believe after a minute or two of carefully attending to one's evidence and which are dissimilar enough that one could not attend to the evidence for each of them at the same time. For example, suppose I currently have excellent evidence that I have taken more than 30 plane rides in my life. To come to believe that proposition, I would need a few moments to take stock of the various flights I've taken at various points in my life. I also have excellent evidence that the car my grandfather had when I was a toddler wasn't a Prius. However, it would take me a moment or two to work out which years he would have had the car and recall the Prius was not released until the late-90s, and thereby form the belief. I have excellent evidence for each of these, but I cannot form them at the same time because I cannot review the evidence for both at the same time. This seems like a perfectly plausible case. However, to avoid strong incoherence within the epistemic by denying the case, one would have to deny that such a case is possible. This, to me, is a

high cost and I'll argue in the next two sections that rejecting Incompatible Alternatives or ER are more promising options.

3.2 REJECT INCOMPATIBLE ALTERNATIVES

The next option we might take to avoid strong incoherence is to reject Incompatible Alternatives for belief formation.³⁶ Denying Incompatible Alternatives for belief formation amounts to saying that if you ought to ϕ at t , and you cannot both ϕ and ψ at t , then ψ -ing may well be permissible at t (where ϕ -ing and ψ -ing are forming some belief or other). If ER is correct, then we will have a vast number of propositions that we ought to come to believe, but which we cannot come to believe at the same time. Indeed, we cannot come to believe all of the propositions we have excellent evidence for at all, let alone at the same time. If we reject Incompatible Alternatives, it will not be impermissible to come to believe any of these propositions (at least not for the reason that doing so would be incompatible with believing the others).

One reason some epistemologists might favor rejecting Incompatible Alternatives for belief formation is that it handles requirements to believe that go beyond an agent's current abilities rather awkwardly. So long as there is one proposition that one ought to come to believe, Incompatible Alternatives will deliver the result that one is not permitted to believe any proposition which one cannot believe. Call the proposition that one ought to believe P and an arbitrary proposition which one cannot believe Q . One cannot believe P and Q at the same time because one cannot believe Q at all. So, by Incompatible Alternatives, one is not permitted to believe Q .

While it may seem harmless to forbid believing propositions that an agent cannot in fact believe, this feature of Incompatible Alternatives is an uneasy fit with the way

³⁶ Note that this is perfectly consistent with maintaining a version of the principle where actions are concerned.

many epistemologists think of the norms of rationality. For instance, some think that rationality has something to say about which propositions are deserving of belief, even beyond the capabilities of human agents.³⁷ On this view, it is extremely odd to say that those beliefs are not just out of reach, but rationally impermissible, for human agents. It is one thing to say that rationality does not require more than we are capable of. It's quite another to say that believing anything we're in fact unable to believe is rationally impermissible!

3.3 REJECT ER OUTRIGHT

The norms that produce strong incoherence when paired with the Case and Incompatible Alternatives are dynamic requirements (though not all dynamic requirements are incoherent). The Case is one in which one cannot form a belief that P and a belief that Q at the same time (where one has excellent evidence for P and Q). The static analogue of the Case would be that one cannot both have a belief that P and have a belief that Q. This is wildly implausible. Whatever difficulty we may have in forming two beliefs about different subjects at the same time, it does not extend to merely having those beliefs.³⁸ Similarly, permission norms do not generate incoherence on their own, as ER does. This is because strongly incoherent norms declare both that one ought to ϕ and that it is not permitted to ϕ , while weakly incoherent norms say that one is permitted and not permitted to ϕ . Permission norms do not say what one *ought* to do, nor do they say what is not permissible.

³⁷ Declan Smithies (2015) and David Christensen (2004) argue for this claim.

³⁸ This is not to say that a set of static norms cannot be strongly or weakly incoherent. What I mean here is that static norms will not, as ER does, induce incoherence *on their own* without any other norm. If a static norm is implicated in incoherence (on its own), it will be by way of a bridge principle that generates dynamic requirements from static requirements.

The observation that it's only dynamic requirements that cause trouble suggests several principled ways to do away with ER and other incoherent norms. The first is to endorse a version time-slice rationality and deny that there are dynamic norms in the epistemic domain.³⁹ This route has a further advantage. Friedman suggests that some static norms will be in tension with ZIP because they generate dynamic requirements which conflict with ZIP. Because time-slicers deny that static norms generate any kind of dynamic norm, their norms will not be in tension with ZIP. Advocates of time-slice rationality can admit zetetic norms as genuinely epistemic without rendering the epistemic incoherent, though I expect such a proposal is not to the taste of many of them.

Denying that there are epistemic requirements is another principled way to reject ER and other incoherent norms.⁴⁰ Permissions alone will not generate strong or weak incoherence. On this view, however, ZIP will still generate weak incoherence with the other epistemic norms, so we would have to either accept weak incoherence or banish the zetetic from the epistemic domain.

Time-slice rationality and denying epistemic requirements both have some purchase among contemporary epistemologists (though perhaps the former more than the latter). However, for many, adopting one of these proposals would be just the sort of radical revision of the epistemic that we were hoping to avoid by cleaving the zetetic from the epistemic. An approach that retains at least some of the spirit of incoherent norms, without reintroducing incoherence, is worth considering.

3.4 RESTRICT ER

³⁹ Brian Hedden (2015) articulates time-slice rationality as (in part) the claim that there are only static norms that govern our doxastic attitudes. This leaves open that actions or non-doxastic attitudes like hopes might be subject to dynamic norms, though such norms would not be in the spirit of the proposal.

⁴⁰ Chase Wrenn (2007) and Mark Nelson (2010) both defend something along these lines, although Nelson allows that there may be negative epistemic duties.

Our second option is to reject ER and replace it with a more modest evidential requirement (and doing the same, *mutatis mutandis*, for other incoherent norms). Not just any more modest norm will do. For instance, requiring that one come to believe propositions for which they have excellent evidence *and* which they are interested in would not suffice. I might be interested in and excellent evidence for several propositions which I cannot come to believe at the same time. A more promising approach focuses on the agent's recognition or appreciation of their evidence:

Recognized Dynamic Evidential Requirement (RER): If one has excellent evidence for P and one recognizes that one has excellent evidence that P at t, then one ought to come to believe that P at t.

The thought behind RER is that to generate a requirement to form a belief, one must not only have excellent evidence, but recognize that one does. Why might RER avoid incoherence? Recognizing that one has excellent evidence for a proposition will often require some amount of attention and consideration of one's evidence. One cannot attend to two very disparate subsets of one's evidence at the same time. So, if one can recognize that one has excellent evidence for P and Q at the same time, there must be substantial overlap in one's evidence for P and one's evidence for Q. If P and Q are sufficiently related that one can recognize that one has excellent evidence for P and Q at the same time, then one can come to believe P and Q at the same time. Coming to believe P and coming to believe Q, then, will not be incompatible alternatives, and so incoherence will not arise.

This tack is unlikely to be successful. In the case of RER, one might have excellent evidence for P but nonetheless be unable to form the belief because believing that P would be uncomfortable or distressing. We can easily imagine that in such a case, one needs to turn P over in one's mind and slowly accustom oneself to the idea. In short, even after recognizing that If in such a situation, one recognized that one had excellent evidence for Q, as well, then coming to believe Q would be required, by RER, yet

impermissible by Incompatible Alternatives, because it is incompatible with coming to believe P. So, restricting ER to *recognized* excellent evidence will not help escape incoherence.

There's reason to think that other more modest dynamic evidential norms will not fare better. We might try to remedy the situation by looking for a norm with a tighter connection between its triggering condition and belief formation. That is, we might look for a candidate norm whose triggering condition is so tightly connected to belief that an agent who meets the triggering condition cannot fail to form the relevant belief. It's not clear how a norm that cannot be violated puts any real normative pressure on agents. And certainly, the epistemic norms that are the bread and butter of contemporary epistemology are not so trivially met. On the other hand, if our candidate norm leaves any daylight between its triggering condition and belief formation, then, as with RER, it will be incoherent. All we need is a case in which one meets the triggering condition for a proposition P but fails to believe P, then meets the triggering condition for a proposition Q that is not possible to come to believe at the same time as P. The prospects for a more modest version of ER, DKR, or other incoherent norms are dim.

3.5 REFORMULATE ER

Another potential route is to reformulate ER to avoid inconsistency. This may seem a strange suggestion. I've just argued that even more modest versions of ER will be incoherent, so it seems unlikely that a version of the very same norm could. But ER and its more modest counterparts, as formulated, differ in an important respect from similar dynamic epistemic requirements as they're typically discussed. Namely, epistemologists typically formulate epistemic requirements *timelessly*. That is, the norm itself does not specify a particular time at which one must comply. For instance, consider the timeless version of ER:

Timeless Dynamic Epistemic Requirement (TER): If one has excellent evidence for P, then one ought to come to believe that P.

While TER does not make *explicit* reference to a particular time, there is a sense in which TER requires exactly what ER does. One will come to have excellent evidence for P at a time, which we can call t. TER then, will imply that, at t, one ought to come to believe that P. Is this not exactly what ER says, that one ought to come to believe that P at t? It would seem, then, that there's little difference between the demands of TER and those of its timestamped counterpart.⁴¹

However, this may be too quick. It's true, at t, that one ought to come to believe that P. Does it follow from that fact that one ought to come to believe P *at t*? Or would t+1 do? Suppose that it's perfectly true, right now, that I ought to give part of my income to charity. This leaves open whether I ought to give part of my income to charity *right now*. Giving part of my income to charity tomorrow, or in a few weeks when I balance my budget, would do just as well to meet this obligation, and there may be strong reasons not to do so right now. The obligation to give part of my income to charity appears to be timeless, in that it does not specify a particular time at which I must comply.

We might insist that the timeless obligation is simply a series of timestamped obligations, according to which, for each time t_i , I ought to give part of my income to charity at t_i . However, the difference between the timeless obligation and a collection of successive timestamped obligations is not trivial. When I give part of my income to charity two weeks from now, we want to say that I have fully discharged the obligation- I'm in good standing as far as my obligation to give to charity is concerned. However, if we have a successive series of timestamped obligations, then while I fulfilled my obligation to give to charity at that the time that I did, I have still failed to meet a number of my obligations. It's a serious cost to have to have to say that, after donating part of my

⁴¹ I use "timestamped" to refer to a norm or obligation that does specify exactly when it is to be met. That is, "timestamped" is just the converse of "timeless".

income, I failed to fulfill many of my obligations to donate but I did fulfill one. It's seems clear that there is only the one obligation which is relatively flexible about when and how it is fulfilled.

If dynamic epistemic requirements are timeless, like TER, then they will require that we form certain beliefs, but not specify exactly when we must do so. Such norms will not produce the strong incoherence that their timestamped counterparts do. Take TER as an example. Suppose I, at *t*, have excellent evidence for *P* and for *Q*. By TER, I ought to come to believe that *P* and I ought to come to believe that *Q*, but not necessarily at *t*. These obligations do not trigger Incompatible Alternatives, because it's not the case that I ought to believe *P at t*, rather than a few moments or hours or even days later. To put the point differently, I will not have violated TER if I come to believe that *P* and *then* come to believe that *Q*. There is no normative pressure to come to believe both at one time, as there is with ER. Coming to believe that *P* (at some point) and coming to believe that *Q* (at some point) are not incompatible alternatives at all.

The separationist, then, can simply insist that the dynamic requirements are timeless in order to avoid the puzzle of the previous section. In addition to avoiding incoherence, positing timeless norms has a number of other advantages. It allows us to say that an agent has fully discharged their obligation when they form the requisite belief. With timestamped requirements, on the other hand, we have to maintain in many cases that such an agent has still failed to meet most of the relevant requirements because they didn't form the belief at each time between when they first met the triggering condition and when they did, in fact, form the belief. This strategy for avoiding incoherence is appealing. We need not reject an everyday case, deny the intuitive Incompatible Alternatives, or impose a serious revision on our favored epistemic norms. In fact, given that these norms are most often formulated timelessly, it may be more in line with the norms that philosophers do, in fact, favor.

4: Dissolving the Tension

A tempting response to the incoherence between ZIP and other epistemic norms is to try to separate the epistemic from the zetetic. However, this separationist strategy faces a puzzle: dynamic requirements generate incoherence alone, without taking any other norms into account. Separating the epistemic from the zetetic, then, will be insufficient to restore coherence to the epistemic. If we want to separate the epistemic and the zetetic without accepting incoherence or radically revising the epistemic, we'll have to reject one of the principles that gave rise to the incoherence.

I've surveyed the three basic responses to the puzzle: rejecting the case, denying Incompatible Alternatives, and denying ER. The case is highly plausible, so rejecting it is not a satisfactory solution to the puzzle. Rejecting Incompatible Alternatives for beliefs is somewhat more promising. Rejecting ER and other dynamic requirements outright suffices to avoid the puzzle at the cost of fairly radical revision. Restricted versions of ER fail to avoid the puzzle. A more promising solution is to deny ER as stated and reformulate it timelessly. However, solving the puzzle does not necessarily vindicate the separationist strategy. All of the plausible solutions to the puzzle, I'll argue, also dissolve the tension between ZIP and ER that motivated separationism in the first place.

The tension between ZIP and the dynamic epistemic requirements and permissions relies on Incompatible Alternatives. If we reject Incompatible Alternatives, ZIP will not deliver the consequence that it's impermissible to form the beliefs that ER requires and EP permits. In other words, ZIP will no longer be strongly incoherent with the dynamic requirements or weakly incoherent with the dynamic permissions. However, a milder sort of tension remains. It is still the case that the dynamic requirements and ZIP together put agents in a bind. If one cannot both Φ and ψ , but one is required to do both, then inevitably, one will fail to do as they ought. This, however, is a more familiar sort of tension: the tension between the demands of epistemic norms and our limited capabilities. While some, of course, think this is reason to revise the

epistemic, many, perhaps most, epistemologists are comfortable with the idea that rationality will ask more of us than we can deliver.

Whether rejecting ER dissolves the tension depends on whether we replace it and what we replace it with. A defender of time-slice rationality, who rejects call dynamic norms including ER, will face no incoherence between ZIP and the norms he endorses. However, he may still have reason to be dissatisfied with ZIP. While ZIP does not require a belief change, it does require temporally extended activities for the purpose of belief change. Whether a time-slicer will want to accept ZIP depends on how we think of the commitments of time-slice rationality. Hedden formulates time-slice rationality as a claim about the rational evaluation of attitudes (namely, that it doesn't depend on an agent's attitudes at other times). ZIP doesn't require attitudes at all, and so is not inconsistent with the letter of time-slice rationality. It may, however, be in tension with the spirit. Time-slicers would like our evaluation of an agent at a time to be solely a function of her doxastic attitudes at that time. We can determine whether a time-slice of an agent meets ZIP, but only by turning our attention away from her doxastic attitudes towards the activities she is engaged in. Another way we might reconcile ZIP with time-slice rationality is to say that rationality isn't all there is to the epistemic domain— even if rationality concerns our attitudes at a given time, other epistemic evaluations, like zetetic evaluations, might take more into account.

Instead of going time-slice, suppose we replace ER with a timeless norm, TER. TER does not require that an agent with excellent evidence that P believe P *at t*, but it does require that she believe P. Supposing P is a proposition which requires some deliberation to believe (in the normal way), Incompatible Alternatives and ZIP will say that it is impermissible to form a belief that P *at t*. While there isn't strong incoherence between ZIP and TER as I've defined it, ZIP and TER imply a sentence of the following form:

(T) One ought to ϕ and one is not permitted to ϕ at *t*.

The question, then, is whether this feature, which we can call temporary incoherence, is an unacceptable kind of tension. Suppose I ought to give part of my income to charity—not necessarily *at t* or any particular time. I also ought to keep my promise to pick my friend up from the airport *at t*. I can't drive to the airport and donate money at the same time, so Incompatible Alternatives will tell us, plausibly, that I am not permitted to give money to charity *at t*. Is it plausible that I ought to give money to charity *and* that I am not permitted to give money to charity *at t*? It's clear to me, at least, that temporary incoherence is commonplace and hardly gives us reason to reject either of the temporarily incoherent norms. This is not to say that there's no tension at all, but only that it's not the sort of tension that motivates radical revision of the epistemic domain.

What does all of this mean for the separationist strategy? The two most promising ways to avoid the puzzle for the separationists—rejecting Incompatible Alternatives and moving to a timeless dynamic norm—both leave us with no tension, or an unproblematic sort of tension, between ZIP and familiar dynamic epistemic norms. Any way of solving the puzzle then, strips away the reason to pursue the separationist approach in the first place: to eliminate incoherence within the epistemic domain. Separationism then fails as a way to preserve coherence within the epistemic. If we keep our time-stamped dynamic requirements and Incompatible Alternatives, the epistemic will be incoherent even without ZIP. If we deal with this incoherence by rejecting Incompatible Alternatives or opting for a timeless dynamic norm, then we don't need to separate the epistemic and the zetetic to preserve coherence.

5: Conclusion

Separationism seems, at first blush, to be a promising answer to the *reaction question*: how should we, as theorists, respond to the tension between the epistemic? ZIP is, after all, an instrumental norm that is highly sensitive to agents' desires and goals. Relegating it and its ilk to the practical is perhaps not such a steep cost to pay to eliminate

incoherence from the epistemic domain. However, as I argue in Section 2, cleaving the zetetic from the epistemic won't suffice to eliminate incoherence. This is a puzzle for the separationist—the dynamic requirements that are strongly incoherent with ZIP are also strongly incoherent alone. The most promising ways of addressing this puzzle will also dissolve the tension between ZIP and the dynamic norms, leaving separationism unmotivated.

Incoherence, however, is not the only reason one might insist that ZIP is not epistemic. For some, the instrumental nature of the norm suffices to make it practical, and therefore not epistemic. This involves denying, contra Friedman, that a norm can be both practical and epistemic. Whether or not this route is appealing, the upshot for theorists of the epistemic is interesting. If we are to separate the epistemic from the zetetic, it will be because of our commitment to a particular characterization of the epistemic according to which zetetic norms aren't epistemic, not because a conflict between them forces our hand.

CHAPTER 3: CAN COMPARATIVISM ABOUT CREDENCE DELIVER PROBABILISM?

If I run into a skeptic about precise real-valued credences, one thing I might offer her is that surely, she's more confident in some propositions than others. If she can order propositions according to how confident she is in them, then it's not too great a jump to assigning numbers to them. These comparative beliefs have an edge on precise, real-valued credence when it comes to convincing the skeptic. Unlike precise, real-valued credences, we can reliably learn what our comparative beliefs are by introspection and comparative beliefs are much less fine-grained.

We might write off this talk of comparative belief as just a way of motivating precise credence. But we might also understand it as saying something about what credences *are*. There have been a number of approaches to characterizing credence. Some have aimed to explain credences in terms of attitudes that aren't graded belief: preferences or evaluative attitudes. Others have insisted that credences are psychologically real and help explain our actions, preferences, and choices. I focus in this chapter on a third option: comparativism about credences, the view that comparative belief, not credences, is psychologically real and primitive.

Credences, according to comparativism, are simply a tool for representing comparative belief. In order to position comparativism as a genuine theory of credence, the comparativist has two tasks. First, he must characterize what credence *is* on the comparativist picture and how it arises from comparative beliefs. Second, he must show that probabilism, the dominant theory of rational credence, will still hold on a comparativist notion of credence. This task is somewhat tricky because comparativists typically see the norms on comparative belief as normatively fundamental and norms on credence as deriving from them. The comparativist must show that probabilism falls out of constraints on rational comparative belief and the comparativist notion of credence.

In Section 1, I explain comparativism and the motivations behind it. In Section 2, I present the comparativist argument for probabilism. In Section 3, I argue that the argument fails because it relies on the claim that a rational agent will satisfy the Totality condition, which requires that they have a comparative belief about every pair of propositions. In Section 4, I explore the prospects for a comparativist argument for probabilism that does not rely on Totality. I argue that the weaker version of the argument fails as well. In Section 5, I canvas some possible responses to the failure of the comparativist argument for probabilism.

1: Comparativism

Comparativism is the view that comparative belief is (i) psychologically real and (ii) a primitive notion in decision theory and epistemology.⁴² Real-valued credences or other numerical degrees of belief, on this view, are merely “a useful fiction”(Stefánsson 2018, p.1). This descriptive claim goes along with the normative claim that norms on comparative belief, rather than norms on credence, are normatively fundamental.

To flesh out the comparativist position, it’s useful to contrast it with its competitors. Stefánsson sees comparativism as an alternative to two broad camps in interpretations of classical decision theory and the notions of credence and utility which it utilizes: behaviorism and mentalism.⁴³ On a behaviorist interpretation, an agent’s credence and utility derive from her preferences. Preferences, in turn, are typically reduced to behaviors or behavioral dispositions. Mentalism, by contrast, takes credences

⁴² The version of comparativism that I’ll discuss here, sometimes referred to as “primitivist comparativism” is defended by H. Orri Stefánsson (2017, 2018). Lyle Zynda (2000) offers a somewhat similar approach, which Stefánsson dubs “pragmatic comparativism.” However, Zynda ultimately sees comparative beliefs not as a primitive doxastic attitude, but a tool for interpreting preferences.

⁴³ See Samir Okasha, “ON THE INTERPRETATION OF DECISION THEORY,” *Economics & Philosophy* 32, no. 3 (November 2016): 409–33, <https://doi.org/10.1017/S0266267115000346>; Lara Buchak, *Risk and Rationality* (Oxford University Press, 2013). for discussion of these broad classes of interpretations of decision theory.

(and utilities) to be psychologically real mental entities. Stefánsson sees comparativism as a third option.⁴⁴ Comparativism rejects the mentalist claim that credences are psychologically real. Unlike the behaviorist, however, comparativists do not reduce credence to conative attitudes or outward behavior.

Stefánsson argues that comparative belief should be taken to be primitive because it is indispensable in explaining intentional action. In order to explain why Karl went to the bar, to borrow Stefánsson's example, it's not enough to know his degree of belief that the bar is a good place to make a friend.⁴⁵ Rather, we need to know that he thinks he's more likely to make a friend at the bar than at the movie theater or at home as he watches Netflix. Indeed, this suggests that comparative belief is *necessary* for explaining intentional action and that numerical degrees of belief can do so only because they encode comparative beliefs.⁴⁶ The explanatory role they play in explaining agent's behavior Stefánsson argues, warrants taking comparative beliefs as a primitive in our theory of intentional action. We have reason to take the primitives of our best theories of intentional action to be psychologically real.

Because it is a primitive, the comparativist will not offer a definition or analysis of comparative belief in terms of preference, desire, full belief, or other notions. However, we can characterize comparative belief by pointing to how it will typically, but not unfailingly, interact with other attitudes. For instance, comparative belief has a role in explaining preferences. If I am more confident in A than B (that is, I have the relevant

⁴⁴ While I focus on comparativism about credences here, Stefánsson also favors comparativism about utilities, according to which comparative desires, rather than numerical utilities, that are psychologically real and fundamental.

⁴⁵ We can suppose, for simplicity, that Karl's sole desire was to make a friend.

⁴⁶ Stefánsson argues further that comparative belief (along with comparative desire) is both necessary and sufficient for explaining and rationalizing intentional action, while numerical degree of belief (and utility) is neither necessary nor sufficient. However, the full argument is not necessary to spell out the view. See Bradley and Stefánsson (2016) for an defense of the view that comparative desire (rather than should be taken as a primitive notion in our theories of intentional action.

comparative belief), then I'll typically prefer a bet that pays \$1 if A over one that pays \$1 if B. Similarly, should I learn that in fact A is false and B is true, my comparative belief that A is more likely than B will explain why I feel surprised (and why my surprise is appropriate).

It's worth clarifying some of the language that I will use to discuss comparative belief. I say that someone 'compares P and Q' when she has some comparative belief or other about P and Q. I'll attribute particular comparative beliefs by saying that someone 'believes P is more likely than Q', 'thinks P is at least as likely as Q' or 'is more confident in P than in Q', and variation on these locutions. It's important to note, however, that this does not imply that that person has a *full belief* about the content 'P is more likely than Q'.⁴⁷

To my knowledge, no comparativist explicitly discusses the relationship between comparative belief and full belief. It seems to me, however, that the comparativist has three basic options. First, he might reduce full belief to comparative belief. In this case, we would expect full beliefs about relative likelihood and comparative beliefs to march in lockstep. Second, he might eliminate full belief entirely. Finally, he could admit two, distinct (though no doubt connected) notions: comparative and full belief.⁴⁸

⁴⁷ Comparativists like Stefánsson and their ilk are typically silent on full belief. This is some indication that, like many of those who focus on credence over full belief, they see full belief as something to be reduced to their preferred primitive notion or eliminated entirely. However, comparativism as stated is neutral on whether full belief should be eliminated, reduced to comparative belief, or considered to be a wholly distinct doxastic state. The one possibility that *is* ruled out by comparativism is that comparative belief would be reduced to a full belief, say, that A is more likely than B. Konek (2019) considers how comparative belief might interact with full belief, although unfortunately he does not discuss comparativism, the view that comparative belief is a primitive doxastic attitude.

⁴⁸ This option may have similar motivations to belief-credence dualism, the view that there are full beliefs and credences, and neither reduces to the other. See Jackson (2020) for an excellent overview of the possible views about how full belief and credence relate to one another.

In order to better characterize comparative belief, it's worth engaging in some speculation about how the comparativist might distinguish full beliefs about relative likelihood and comparative beliefs, should he choose the final option. Notice that when Karl has a comparative belief, he simply has a single attitude towards a pair of propositions: that he will find a friend if he goes to the bar and that he will find a friend if he goes to the movies. However, if Karl also has a full belief that he is more likely to find a friend in the bar than at the movies, then he has an attitude (full belief) towards a content that includes explicit reference to likelihood. Depending on which notion of likelihood appears in Karl's full belief, he may believe one of a number of things. Suppose we impose a frequentist interpretation, so that what Karl believes is something like the claim that in the relevant class of similar cases, he's found a friend more often at the bar than at the movies. Such a full belief may provide some support for a comparative belief that he's more likely to find a friend at the bar than at the movies, but it seems clear they are not the same state. We could imagine Karl having the same full belief about relative frequencies in the relevant class of cases, but being struck with a bout optimism about finding a friend at the movies, and so having a (perhaps unreasonable) comparative belief that he's more likely to find a friend at the movies than the bar. Similarly, suppose Karl's full belief is about the objective chance that he'll find a friend at the bar or the movies. We can imagine him sincerely claiming ignorance of the objective chance of his finding a friend each place or even saying that, for all he knows, the objective chance of his finding a friend at the movie is higher after all, while still having a comparative belief that he's more likely to find a friend at the bar. A full belief about the relative likelihood of two propositions, where likelihood is interpreted as frequency or objective chance, can come apart from the corresponding comparative belief.

However, it's clear that the attitude of comparative belief is intimately connected to probability, in particular to subjective probability. Suppose that it's a subjective interpretation of "likely" at play in Karl's full belief that he's more likely to find a friend at the bar than the movies. Typically, we would cash this out by reference to Karl's

degrees of belief that he'll find a friend at the bar or that he'll find a friend at the movies. But on a comparativist framework, it's comparative belief that underwrites degrees of belief. So, our hypothetical comparativist might say that Karl's full belief will be true just in case he has the comparative belief that he's more likely to find a friend at the bar than the movies. On this picture, what it is for P to be subjectively more likely than Q (relative to a given agent) is for the agent to have a comparative belief that P is more likely than Q.

1.1: MOTIVATIONS FOR COMPARATIVISM

Stefánsson sees comparativism as a way to get the theoretical benefits of credences without making arbitrary or otherwise implausible claims about agent's mental states:

The general motivation underlying Comparativism about belief is the thought that there is no reason to accept realism about any particular type of numerical degree of belief, since, first, comparative belief can play the role we want numerical degree of belief to play, and, second, no possible evidence could justify attributing one type of numerical degree of belief rather than another to an agent. (2018, p.2).

The two key claims here are that we have reason to be dissatisfied with credences (and other numerical degrees of belief) and that comparativism is a genuine alternative that can play the theoretical role in epistemology and decision theory that credences typically do. I'll discuss each of these claims in turn.

Zynda (2000, p.51-54) argues that probability functions are but one of a number of non-equivalent ways of numerically representing degrees of beliefs. For instance, believability functions map each proposition to a real number between 1 and 10, which is its believability. Believability differs from probability in more than just its scaling convention. Unlike probability functions, a believability function is not additive. Instead,

it is sub-additive: when $[X \cap Y = \emptyset, B(X \cup Y) < B(X) + B(Y)]$, where $B(X)$ is the believability of X . The upshot is that there are at least two coherent but incompatible numerical representations of degrees of belief.⁴⁹

The question, then, is why we should attribute credences, rather than believabilities, to agents. Stefánsson suggests there is no fact about the agent that would justify saying that her degrees of belief are *really* on a scale from 0 to 1 rather than 1 to 10. Instead, all we can say is that she her degree of belief is .5 relative to a 0-to-1 scale but 5.5 relative to a 1-10 scale. His conclusion is that numerical degrees of belief, be they credences or believabilities, are not psychologically real. The sense in which they *are* real is that they encode information about something that *is* psychologically real. The features that are preserved between believability and credence representations of an agent's degrees of belief are the comparative facts.

Beyond worries about the incompatible systems for numerically representing degrees of belief, taking comparative belief, rather than precise, real-valued credence, to be psychologically real can assuage many of the worries that cause some to balk at credence talk. For one, we can discover quite a lot about our comparative beliefs by introspection. By contrast, we cannot typically learn what our precise, real-valued credences are by introspection. Of course, introspection is not a perfect guide to comparative beliefs or to credence. If we think credence is determined by dispositions to bet, or to evaluate bets as fair, or even to reason in certain ways, then we should expect our exact degree of belief to often be opaque to us. However, the fact that we can almost never discover our precise, real-valued credences by introspection should give us pause.

A related worry is that precise, real-valued credences are unrealistically precise. What fact about me would justify attributing to me a .9987 credence, rather than a .9986

⁴⁹ One might worry that having believabilities rather than credences is incompatible with practical rationality. If we imagine one's preferences as resulting from degrees of belief and utilities by an expected value maximization rule, then an agent with believabilities will have incoherent preferences. However, Zynda argues that an agent with believabilities can still have coherent preferences, but her preferences will arise from her believabilities and utilities by a maximizing the "valuation" of various actions or options.

credence, that prickly pears are native to Texas? Attributing one credence rather than a very similar one, the thought goes, seems to suppose that I have very fine-grained doxastic states.

A second further worry along these lines is that having precise credences goes beyond our evidence. Our evidence is often sparse, vague, and imprecise. Evidence like that doesn't support a particular precise credence, but rather a more imprecise or coarse-grained attitude.⁵⁰ On this view, having precise credences is not just unrealistic, it's rationally subpar.

A final worry, and a particularly salient one in discussion of comparativism, is that a precise credence function (even an improbabilistic one) represents us as being much more committal than we in fact are. If I have a credence in P and one in Q, then according to my credence function, I am more confident that P than that Q, more confident that Q than that P, or equally confident in each. This is an artifact of numerical representation—numbers are all comparable. However, it seems an ordinary human agent will not be so opinionated. For instance, I am very confident that escarpment live oaks are native to Texas and that Cora is a good dog. In fact, I fully believe each. But I cannot tell you which I am more confident, or whether I'm equally confident in both.⁵¹ Insofar as normal human agents lack such comparative attitudes, then credences represent them as having more opinions than they do.

There's at least some reason, then, to be dissatisfied with taking precise, real-valued credences to be psychologically real. But if there are no genuine alternatives that can play the role that credence does in our theories, then perhaps we need to accept or explain away these concerns. Stefánsson, however, claims that comparative belief can do

⁵⁰ See Joyce (2005) and Sturgeon (2008).

⁵¹ Of course, my failure to introspect a comparative judgment does not settle the question about whether I in fact compare these two propositions. Still, it seems plausible that in at least some cases, I cannot find a comparative belief about P and Q by introspection because I genuinely don't compare P and Q.

everything that credences do. That is, it can explain intentional action and it can play the same (or a similar) role in probabilism and decision theory.

1.2: CREDENCES, COMPARATIVISM, AND PROBABILISM

One might wonder why the comparativist is interested in a numerical representation of comparative belief at all. After all, we can perfectly well discuss an agent's comparative beliefs and which comparative beliefs they ought to have without resorting to credence talk at all. We've seen in the case of Karl's trip to the bar that comparative belief can explain intentional action without reference to numerical degrees of belief. However, if comparativism is to be a genuine alternative, the comparativist must show that comparative belief can play the role that credence does in classical decision theory and probabilism.

Of course, the comparativist could claim that we ought to abandon any decision theory or epistemic norms that rely on the mistaken premise that agents have credences. However, classical decision theory has been enormously productive and there are a host of arguments for probabilism. With this in mind, the comparativist had better leave probabilism and decision theory intact or, failing that, explain away their appeal.

Stefánsson's answer to the question of whether comparative belief can do everything that credence does in decision theory and probabilism is yes. The reason is that comparativism can deliver credences. Credences, in the comparativist sense, are merely a tool for representing comparative beliefs. However, if credences can usefully and accurately represent comparative beliefs, then the comparativist can bring decision theory and probabilism on board.

The remainder of this paper will ask whether comparativism can vindicate probabilism. The comparativist strategy makes use of a representation theorem which shows that comparative beliefs that meet certain conditions will be representable by a probability function. If those conditions are *constraints on rational* comparative beliefs,

then any rational agent's comparative beliefs will be representable by a probability function. That is, probabilism will be true.

2: The Comparativist Argument for Probabilism

The comparativist argument for probabilism uses a representation theorem that shows that if one's comparative beliefs meet certain conditions, then they can be represented as a probability function. In 2.1, I briefly overview the result that appears in the comparativist argument for probabilism. In 2.2, I give the argument and contrast it with similar arguments for probabilism based on the decision-theoretic representation theorem. In 2.3, I motivate crucial premise of the argument by articulating what is required for credence to play the theoretical role that the comparativist sets out for it.

2.1: COMPARATIVE BELIEF REPRESENTATION THEOREMS

Scott (1964) proves that a set of comparative beliefs that meets certain conditions will be probabilistically representable. Call \mathcal{S} the set of possible worlds (or basic possibilities), and \mathcal{Q} a Boolean algebra on \mathcal{S} . The elements of \mathcal{Q} will be subsets of \mathcal{S} , which I'll refer to as propositions. Call \succsim the comparative belief relation, which holds between two propositions. A comparative belief relation \succsim over \mathcal{Q} is *fully probabilistically representable* by some probability function p just in case:

$$X \succsim Y \Leftrightarrow p(X) \geq p(Y)$$

The most straightforward presentation of Scott's theorem relies on the notion of a balanced pair of sequences of propositions. A pair of sequences of propositions $\langle X_1, \dots, X_k \rangle$ and $\langle Y_1, \dots, Y_k \rangle$ is *balanced* just in case, at each possible world $s \in \mathcal{S}$, the

cardinality of $\{i|s \in X_i\}$ is equal to the cardinality of $\{i|s \in Y_i\}$. In less formal terms, a balanced pair of sequences of propositions will contain the same number of truths at each possible world. We can think of $\{i|s \in X_i\}$ as a way of counting the truths in the sequence $\langle X_1, \dots, X_k \rangle$ at a given possible world s . To make this more concrete, consider the following model and assignments.

$$\begin{aligned} S &= \{s_1, s_2, s_3\} \\ X &= \langle \{s_1\}, \{s_2, s_3\} \rangle \end{aligned}$$

1 will be an element of $\{i|s_1 \in X_i\}$, because X_1 is $\{s_1\}$ which is true at s_1 (and therefore has $\{s_1\}$ as an element). 2 will not be an element of $\{i|s_1 \in X_i\}$ as well, because X_2 is $\{s_2, s_3\}$, which is not true at s_1 . $\{i|s_1 \in X_i\}$ will be $\{1\}$ and have a cardinality of 1. In the same way, we can find that the cardinality of $\{i|s_2 \in X_i\}$, which counts the truths in X at s_2 , is 1 and the cardinality of $\{i|s_3 \in X_i\}$ is 1. In other words, at each possible world, exactly one proposition from X is true. Consider the following two sequences of propositions:

$$\begin{aligned} Y &= \langle \{s_2\}, \{s_1, s_3\} \rangle \\ Z &= \langle \{s_1\}, \{s_1, s_3\} \rangle \end{aligned}$$

X and Y are balanced because the cardinality of $\{i|s \in X_i\}$ is equal to the cardinality of $\{i|s \in Y_i\}$ at each possible world in S . X and Z , however, are not balanced. The cardinality of $\{i|s_1 \in X_i\}$ is 1 while the cardinality of $\{i|s_1 \in Z_i\}$ is 2. That is, the number of truths in Y and Z differs at s_1 . With the notion of a balanced pair of sequences of propositions in hand, we can turn to Scott's representation theorem. Scott (1964) proves the following theorem:

A comparative belief relation \succsim over Ω is fully probabilistically representable by some probability function p if, and only if, \succsim satisfies the following axioms:

1. Non-Triviality: $\Omega \succcurlyeq \emptyset$
2. Non-Negativity: $X \succcurlyeq \emptyset$
3. Totality: $X \succcurlyeq Y \text{ or } Y \succcurlyeq X$
4. Isovalence: For all balanced pairs of sequences of propositions $\langle X_1, \dots, X_k \rangle$ and $\langle Y_1, \dots, Y_k \rangle$, if $X_i \succcurlyeq Y_i$ for all $i \leq k$, then $Y_i \succcurlyeq X_i$ for all $i \leq k$, as well.⁵²

Non-triviality says that the entire state space is at least as likely as the contradiction (that is, the empty set). Non-negativity says that every proposition is at least as likely as the contradiction. Totality says that any two propositions X and Y , either X is at least as likely as Y or Y is at least as likely as X (or both). Isovalence says that for any two balanced sequences of propositions, if each X_i is at least as likely as each Y_i , then X_i is at least as likely as Y_i . That is, X_i and Y_i are equally likely. The thought is that if the two sequences of propositions contain the same number of truths, no matter which world is the actual world, and each X_i is at least as likely as each Y_i , then each X_i and Y_i should be equally likely.

Any agent whose comparative belief relation that meets Non-Negativity, Non-Triviality, Totality, and Isovalence will be fully probabilistically representable by some unique probability function.⁵³ However, this alone does not get us probabilism.

⁵²This presentation is largely adapted from Konek (2019), but I recast Isovalence in terms of balanced sequences of propositions, a notion I borrow from Harrison-Trainor, Holliday, and Icard (2016). I suppress one formal detail here; it's the comparative belief structure, a triple of S , Ω , and the comparative belief relation \succcurlyeq that is probabilistically representable. However, we're holding S and Ω fixed for present purposes, so this won't make a difference.

⁵³ Strictly, Scott's theorem shows that a comparative belief relation will be unique up to a positive affine transformation.

2.2: THE ARGUMENT FOR PROBABILISM

There are two gaps we must fill to leverage the representation theorem for an argument for probabilism. First, probabilism requires that a rational agent's credences be probabilistic. The representation theorem shows that agents that meet the axioms will be probabilistically representable. So, we need to claim that any rational agent will satisfy the axioms.

Second, there is a gap between an agent's being *representable* by some probability function and her actually having probabilistic credences. What we need then, is a premise connecting probabilistic representability with the agent's actual credences.

1. Any rational agent's comparative belief relation \geq meets Non-Negativity, Non-Triviality, Totality, and Isovalence.
 2. Representation Theorem: If the agent's comparative belief relation meets Non-Negativity, Non-Triviality, Totality, and Isovalence, then it is fully representable by some probability function p .
 3. If an agent's comparative beliefs relation is fully representable by some probability function p , then p is her credence function.
- C. Probabilism: Any rational agent has probabilistic credences.

Premise 1 says that each axiom of the representation theorem is a constraint on rational comparative belief. Premise 3 says, in effect, that the probability function picked out by the representation theorem is actually the agent in question's credence function.

The gap between representability and reality is much discussed the literature on the decision-theoretic representation theorem.⁵⁴ The decision-theoretic representation theorem shows that if an agent's preferences meet certain conditions, then her preferences can be represented as resulting from a utility function and probabilistic

⁵⁴ See Christensen (2001), Eriksson and Hájek (2007) and Meacham and Weisberg (2011) for discussion of the analogous issue concerning the decision-theoretic representation theorem.

credences. The question, then, is whether the probability and utility function that the decision-theoretic representation theorem picks out agrees with the independent facts about the agent's credences and desires. That is, the question is whether the representation is *accurate*.

By contrast, the comparativist's task is easy. According to comparativism, credences are useful fiction—a tool for representing comparative beliefs. There are no comparative-belief-independent facts about an agent's credences that a probabilistic representation could fail to capture. Whether a probability function (or an improbabilistic credence function) is an agent's credence function is a matter of whether it represents her comparative beliefs sufficiently well. Whether premise 3 is true, then, depends on whether full probabilistic representation is enough to play the role that the comparativist has carved out for credence.

2.3: REPRESENTING COMPARATIVE BELIEFS

Representations are adequate only with respect to a given goal. Consider a topographical map and a road map. For a postal worker delivering mail in a new section of the city, the road map will give her all the information she needs, while the topographical map will be thoroughly useless. A city planner trying to figure out where to place new storm drains and sewer lines, on the other hand, will find the topographical map indispensable. The road map adequately represents the city with respect to the postal worker's goal of navigating the city, but not with respect to the city planner's goal.

The goal of representing comparative beliefs with credence functions is to evaluate them epistemically and to predict and explain the actions they motivate.⁵⁵ If we

⁵⁵ Epistemic evaluation is slightly too general. We want to evaluate the structural features that comparative belief relations share with probability functions: P is more likely than Q, P is equally likely as Q, or P is less likely than Q, etc. all. We are not trying to evaluate whether the comparative beliefs were responsibly or virtuously formed.

want to know whether the probability function picked out by the representation theorem is a sufficiently *good* representation of the agent's comparative beliefs, then we have to answer two questions. First, what is required to represent some object well with respect to some task? Second, does a probability function that fully represents a set of comparative beliefs represent them well?

Two features seem indispensable to representing something well with respect to a given goal. First, the representation should be thorough, in the sense that it captures all of the relevant sort of feature. The postal worker navigating a new route will find a topographical map or an out-of-date road map inadequate, because both of these maps fail to include the sort of feature that is relevant to her goal: roads, their names, and their location relative to one another. Of course, a representation can be more or less thorough. An out-of-date road map might miss only a handful of small streets or a freeway exit here or there. The topographical map, on the other hand, will miss almost all of the features that she needs to carry out her task.

A second feature that a good representation has is faithfulness, in the sense that it does not represent the object as having features (of the relevant sort) that it does not in fact have. Consider a more fanciful map, perhaps stolen off a city planner's desk, which includes a number of roads that don't currently exist, in addition to all of the roads that do. This, too, will frustrate the postal worker as she tries to use it to navigate. Routes that seemed possible will fail to materialize and leave her scrambling to find her destination.

It's important to clarify here that almost all representations will include *some* features that do not correspond to features of the represented object. This is not a bad thing. For instance, a political map of the world might show Argentina in pink, Brazil in yellow, and Uruguay in purple, although these colors don't correspond to any actual feature of the world. The colors of countries on a world political map are *representational artifacts*—features of the representation that do not correspond to actual features of the represented object. The important thing is that a representation be thorough and faithful with respect to the features that are relevant to the goal at hand. A

map can be perfectly good for the postal worker's purposes if it fails to include storm drains and fire stations or if it includes some non-road representational artifacts, like the colors of various features.

The relevant features for epistemically evaluating comparative beliefs are comparative-belief facts. So, a perfect representation for our purposes will capture *all* of the agent's comparative beliefs and not represent the agent as having comparative beliefs that she does not have.

The question we began with was whether the probability function picked out by the representation theorem represents the relevant agent's comparative beliefs sufficiently well to count as her credences. Recall that an agent who meets the axioms will have fully probabilistically representable comparative beliefs in the following sense:

$$\boxed{X \succcurlyeq Y \Leftrightarrow p(X) \geq p(Y)}$$

We can see that a probability function that agrees with an agent's comparative beliefs in this sense will be thorough—whenever $\boxed{X \succcurlyeq Y}$, the probability function reflects that fact. Similarly, it's clear that the probability function that the representation theorem picks out will be faithful. Any comparative fact that appears in the probability function will also appear in the agent's comparative beliefs. Strong representability, then, suffices for good representation of comparative beliefs with the goal of epistemic evaluation. A credence function that strongly represents a set of comparative beliefs can play the theoretical role that the comparativist sets out for credence. This means that Premise 3 of the comparativist argument for probabilism is true. However, I argue in the next section, that the argument is unsound because Premise 1 is false.

3: The Totality Condition

Premise 1 of the comparativist argument for probabilism says any rational agent will meet the axioms of the representation theorem— Non-Negativity, Non-Triviality, Totality, and Isovalence. I argue in this section that a rational agent need not meet Totality. The Totality axiom says that the agent compares each pair of propositions, in the sense that she has a comparative belief about them. If Totality is a rational requirement, then any rational agent will compare every pair of propositions. In this section, I canvas two arguments in favor of thinking that rational agents will satisfy Totality and argue that they are not persuasive. Then, I present reasons to think that Totality is not a rational requirement, and therefore that comparativist argument for probabilism presented in the previous section fails.

3.1: ARGUMENTS FOR TOTALITY

Cian Dorr, Jake Nebel, and Jake Zuehl (ms.) argue for a comparability principle for gradable adjectives like “tall,” “probable,” or “valuable”.⁵⁶ They further claim that this comparability principle, with a few plausible assumptions, entails the striking conclusion that all agents have precise credences in every propositions which they are confident to any degree (forthcoming).⁵⁷ It stands to reason, then, that the comparability principle may provide some support for the claim that rational agents will satisfy Totality.

⁵⁶ They also claim that comparability holds for certain adverbs and gradable adjectives that take an internal argument like “fond (of x)”.

⁵⁷ It might seem that this conclusion should suffice to motivate the claim that rational agents will meet Totality. If agents have precise credences, and these credences fully represent their comparative beliefs, then they must have complete comparative beliefs as well. However, there are two obstacles to this line of reasoning. First, it’s possible that there are propositions that an agent is not confident in to *any* degree. (That is, they have no degree of confidence about the proposition, not that they have an extremely low degree of confidence.) So, their claim does not quite suffice to get us complete comparative beliefs. Second, as I’ll discuss in the next section, there are several notions of credence that the comparativist can draw on that fall short of *full* representability. According to some of these weaker notions, an agent might have complete credences, but nonetheless have incomplete comparative beliefs.

Although Dorr, Nebel, and Zuehl do not discuss comparative belief directly, I'll attempt to fill in an argument from their comparability principle to the claim that rational agents will satisfy Totality. The principle they discuss is:

Comparability: If x is at least as F as x and y is at least as F as y , then either x is at least as F as y or y is at least as F as x .

The antecedent of this principle may look rather odd. The idea is to address values for x and y to which the graded adjective in question does not apply to any degree. So, if we think that sentences like “magenta is as tall as Tuesday” will be false, then so too will “magenta is as tall as magenta”.⁵⁸ Comparability will be trivially true in these cases. For our application, we can assume that x and y will be propositions—subsets of S —and that any rational agent will believe $\boxed{X \geq Y}$. With these assumptions we can work with a simplified version of Comparability:

Simplified Comparability: Either x is at least as F as y or y is at least as F as x .

While there may be concerns about whether Comparability is true, I'll grant both Comparability and Simplified Comparability (with our simplifying assumptions). Instead, I'll argue that they're insufficient to show that any rational agent will meet Totality. Recall that an agent meets Totality if they have a comparative belief that $\boxed{X \geq Y}$ or that $\boxed{Y \geq X}$ for each X and Y in S . Comparative belief itself is an attitude, not a gradable adjective. However, it is closely connected with subjective likelihood and “likely” is a gradable adjective. The strategy for arguing from Simplified Comparability to the claim that rational agents meet Totality, then, is to leverage comparability for “likely”.

⁵⁸ Dorr et al. do not claim that such sentences will be false. Nonetheless, they want to claim only that the consequent of Comparability will be true of things that are F to some degree.

Subjective Likelihood Comparability: Either X is at least as likely as Y or Y is at least as likely as X.

If Subjective Likelihood Comparability is right, and an agent has a comparative belief that P at least as likely as Q just in case P is at least as subjectively likely as Q, then every agent will have complete comparative beliefs. This is a strong conclusion. The claim here is *not* that any rational agent will have complete comparative beliefs. Rather, the conclusion is that all agents, even the wildly irrational, do *in fact* have complete comparative beliefs.

While Subjective Likelihood Comparability looks like an instance of Comparability, understood correctly, it is not. I'll argue that we can accept Comparability and still reject Subjective Likelihood Comparability.

"Likely", on its face, looks like any other gradable adjective, like "tall" or "heavy". But when we are considering *subjective likelihood*, there is an implicit reference to an agent. A proposition is subjectively likely *to* a given agent. What we mean when we say that P is at least as subjectively likely than Q is that the agent in question finds P to be at least as subjectively likely as Q. Subjective Likelihood Comparability, then, is elliptical for something like the following:

Either A finds X at least as likely as Y or A finds Y at least as likely as X.

This is *not* a consequence of Simplified Comparability. To see this, consider an easy and plausible instance of Simplified Comparability: either X is at least as tall as Y or Y is at least as tall as X (assuming X and Y are at least as tall as themselves). This seems undeniable. It doesn't however, imply that Hannah believes that X is at least as tall as Y or believes that Y is at least as tall as X. Nor does it imply that they *suspect* or *hope* that X is at least as tall as Y, or vice versa. That is, the comparability of tallness doesn't tell us anything at all about Hannah's attitudes about the relative tallness of two individuals.

Finding X to be at least as likely as Y is an attitude that an agent can have. The comparability of gradable adjectives does not, on its own, tell us much about the attitudes that an agent takes towards them. If I'm correct that subjective likelihood claims are really telling us about an agent's attitudes, then Subjective Likelihood Comparability does not follow from Simplified Comparability at all. Comparability, then, does not give us reason to suppose that rational agents will meet Totality.

Another reason to think that comparative beliefs should be complete is that having incomplete comparative beliefs might be practically irrational. For instance, suppose I fail to compare "Cora is a good dog" and "it's raining outside right now". Suppose I am offered a bet that pays out \$1 if Cora is a good dog or a bet that pays out \$1 if it's raining outside. I would have no preference between them or, at most, a preference that cannot be rationalized by my comparative beliefs. Supposing I should only take an option if I believe it has a higher expected value than the alternatives, I should decline to take either bet. This is practically irrational- each bet has positive expected value, so I would be better off choosing either one than declining to choose.

Setting aside concerns about the relationship between practical and epistemic rationality, this argument does not cut much ice. If I have no idea which bet is more likely to pay out and their payoffs are equal, then there is no fact of the matter about which bet has the highest expected value from my perspective. It is clear, however, that each bet has a higher expected value than refusing to bet (assuming I have some confidence in each proposition), which has an expected value of \$0. This suggests that the problem lies with the normative supposition that I should only choose an option that has higher expected value than the alternatives, rather than with my comparative beliefs. It's true that expected value won't settle the question of which I should choose, but it seems perfectly fine to, say, choose arbitrarily, as I might if the options had equal expected value.

3.2: ARGUMENTS AGAINST TOTALITY

Beyond the lack of compelling arguments for the claim that a rational agent will meet Totality, there are a number of further reasons to doubt the claim. A preliminary reason is the bare intuition that there are pairs of propositions which an agent would not be rationally better off for comparing. For instance, I fail to compare “Cora is a good dog” and “it’s raining right now”. I have a great deal of confidence in each of these propositions, but I cannot say which I consider more likely, or whether they are equally likely. Of course, introspection may not be a perfect guide to which comparative beliefs I have or lack. However, it’s plausible that in at least some cases in which I cannot find a comparative belief by introspection, I in fact don’t have a comparative belief. The question, then, is whether I would be better off, epistemically speaking, for comparing “Cora is a good dog” and “it’s raining right now” or “I would enjoy being a florist” and “the first person to enter the post office in Austin, Texas on July 10, 1924 was taller than 5’5””. It does not seem to me that I would.

A more principled reason to doubt that rational agents will have complete comparative beliefs is that many such comparative beliefs would be unsupported by their evidence. In such cases, to form a comparative belief is to make a baseless guess, something rational agents ought not do.

Because comparative belief is a primitive, according to comparativism, it is somewhat tricky to characterize what constitutes sufficient evidence for a comparative belief. A comparative belief is not simply a full belief that X is more likely than Y. Nor, on the comparativist picture, is it the state of having a higher credence in X than in Y. We have well developed theories of what would justify being in either of those states, but it’s not obvious to what degree these extend to comparative belief. I’ll assume going forward that sufficient evidence for a full belief that X is more likely than Y is sufficient evidence for the corresponding comparative belief.

Our question is whether a rational agent will have sufficient evidence for a comparative belief about any two propositions whatsoever. Consider the following case.

Bob tells me that he will probably go to the party tomorrow. Mary tells me that it's pretty likely that she'll go. Their testimony about their plans is the only evidence I have about whether they'll be at the party. If Totality is a rational requirement, then I ought to believe that it's at least as likely that Mary will go to the party as that Bob will, that it's at least as likely that Bob will go as that Mary will, or both, if I think they're equally likely to go. Is either of these supported by my evidence?

I would certainly be justified in saying that it is *roughly* as likely that Bob will go to the party as that Mary will. Being roughly equally likely, however, is consistent with one being (strictly) more likely than the other or with their being equally likely. Something a great deal more precise would be required to support the claim that one was at least as likely as the other and therefore not even a hair less likely. My evidence in this case, however, is too coarse-grained to support such a claim. Natural language likelihood terms, like "probably" and "pretty likely", are notoriously hard to pin down to even rough ranges of probabilities.⁵⁹ It's not unlikely that Bob and Mary themselves didn't have anything much more precise than "probably" and "pretty likely", respectively, in mind. That is, it's plausible that, even on reflection, they would not have been able to produce, say, a precise probability that they would go to the party.

The relatively coarse-grained evidence I have does not support the kind of fine-grained judgement required to say that it's at least as likely that Mary will go to the party as that Bob will (or vice versa). Given my evidence, it is open to me that it's a hair more likely that Bob will, or a hair more likely that Mary will. This does not seem to be a rational failing—in fact, it seems to be just what my evidence would dictate. If this is right, then a rational agent need not have complete comparative beliefs and therefore Premise 1 of the comparativist argument for probabilism is false.

⁵⁹ Track down Mandel & Irwin citation.

4: Representability without Totality

Perhaps unsurprisingly, there has been a push to abandon Totality in favor of representation theorems that show that incomplete comparative beliefs can be represented by probability function in some sense weaker than full representability.⁶⁰ In this section, I consider the weaker forms of probabilistic representability that we get without Totality and the version of the argument for probabilism that results. I argue that this version fails as well, though for a different reason.

4.1: Strong Representability

Totality is necessary for *full* probabilistic representability. That is, it's necessary for representability in the following sense:

$$\boxed{X \succcurlyeq Y \Leftrightarrow p(X) \geq p(Y)}$$

However, there are weaker forms of probabilistic representability. For instance, strong representability:

$$\boxed{X \succcurlyeq Y \Rightarrow p(X) \geq p(Y)}$$

$$X \succ Y \Rightarrow p(X) > p(Y)$$

Strong representability requires that one's credence function preserve the structure of one's comparative beliefs. That is, if one has the comparative belief that $P \succ Q$, then that fact must be reflected in one's credences. But the converse is not required, as it is for full probabilistic representability. Not every comparative fact about

⁶⁰ See Ríos Insua (1992) and Alon and Lehrer (2014). Konek (2019) provides an overview of varieties of representability and the axioms required to guarantee them.

one's credences need represent an actual comparative belief. If an agent has incomplete comparative beliefs, then there will be multiple probability functions that strongly represent her comparative beliefs.⁶¹

While full probabilistic representability requires Totality, strong probabilistic representability does not.⁶² A comparative belief structure is strongly representable just in case it meets the following conditions:

1. Non-triviality: $\Omega \succcurlyeq \emptyset$
2. Non-negativity: $X \succcurlyeq \emptyset$
3. Isovalence: If $X_1 + \dots + X_n = Y_1 + \dots + Y_n$ and $X_i \succcurlyeq Y_i$ for all $i \leq n$, then $Y_i \succcurlyeq X_i$ for all $i \leq n$, as well.⁶³

These are the same conditions, save for Totality, that are required for full probabilistic representation. In order to vindicate probabilism, the comparativist still must argue that Isovalence, Non-triviality, and Non-negativity are conditions on rational comparative belief. But they need not make the implausible claim that Totality is as well.

4.2: A REVISED COMPARATIVIST ARGUMENT FOR PROBABILISM

⁶¹ However, this is not to say that her comparative beliefs is strongly represented only by a set of probability functions. Rather, each individual function strongly represents her comparative beliefs.

⁶² There are other forms of representability that are weaker than strong representability: almost representability and partial representability. The worries I raise about strong representability will extend to these as well.

⁶³ Once again, this formulation of Scott's theorem (1964) is due to Jason Konek (2019).

With this new representation theorem in hand, we can formulate a weaker version of the Comparativist argument for probabilism:

1. Any rational agent's comparative belief relation \geq meets Non-Negativity, Non-Triviality, and Isovalence.
2. Representation Theorem: If the agent's comparative belief relation meets Non-Negativity, Non-Triviality, and Isovalence, then it is fully representable by some probability function p .
3. If an agent's comparative beliefs relation is strongly representable by some probability function p , then p is her credence function.
- D. Probabilism: Any rational agent has probabilistic credences.

This version of the argument does away with the implausible claim that rational agents will have complete comparative beliefs. This version of Premise 1 is a great deal more plausible.⁶⁴ However, strong representability is a weaker standard than full representability. I argued in Section 2.3 that *full* representability ensures that a credence function will represent an agent's comparative beliefs well enough to play the role the comparativist needs it to: a useful tool for representing comparative belief. In other words, Premise 3 of the original argument for probabilism is true.

Strong representability, on the other hand, is too weak to motivate Premise 3 of the revised argument. A probability function that only strongly represents an agent with incomplete comparative beliefs will represent her as having comparative beliefs that she does not in fact have. For instance, imagine two agents, Alice and Burt. Each of their comparative beliefs meet Non-triviality, Non-negativity, and Isovalence. Alice does not compare P and $P \vee Q$, but Burt does. Burt thinks $P \vee Q$ is at least as likely as P . Burt's credences (in the sense of strong representability) will have $p(P \vee Q) \geq p(P)$. Alice's will as

⁶⁴ Some worries may remain about why a rational agent would meet Isovalence. I won't address these here, but refer the reader to Konek (2019, find page) for a discussion of Isovalence.

well, despite the fact that she has no opinion on the matter. This comparative fact about Alice's credences is a mere representational artifact.

One might argue that a representation need not be *perfect* to be useful or to be sufficiently accurate to epistemically evaluate the agent in question. Some features of excellent and useful representations have representational artifacts. However, the failure in Alice's case is not so innocuous. First, the features that the credence function misses might be normatively important features. For instance, Alice fails to compare P and $P \vee Q$. But it's plausible that she ought to, and specifically, she ought to believe as Burt does that $P \vee Q$ is at least as likely as P . One goal of representing an agent's comparative beliefs is epistemic evaluation. If a representation misrepresents the very features we want to evaluate, then it is not adequate for the task.

A similar issue arises when we look at another goal of representing comparative belief: predicting and explaining intentional action. Suppose Burt is offered a choice between \$1 if A or \$1 if B . Burt's credences have $p(A) \geq p(B)$, but this is a mere artifact. In fact, Burt does not compare A and B at all. We would predict based on Burt's credences that he would choose \$1 if A . But this prediction would not be based on any real fact about Burt's doxastic states. If we look at his comparative beliefs, we should be unsure which he would choose or perhaps expect that he would choose arbitrarily.

Further, it's impossible to tell by looking at a credence function which comparative facts are representational artifacts and which genuinely represent a feature of the agent's comparative beliefs. That $p(P \vee Q) \geq p(P)$ in Alice's credences is a mere representational artifact—there is no such comparative belief. But other features of the same sort *do* represent features of Alice's comparative beliefs— for instance $p(\Omega) \geq p(\emptyset)$ represents a comparative belief that Alice actually has. Similarly, the very same feature, $p(P \vee Q) \geq p(P)$, can be an artifact in Alice's credence function, but in Burt's credences, it represents an actual features of Burt's comparative beliefs. On a political map, we can tell representational artifacts from features that genuinely represent some part of the represented object by referring to what *type* of feature it is. Border lines and stars represent borders and capitals, but the particular colors of a country and the particular

location of the country's name over the country are mere artifacts. When it comes to credences that strongly (not fully) represent comparative beliefs, on the other hand, there is no way to know which features are *actual* features of the agent's comparative beliefs and which are artifacts.

The inaccuracy of credence functions that merely strongly represent will, for some agents, be extremely widespread. It is possible that *almost all* of the comparative facts that fall out of an agent's credences are mere artifacts. Consider Cindy, who meets Non-triviality, Non-negativity, and Isovalence. Cindy believes $\Omega \succcurlyeq \emptyset$ and $X \succcurlyeq \emptyset$ for every proposition, and nothing else. She meets Isovalence vacuously—there is no pair of sequences of propositions which contain the same number of truths at each possible world for which Cindy believes $X_i \succcurlyeq Y_i$ for each pair.

Strong representability, then, is not sufficient to ensure that credences can play the theoretical role the comparativist needs them to. Credences, on the comparativist view, are supposed to be a useful fiction— a tool for representing comparative beliefs in a way that allows us to evaluate an agent epistemically. However, strong representability, unlike full representability, allows for credence functions that are partially fabricated and offer no guide to sort the meaningful features from the mere artifacts. Without full representability, then, it seems that the probability function picked out by the (strong) representation theorem won't play the role that credences are meant to play according to comparativists. The revised comparativist argument for probabilism, then, fails because Premise 3 is false. A probability function that only strongly represents an agent's comparative beliefs is such a poor representation that it could not be what the comparativist claims credences are: a useful tool for representing comparative beliefs.

5: Conclusion

If we require full representability, then we must accept the implausible conclusion that any rational agent will satisfy Totality. If we weaken the sort of representability we

require, then credence are a poor representation of comparative beliefs. According to comparativism, we only have credences insofar as they represent our comparative beliefs well. However, for agents with incomplete comparative beliefs, credence functions cannot represent comparative beliefs well.

Where does this discussion leave comparativism? In order to establish comparativism as a genuine alternative to other approaches to credence, the comparativist aimed to show that 1) credence arises out of comparative belief and 2) that probabilism is true on a comparativist framework. The comparativist shows (1) by articulating a theoretical role for credence—a tool for representing comparative belief— and arguing that if a probability function (or set of functions) plays that role sufficiently well, then it counts as an agent's credence function. However, rational agents need not have complete comparative beliefs (to say nothing of irrational agents). Strong representability, which does not require Totality, is not enough to ground the claim that a probability function (or even an improbabilistic credence function for that matter) will represent an agent well enough to count as her credence function.

The comparativist establishes (2) with the comparativist argument for probabilism. I've presented two versions of this argument, one using full representability, the other using strong representability. The first fails because it has as a premise the claim that a rational agent will satisfy Totality. The second fails because its third premise —if an agent's comparative beliefs relation is strongly representable by some probability function p , then p is her credence function — is false.

Where does this discussion leave comparativism? If the conclusions of this paper are correct, then there are two options worth discussing. The first is to develop further notions of representability and representation theorems to go with them. The hope would be to find a notion of representability strong enough to motivate the claim that a probability function *is* an agent's credence function and a set of conditions on comparative beliefs that a rational agent would plausibly meet. The second is to abandon probabilism and focus instead on the norms of comparative belief themselves.

Comparativism is the view that comparative beliefs, not credences, are psychologically real and primitive. Comparativism is typically offered as an alternative view about what credences are, but strictly, it does not entail that any agents, even perfectly rational ones, be well representable by credences or subject to the requirements of probabilism. Nothing I've said here, then, refutes comparativism. This points to one path forward for the comparativist: focus on the norms that govern comparative belief and abandon probabilism. After all, if the norms on credences are derivative and these comparative belief norms are fundamental, there is no reason we could not simply theorize about them directly.

Focusing on norms on comparative belief does not necessarily mean turning our back on the appeal of probabilism. Each axiom of probability has a comparative belief analogue. Indeed, this was de Finetti's aim in defining a *qualitative probability*, though his conditions do include Totality. The Non-Negativity axiom which says that $X \succcurlyeq \emptyset$, for all X , is the analogue of the axiom of the same name which says that $p(X) \geq 0$ for all X . The spirit of Normality could be captured by a principle that said $\Omega \succcurlyeq X$, for all X .⁶⁵ Similarly, de Finetti offers a Qualitative Additivity axiom, an analogue of Quantitative Additivity: : If $X \cap Y = \emptyset = Z \cap Y$, then $X \succcurlyeq Z$ if, and only if, $X \cup Y \succcurlyeq Z \cup Y$.

Turning towards norms on comparative belief opens up questions that do not arise when we focus on norms on numerical credence. While having complete comparative beliefs is not rationally required, and may even be rationally subpar, there does seem to be some comparative beliefs it's irrational not to have. Alice, for instance, ought to believe that $P \vee Q$ is at least as likely as Q . Even for some pairs of propositions that have no logical connection, it seems like that we ought to compare them. For instance, I can see the rain outside my window and I know that it rarely rains in the Atacama desert. It seems that I ought to think that it's more likely that it's raining outside my house than in the Atacama desert.

⁶⁵ De Finetti (1931) offers a different axiom called Normality: : $\emptyset \succcurlyeq X, \emptyset \succcurlyeq S$. The version I offer here is simply an example of a close analogue of the Normality condition on probability functions.

Comparativists look to representation theorems to show that comparativism can vindicate probabilism, despite considering norms on credence to be derivative. However, comparativist argument for probabilism is not as foolproof as it looks. If full representability is the standard, then the comparativist has to make the implausible claim that any rational agent has complete comparative beliefs. On the other hand, if we settle for strong representability, credences turn out to be too comparative beliefs to play the role the comparativist wants them to. This is not necessarily a death knell for comparativism. It does, however, mean that comparativism is not a genuine alternative theory of credence.

WORKS CITED

- Alon, Shiri, and Ehud Lehrer. "Subjective Multi-Prior Probability: A Representation of a Partial Likelihood Relation." *Journal of Economic Theory* 151 (May 1, 2014): 476–92. <https://doi.org/10.1016/j.jet.2014.02.005>.
- Annesley, Kelly. "[Manuscript]," n.d.
- Barnier, Amanda, Lynette Hung, and Martin Conway. "Retrieval-Induced Forgetting of Emotional and Unemotional Autobiographical Memories." *Cognition & Emotion - COGNITION EMOTION* 18 (June 1, 2004): 457–77. <https://doi.org/10.1080/0269993034000392>.
- Berker, Selim. "The Rejection of Epistemic Consequentialism." *Philosophical Issues* 23, no. 1 (2013): 363–87.
- Bradley, Richard, and H. Orii Stefansson. "Desire, Expectation, and Invariance." *Mind* 125, no. 499 (2016): 691–725. <https://doi.org/10.1093/mind/fzv200>.
- Buchak, Lara. *Risk and Rationality*. Oxford University Press, 2013.
- Campbell, Jamie, Kate Dufour, and Yalin Chen. "Retrieval-Induced Forgetting of Multiplication Facts and Identity Rule." *Memory & Cognition* 43, no. 4 (2015): 672–80. <https://doi.org/10.3758/s13421-014-0483-1>.
- Campbell, Jamie I. D., and Valerie A. Thompson. "Retrieval-Induced Forgetting of Arithmetic Facts." *Journal of Experimental Psychology: Learning, Memory, and Cognition* 38, no. 1 (2012): 118–29. <https://doi.org/10.1037/a0025056>.
- Chan, Jason C. K. "When Does Retrieval Induce Forgetting and When Does It Induce Facilitation? Implications for Retrieval Inhibition, Testing Effect, and Text Processing." *Journal of Memory and Language* 61, no. 2 (August 1, 2009): 153–70. <https://doi.org/10.1016/j.jml.2009.04.004>.
- Chan, Jason C. K., Kathleen B. McDermott, and Henry L. Roediger. "Retrieval-Induced Facilitation: Initially Nontested Material Can Benefit from Prior Testing of Related Material." *Journal of Experimental Psychology. General* 135, no. 4 (November 2006): 553–71. <https://doi.org/10.1037/0096-3445.135.4.553>.
- Christensen, David. "Preference-Based Arguments for Probabilism." *Philosophy of Science* 68, no. 3 (2001): 356–76.
- . *Putting Logic in Its Place: Formal Constraints on Rational Belief*. Oxford University Press, 2004.
- Crimmins, Mark. *Talk About Beliefs*. Stanford University, 2002.
- Dancy, Jonathan. *Practical Reality*. Oxford: University Press, 2000.
- . *Practical Reality*. *Practical Reality*. Oxford University Press. Accessed April 9, 2020. <https://www.oxfordscholarship.com/view/10.1093/0199253056.001.0001/acprof-9780199253050>.

- . *Practical Shape: A Theory of Practical Reasoning*. First edition. Oxford ; Oxford University Press, 2018.
- Dennett, Daniel C. “Brain Writing and Mind Reading.” *Minnesota Studies in the Philosophy of Science* 7 (1975): 403–15.
- DePaul, Michael R. “Truth Consequentialism, Withholding and Proportioning Belief to the Evidence.” *Philosophical Issues* 14, no. 1 (2004): 91–112. <https://doi.org/10.1111/j.1533-6077.2004.00022.x>.
- Dogramaci, Sinan. “Solving the Problem of Logical Omniscience.” *Philosophical Issues* 28, no. 1 (2018): 107–28. <https://doi.org/10.1111/phs.12118>.
- Dorr, Cian, Jacob M. Nebel, and Jake Zuehl. “The Case for Comparability,” n.d.
- Dorr, Cian, Jake Nebel, and Jake Zuehl. “Consequences of Comparability.” *Philosophical Perspectives*, forthcoming.
- Douven, Igor. “The Pragmatics of Belief.” *Journal of Pragmatics* 42, no. 1 (January 1, 2010): 35–47. <https://doi.org/10.1016/j.pragma.2009.05.025>.
- Dudai, Y. “How Big Is Human Memory, or on Being Just Useful Enough.” *Learning & Memory* 3, no. 5 (January 1, 1997): 341–65. <https://doi.org/10.1101/lm.3.5.341>.
- Eriksson, Lina, and Alan Hájek. “What Are Degrees of Belief?” *Studia Logica: An International Journal for Symbolic Logic* 86, no. 2 (2007): 183–213.
- Field, Hartry. “Mental Representation.” *Erkenntnis* 13, no. July (1978): 9–61.
- Finetti, Bruno De. “Sul Significato Soggettivo Della Probabilità.” *Fundamenta Mathematicae* 17 (1931): 298–329.
- Friedman, Jane. “Junk Beliefs and Interest-Driven Epistemology.” *Philosophy and Phenomenological Research* 97, no. 3 (2018): 568–83.
- . “The Epistemic and the Zetetic.” *Philosophical Review*, forthcoming.
- Goldman, Alvin I. *Knowledge in a Social World*. Oxford University Press. Accessed April 9, 2020. <https://www.oxfordscholarship.com/view/10.1093/0198238207.001.0001/acprof-9780198238201>.
- Gómez-Ariza, Carlos J., Angel Fernandez, and M. Teresa Bajo. “Incidental Retrieval-Induced Forgetting of Location Information.” *Psychonomic Bulletin & Review* 19, no. 3 (June 1, 2012): 483–89. <https://doi.org/10.3758/s13423-012-0228-8>.
- Hall, Richard J., and Charles R. Johnson. “The Epistemic Duty to Seek More Evidence.” *American Philosophical Quarterly* 35, no. 2 (1998): 129–39.
- Harman, Gilbert. *Change in View*. MIT Press, 1986.
- Harrison-Trainor, Matthew, Wesley H. Holliday, and Thomas F. Icard. “A Note on Cancellation Axioms for Comparative Probability.” *Theory and Decision* 80, no. 1 (January 1, 2016): 159–66. <https://doi.org/10.1007/s11238-015-9491-2>.
- Hedden, Brian. “Time-Slice Rationality.” *Mind* 124, no. 494 (April 1, 2015): 449–91. <https://doi.org/10.1093/mind/fzu181>.
- Hughes, Nick. “Dilemmic Epistemology.” *Synthese* 196, no. 10 (2019): 4059–90. <https://doi.org/10.1007/s11229-017-1639-x>.
- . “Epistemic Dilemmas Defended.” In *Epistemic Dilemmas*, edited by Nick Hughes, forthcoming.

- . “Who’s Afraid Of Epistemic Dilemmas?” In *Epistemic Dilemmas: New Arguments, New Angles*, edited by Scott Stapleford, Mathias Steup, and Kevin McCain, forthcoming.
- Insua, David Rios. “On the Foundations of Decision Making under Partial Information.” *Theory and Decision* 33, no. 1 (July 1, 1992): 83–100. <https://doi.org/10.1007/BF00133984>.
- Jackson, Elizabeth G. “The Relationship Between Belief and Credence.” *Philosophy Compass* 15, no. 6 (2020): 1–13. <https://doi.org/10.1111/phc3.12668>.
- Joyce, James M. “How Probabilities Reflect Evidence.” *Philosophical Perspectives* 19, no. 1 (2005): 153–78. <https://doi.org/10.1111/j.1520-8583.2005.00058.x>.
- Kaplan, Stephen, and Marc G. Berman. “Directed Attention as a Common Resource for Executive Functioning and Self-Regulation.” *Perspectives on Psychological Science* 5, no. 1 (2010): 43–57.
- Konek, Jason. “Comparative Probabilities.” In *The Open Handbook of Formal Epistemology*, edited by Richard Pettigrew and Jonathan Weisberg, 267–348. PhilPapers Foundation, 2019.
- Macfarlane, John. *In What Sense (If Any) Is Logic Normative for Thought?*, 2004.
- Meacham, Christopher J. G., and Jonathan Weisberg. “Representation Theorems and the Foundations of Decision Theory.” *Australasian Journal of Philosophy* 89, no. 4 (December 1, 2011): 641–63. <https://doi.org/10.1080/00048402.2010.510529>.
- Michaelian, Kourken. “The Epistemology of Forgetting.” *Erkenntnis* 74, no. 3 (2011): 399–424.
- Nelson, Mark T. “We Have No Positive Epistemic Duties.” *Mind* 119, no. 473 (2010): 83–102.
- Okasha, Samir. “ON THE INTERPRETATION OF DECISION THEORY.” *Economics & Philosophy* 32, no. 3 (November 2016): 409–33. <https://doi.org/10.1017/S0266267115000346>.
- Sainsbury, R. M. “What Logic Should We Think With?” *Royal Institute of Philosophy Supplements* 51 (March 2002): 1–17. <https://doi.org/10.1017/S1358246100008055>.
- Skorupski, John. *The Domain of Reasons*. Oxford ; Oxford University Press, 2010.
- . *The Domain of Reasons*. The Domain of Reasons. Oxford University Press. Accessed April 9, 2020. <https://www.oxfordscholarship.com/view/10.1093/acprof:oso/9780199587636.001.0001/acprof-9780199587636>.
- Smithies, Declan. “Ideal Rationality and Logical Omniscience.” *Synthese* 192, no. 9 (2015): 2769–93.
- Stefánsson, H. Orri. “On the Ratio Challenge for Comparativism.” *Australasian Journal of Philosophy* 96, no. 2 (2018): 380–90. <https://doi.org/10.1080/00048402.2017.1349159>.
- . “What Is ‘Real’ in Probabilism?” *Australasian Journal of Philosophy* 95, no. 3 (2017): 573–87. <https://doi.org/10.1080/00048402.2016.1224906>.

- Sturgeon, Scott. "Reason and the Grain of Belief." *Noûs* 42, no. 1 (2008): 139–65.
<https://doi.org/10.1111/j.1468-0068.2007.00676.x>.
- Thorstad, David. "Inquiry and the Epistemic." *Philosophical Studies*, January 2, 2021.
<https://doi.org/10.1007/s11098-020-01592-y>.
- Wallace, Jay. "Normativity, Commitment, and Instrumental Reason." *Philosophers' Imprint* 1 (2001).
- Wedgwood, Ralph. "Justified Inference." *Synthese* 189, no. 2 (2012): 273–95.
- White, Roger. "Epistemic Permissiveness." *Philosophical Perspectives* 19, no. 1 (2005): 445–59.
- Woodward, Elise, and Carolina Flores. "Epistemic Vigilance: In Defense of Epistemic Norms on Evidence-Gathering." [Manuscript], n.d.
- Wrenn, Chase B. "Why There Are No Epistemic Duties." *Dialogue: Canadian Philosophical Review / Revue Canadienne de Philosophie* 46, no. 1 (ed 2007): 115–36. <https://doi.org/10.1017/S001221730000158X>.