

What Does Weak Belief Tell Us about Full Belief?

Ben Holguín

April 2025

“I hope that some people see some connection
between the two topics in the title.”

–Saul Kripke, *Naming and Necessity*

1 Introduction

Epistemologists have recently begun to self-consciously distinguish two notions of belief: full and weak. Though there is some controversy over the question of which (if either) corresponds to the “ordinary” notion of belief—i.e., which is expressed by the natural language attitude verb ‘believe’—there is much that can be said about them independently of one’s preferred answer to this question.¹

For starters, the distinction between weak belief and full belief seems to track closely a distinction familiar from natural language: namely that between *thinking* that something is true and *being sure* that it is (Goodman & Holguín 2022). Here there is an asymmetric relation of entailment: anything I’m sure is true I think is true, but much of what I think is true I’m not sure of. I’m sure (and thus think) that Trump completed his first term, but I’m unsure whether he will complete his second—he’s pushing 80 and seems not to be in the best of health, after all. Still, I *think* he’ll complete his term, since I estimate it’s more likely than not that he does. Likewise, if I know that Jane has 70 of the 100 tickets to the upcoming lottery (while every other entrant has only 1), then although I’m not sure that Jane will win, I do at least think she will. Or consider a multiple-choice exam. Although I’ll often reject the presuppositions of a question like ‘What made you sure the answer to the question was A?’, typically I’m more than happy to answer a question like ‘What made you think the answer to the question was A?’.

With respect to weak belief, much of the theorizing about it has been done explicitly in terms of ordinary language ‘thinks’-reports. That is: the study of weak belief has self-consciously been the study of the attitude denoted by ‘thinks’ in sentences of the form ‘S thinks that p’ (Hawthorne *et al.* 2016, Dorst 2019, Rothschild 2020, Holguín 2022). By contrast, few contemporary epistemologists explicitly conceive of the study of full belief as the study of the attitude denoted by ‘is sure’ in sentences of the form ‘S is sure that p’.

¹ For discussion, see, e.g., Hawthorne *et al.* (2016), Moss (2019), Rothschild (2020), Nagel (2021) Clarke (2024), Williamson (forthcoming).

But as far as I can tell this is just a historical accident: once one distinguishes the ordinary notion of *being sure* from the sorts of quasi-technical notions one finds in discussions of Cartesian skepticism and the like, it's not clear that the epistemologist's notion of full belief is of something other than what is denoted by 'is sure' in ordinary surety-reports (cf. [Beddor 2020](#), [Goodman & Holguín 2022](#)).² So for the sake of vivacity and concreteness, I will treat as synonymous 'full believing' and 'being sure', on the one hand, and 'weakly believing' and 'thinking', on the other, and I will for the most part move freely between these ways of talking.

Full belief is meant to be the "traditional" epistemological notion of belief.³ Some widely held views about it include: that it is the inner analog of assertion; that it has a central role in folk-psychological explanations of action and inference; that it is the aim of inquiry; that it is subject to a knowledge norm ("Don't fully believe that p unless you know that p"); that it partly characterizes the differences between wanting, hoping, and wishing; and that we do not (and perhaps cannot) form full beliefs voluntaristically. Other questions about it engender more controversy: its relationship to "degrees of belief" (credence); its logic; and, in particular relevance to this paper, certain issues about its norms: can two agents in identical epistemic situations differ over which propositions they fully believe? Are moral and/or pragmatic factors relevant in the assessment of the rationality of full belief?

Weak belief, by contrast, has become a topic of interest to epistemologists only somewhat recently (at least under this guise). To the extent weak belief has interesting theoretical roles, they are quite unlike those of full belief: it is not the inner analog of assertion;⁴ it does not have the same central role in folk-psychological explanations of action and inference; it is not subject to a knowledge norm; and so on ([Hawthorne et al. 2016](#)). On the theory of weak belief I have developed (2022; see also [Dorst & Mandelkern 2023](#)), this much is unsurprising. For according to this theory, weak beliefs are *guesses*, and guesses clearly have a very different role in our psychological lives than that which we are sure of. Merely having a guess does not suffice for having the kind of attitude epistemologists take themselves to be targeting in their talk of 'full belief': there is no relevant sense in which one "takes for granted" the truth of one's guesses in the way one does one's full beliefs. So whatever is involved with weak belief, it's less demanding than what's involved with full belief.

Given the apparent conceptual autonomy of full and weak belief, one might reasonably expect that most of the questions of interest about one of the notions can be asked and answered without attending much to the workings of the other. But I believe weak belief actually has quite a lot to tell us about full belief. As proof of concept, I will discuss two debates about the nature of full belief, and I will argue in each case that there is much insight to be gleaned by considering what happens when the debate's central questions are transposed from full belief to weak belief.

² Compare the way [Gettier \(1963\)](#) presents the views of [Chisholm \(1957\)](#) and [Ayer \(2006\)](#).

³ For helpful background, see, e.g., [Williamson \(2000, forthcoming\)](#), [Wedgwood \(2012\)](#), [Buchak \(2014\)](#), [Staffel \(2016\)](#), [Genin \(2019\)](#), and [Goodman \(2023\)](#).

⁴ Though see [Mandelkern & Dorst \(2022\)](#).

First, I will discuss the debate about what explains the apparent impermissibility of full belief “toggling” (i.e., going back and forth between fully believing something and failing to fully believe it, without any change in evidence). And second, I will discuss the debate about whether (and if so, how) morality encroaches on full belief and, relatedly, whether (and if so, why) full beliefs can wrong. With respect to the first debate, I will argue that toggling seems as impermissible for weak belief as it seems for full belief, and that this has implications both for the debate about whether uniqueness versus permissivism is true of full belief, and for some of the standard, permissivist-friendly accounts of the impermissibility of full belief toggling. And with respect to the second debate, I will argue that *if* morality encroaches on full belief, it is most plausibly by way of weak belief, and in a manner that raises interesting issues about the connection between doxastic wrongdoing and doxastic voluntarism. In the concluding sections of the paper, I will draw some morals—mostly negative—about whether we can expect rational explanations of the norms associated with toggling, moral encroachment, and doxastic wrongdoing.

2 Weak belief

Before getting into this paper’s titular question, I’ll briefly survey some distinctive properties of weak belief.⁵ (Recall that, by stipulation, weak belief is the attitude denoted by ‘thinks’ in ordinary ‘think’-reports—i.e., sentences of the form ‘S thinks that p’. As such, I’ll be moving seamlessly between talk of what agents weakly believe and what they think. Furthermore, I won’t be distinguishing between *rational* weak belief and weak belief *simpliciter*. Unless stated otherwise, all the claims I make about weak belief are meant to hold for fully rational agents.)

Weak belief is weak. You can think something is true without being sure it is. There’s no infelicity whatsoever in speeches like ‘I think it will rain this week, but I’m not sure’. Indeed, you can think something is true while acknowledging there’s a substantial chance it’s not, as in: ‘I think Jane will win—she’s the favorite, after all—though I realize there’s a 30% chance someone else will instead’. When it comes to thinking that something true (as opposed to being sure of it), the evidential standards are weak (hence the name).

In fact, weak belief is *extremely* weak. You can think a proposition is true while being all but certain it is false. Suppose Jane has 30 of the lottery’s 100 tickets, and that every other entrant has only 1. Who do you think will win the lottery? ‘Jane’ remains an appropriate answer, despite the fact that your credence that she will win is substantially below .5. In fact, you can make the odds as low as you want. Answering a question like ‘Who do you think will win?’ with ‘The person who is most likely to win’ is not a way of evincing irrationality. The evidential standards for weak belief are extremely weak.⁶

⁵ For a more thorough treatment of these issues, see [Holguín \(2022\)](#) and [Dorst & Mandelkern \(2023\)](#).

⁶ More speculatively, I think you can think that a proposition is true while having credence 1 that it is false. To see why, imagine you’re at a coin factory that produces coins of real-valued bias between 0 and 1 inclusive. A coin of bias 0 always lands tails, a coin of bias 1 always lands heads, a coin of bias .5 lands heads 50% of the time, and so on. You’re given a coin whose bias has been determined by a random (but fair) process, corresponding to a continuous uniform distribution over all the reals between 0 and 1. You’re then allowed to flip it 1,000 times. It lands heads 666 times. What do you think the bias of the coin is? To my ears, there is exactly one felicitous three significant figure answer: ‘.666’. But your credence that the coin’s bias is .666 is 0.

*Weak belief is question-sensitive.*⁷ Suppose Jane has 30 of the 100 tickets (and everyone else has 1). Though ‘Jane’ is an acceptable answer to the question ‘Who do you think will win?’, it is *not* obviously an acceptable answer to the question ‘Do you think Jane will win?’. Instead, ‘I think Jane won’t win’ is the more natural answer of the two opinionated options.⁸ This suggests that we don’t weakly believe propositions *simpliciter*; rather, we weakly believe propositions *relative to certain questions*. By ‘question’ I officially mean a *partition*—a set of mutually exclusive and exhaustive propositions. Let J be the proposition that Jane will win, E_1 be the proposition that the first of the remaining 70 entrants will win, E_2 that the second will, and so on. In the setting of the lottery case, we can distinguish the following two questions: (i) $\{J, E_1, \dots, E_{70}\}$ and (ii) $\{J, \neg J\}$. The first we can think of as the “default” interpretation of the natural language question ‘Who will win the lottery?’, the second as the “default” interpretation of the natural language question ‘Will Jane win the lottery?’. You can think that Jane will win the lottery relative to the first question, but not relative to the second. So weak belief is question-sensitive.

Weak belief is permissive. You can weakly believe a proposition even when you recognize that there is a (very) substantial chance that it is false (at least relative to the right sorts of questions—more on this in a moment). But you can also *fail* to weakly believe that proposition, precisely on the grounds that there is a substantial chance that it is false. Though there is no infelicity in reporting oneself as thinking Jane will win the lottery (when her 30 tickets are a plurality), there is equally no infelicity in reporting that one fails to think she will win. Speeches like ‘It’s not fair to say I think Jane will win; I merely think she’s most likely to win’ or (more bluntly) ‘I don’t have an opinion on whether Jane will win’ are perfectly fine in response to the question ‘Who do you think will win the lottery?’. When one answers in this way, one is reporting the fact that relative to the question $\{J, E_1, \dots, E_{70}\}$, the strongest thing one thinks is that $J \vee E_1 \dots \vee E_{70}$. This suggests that weak belief is *permissive*. In particular: two agents with identical bodies of evidence can differ over whether they think Jane will win the lottery, with both being perfectly rational in what they weakly believe. One agent can be of the opinion that Jane will win; the other can be agnostic, thinking only that one of the 71 entrants will win.

At minimum, then, weak belief is permissive with respect to the issue of whether one weakly believes some particular answer to the relevant question, or instead fails to have any weak belief on the matter at all (except perhaps the trivial one—viz., that some or other of the question’s answers is correct). But are there cases where weak belief is permissive with respect to two more *complete* answers to the relevant question?

I think so. What convinces me is the following simple kind of example. I show you a fair coin and then launch it hundreds of feet in the air. While the coin is on its ascent I ask you, ‘How do you think the coin will land?’. I submit that any of the following answers are perfectly acceptable: ‘Heads’, ‘Tails’, or some variant on ‘I don’t have a view’. Taking these

⁷ Recently there has been a considerable amount of work arguing that full belief is also question-sensitive—see, e.g., Yalcin (2018), Drucker (2020), Hoek (2022, 2025), and Goodman & Salow (2023, forthcoming) for some representative samples. Note, however, that with the exception of Goodman & Salow’s work, the ways in which full belief has been argued to be question-sensitive are quite unlike the way in which I’m about to argue weak belief is.

⁸ Here it is important to consider the questions in isolation, as if they were each asked discourse-initially.

judgments at face-value, it appears that one can either think that the coin will land heads or think that the coin will land tails (or think neither of these things, and instead merely that it will land some way or other). Heads and tails are incompatible complete answers to the question ‘How will the coin land?’, yet both appear permitted given one’s evidence. So there are cases where for some single body of evidence, proposition p , and question $Q^?$, a person with that evidence can either weakly believe p , weakly believe $\neg p$, or merely weakly believe that $p \vee \neg p$.

The respects in which weak belief is permissive can be captured in general terms (Holguín 2022):⁹

COGENCY

It is rationally permissible for p to be the strongest thing S thinks relative to $Q^?$ just in case:

- (1) p is a union of complete answers to $Q^?$.
- (2) If there is a complete answer to $Q^?$, r_1 , such that r_1 doesn’t entail p , then there is no other complete answer to $Q^?$, r_2 , such that: r_2 entails p , but S ’s evidence supports r_1 more than r_2 .

What this principle says, intuitively, is the following. Suppose n is the number of complete answers to $Q^?$ (if $Q^?$ is a polar question, $n = 2$; if it’s the natural interpretation of ‘Who will win the lottery?’ in the version of the lottery where Jane has 30 tickets, then $n = 71$; etc.). Relative to $Q^?$ you can be as opinionated (or as unopinionated) as you want about its answer, as measured by the number of complete answers that are compatible with what you weakly believe (1 is maximally opinionated, n is maximally unopinionated). But if what you weakly believe rules out all but m answers, those answers must be the m that are best supported by your evidence (including ties). So if $m = 1$ (i.e., you weakly believe a single complete answer to the question), your answer needs to be first or tied for first with respect to its evidential support. If $m = 3$, your answer needs to be a disjunction of the first, second, and third best options (if two answers are tied for first, they are automatically the first and “second” best options). And if $m = n$, then you satisfy the condition trivially.

I will be assuming the correctness of COGENCY in what follows. A more thorough explanation of the principle can be found in the existing literature.¹⁰ There are various worries that have been raised about it in recent years, especially in connection to issues of “inter-question” constraints on weak belief.¹¹ But for the purposes of this paper I set such worries aside, since they are orthogonal to the main points of inquiry.

Weak belief is voluntaristic. Consider again the coin flip case. As we know from COGENCY, you are rationally permitted to do any of the following: think that the coin will land heads, think that the coin will land tails, or merely think that the coin will land either heads or tails. Suppose in fact you think the coin will land heads. In virtue of what is this what you think?

⁹ Note that this is a slight departure from the presentation of the principle in Holguín (2022), in that (i) it talks about evidential support rather than credence and (ii) it allows for tie-breaking. With respect to (i), this difference is immaterial. With respect to (ii), I think it’s just a mistake to have a theory of weak belief that prohibits tie-breaking.

¹⁰ See, e.g., Holguín (2022), Dorst & Mandelkern (2023), Quillien et al. (2023), and Skipper (2023).

¹¹ See, e.g., Teague (2024) and Pearson (forthcoming).

As far as I can tell, the answer has to be something along the lines of: it's how you *made up your mind*. What else could really be going on here? By hypothesis, you know both possible outcomes are equally likely, so it's not like your evidence favors heads. Nor does it seem plausible that you simply "found" yourself with the weak belief that the coin will land heads. (Well, perhaps in some cases this is what happens; what is implausible is that this is *always* what's going on.) Phenomenologically speaking, the experience of coming to weakly believe the coin will land heads feels just like the experience of making up your mind between two identical cans of soup, or of forming an intention to take one of two equally good paths to work, or of choosing an arbitrary aspect of your visual experience to attend to. Each of these things is something you can find yourself doing, as it were; but each is also something you sometimes do at will. I contend that weak belief is no different in this respect. And if I am right about this, then we have two dimensions of voluntarism for weak belief. First, there is voluntarism with respect to the strength of your opinion. And second, there is voluntarism with respect to tie-breaking.

One final point before finally getting to the two debates about full belief. Whatever one makes of the claims this section has made about weak belief (*qua* denotation of the ordinary propositional attitude verb 'thinks'), the analogous claims about the propositional attitude *guessing* strike me as basically platitudinous. If I tell you I'm going to select a number from 1 to 1,000 at random (using a genuinely fair randomizing device) and then ask you to guess what it will be, you do just fine guessing any of 1-1,000. So guessing is extremely weak (and tie-breaking permissive). And if instead I ask you to guess whether it will be (say) 666, you do much better guessing that it won't than that it will, but also seemingly do fine by refusing to have a non-trivial guess. So guessing is question-sensitive (and agnostic permissive). And to the extent I have an idea of what 'voluntaristic' could mean, guessing seems as voluntaristic as anything could be.

In other work I have argued that weak belief *just is* guessing (Holguín 2022; see also Dorst & Mandelkern 2023). But whether or not that view is correct, it's worth emphasizing that this paper equally well could have been titled 'What Does *Guessing* Tell Us about Full Belief?'. Had it been so-titled, it would have advanced the thesis that the plausibility of certain claims about full belief are sensitive to the workings of guessing.

I believe this thesis that is every bit as interesting as the one stated in the idiom of weak belief. Still, I think there are good reasons to center the discussion around the idiomatic thesis. Chief among them is the fact that the natural language expression 'guess' is messy in a way that 'think' is not. In particular, 'guess' is primarily used as a speech act verb, and rational assessments of guesses of *that* sort tend to involve all sorts of pragmatic factors that are irrelevant to the epistemic assessment of one's doxastic state. For example, one might have strategic reasons to answer a trivia question in a way that departs from one's private estimation of which answer is most likely to be correct. In such cases it is natural to interpret 'guess' in such a way that it tracks one's reported answer. But one's reported answer need not be the content of any theoretically interesting psychological state. The more interesting state is the one one would be in if one were asked to guess what the answer to the trivia question is *but to keep it to oneself*—or, as I've argued elsewhere: by

simply looking at what one *thinks* the answer to the question is. So I will continue to focus on weak belief as such in what follows.

With that, we can now turn to the first substantive question of the paper: what does weak belief tell us about toggling (and its relationship to permissivism) for full belief?

3 Toggling and permissivism

Consider the following two “uniqueness” principles for full belief:¹²

FB INTERPERSONAL UNIQUENESS

For any body of evidence E and proposition p , there is exactly one full belief attitude Φ such that: for any agent S , if E is S 's body of evidence, then S is rationally permitted to Φ that p .

FB INTRAPERSONAL UNIQUENESS

For any agent S , body of evidence E , and proposition p , there is exactly one full belief attitude Φ such that: if E is S 's body of evidence, then S is rationally permitted to Φ that p .

(And let FB INTERPERSONAL PERMISSIVISM and FB INTRAPERSONAL PERMISSIVISM be the negations of the corresponding principles.)

What FB INTERPERSONAL UNIQUENESS says, in words, is if you fix a body of evidence, then you fix which full belief attitude is rational to take toward a given proposition for any agent who has that body of evidence. If FB INTERPERSONAL UNIQUENESS is true, then there can be no “reasonable disagreement” about what to fully believe among agents who share the same evidence. What FB INTRAPERSONAL UNIQUENESS says, by contrast, is that if you fix a body of evidence *and an agent who has that evidence*, then there is exactly one full belief attitude that that agent can rationally take toward a given proposition.

FB INTERPERSONAL UNIQUENESS is strictly stronger than FB INTRAPERSONAL UNIQUENESS (by extension, FB INTERPERSONAL PERMISSIVISM is strictly weaker than FB INTRAPERSONAL PERMISSIVISM). Indeed, a number of theorists accept the former while rejecting the latter.¹³ Some hold the view that which full belief attitude is licensed toward p depends not only on an agent's evidence, but also on such things as the agent's ur-prior and/or epistemic values. On this sort of view, two agents with the same evidence can rationally differ over which full belief attitude they take to p , in virtue of differing in these other respects. A proponent of such a view would thus be committed to rejecting FB INTERPERSONAL UNIQUENESS. But they would *not* be committed to rejecting FB INTRAPERSONAL UNIQUENESS, since they could maintain that once you hold fixed the agent, you then fix which particular rational full belief attitude is licensed given their evidence.

Indeed, once one fixes on a particular agent with a particular body of evidence it's at least a bit more puzzling to see how it could be “up for grabs” which full belief attitude that

¹² By ‘full belief attitude’ I intend the tripartite distinction between being sure that p , being sure that $\neg p$, and being unsure whether p (i.e., suspending judgment on p). For a helpful survey of (the wide variety of) “uniqueness” principles, see Kopec & Titelbaum (2016).

¹³ See, e.g., Douven (2009), Kelly (2013), and Schoenfield (2014). Though for criticisms of these positions see Weisberg (2020) and Dandelet (2023).

agent takes toward p .¹⁴ This is especially so if one thinks that the sorts of factors that might undermine FB INTERPERSONAL UNIQUENESS—namely differences between agents in their ur-priors, epistemic values, and so on—are not themselves “up for grabs” for any individual agent. For example: one might think that although different agents may be permitted to be such that they have different ur-priors—and so are permitted to have different full beliefs, even given a shared body of evidence—no single agent is permitted to be such that they have some ur-prior other than the one they actually have.

Things are a bit murkier with so-called “epistemic values”. Consider a standard example: the relative weights assigned to the disvalue of having a false full belief versus the value of having a true full belief (Kelly 2013). It is plausible that different agents give different verdicts on this question (at least at some level of abstraction). Suppose we’re in the frame of mind where the full belief attitude an agent can rationally take toward p depends at least in part on how such weights are assigned: perhaps if the agent goes with the first weighting they ought to suspend on whether p , whereas if they go with the second they ought to be sure that p . Given that the agent in fact “adopts” a certain weighting (say the one that compels them to suspend on p), are they permitted to be such that they adopt a different weighting, and with it a different full belief attitude toward p ? Some theorists think the answer is no: they take epistemic values to be “immodest” in the sense that, by the lights of any particular weighting, *that* particular weighting is the only rationally permissible one for the agent to have (cf. Kelly 2013, Schoenfield 2014). One who has this conception of epistemic values may plausibly deny FB INTERPERSONAL UNIQUENESS while accepting FB INTRAPERSONAL UNIQUENESS.

But it’s also not clear that such a conception of epistemic values is mandatory. To the extent one thinks that agents are sometimes rationally permitted to have different epistemic values, one could plausibly reject FB INTRAPERSONAL UNIQUENESS as well.

However, there is an influential argument, due to White (2005, 2013), for the conclusion that FB INTRAPERSONAL UNIQUENESS must be true. It goes roughly as follows. Suppose FB INTRAPERSONAL PERMISSIVISM is true, and that, in particular, given S ’s evidence, S is rationally permitted to be in either of the following two states: (i) being sure that p (while not being sure that $\neg p$) or (ii) being sure that $\neg p$ (while not being sure that p).¹⁵ In fact, suppose S *knows* that S is in this permissive situation. Suppose in front of S are two pill bottles, one filled with red pills, the other with blue pills.¹⁶ If S takes a red pill, S will become sure that p and not sure that $\neg p$; if S takes a blue pill, S will become sure that $\neg p$ and not sure that p . (Whichever pill is taken last trumps the effects of the previous pill.)

Here’s the core question: epistemically speaking, what’s to stop S from taking one of the pills—say the red one? After all, S knows that the state they will be in after taking the red pill is one of the rationally permitted full belief states. Likewise for the blue pill. So what (epistemic) reason does S have not to take it? In fact, what’s to stop S from just taking pills one after the other, toggling happily between the full belief that p and the full belief that

¹⁴ Though for defenses of this sort of view, see, e.g., Titelbaum (2015), Weisberg (2020), and Jackson (2021).

¹⁵ The argument works just as well if the two options are (i) being sure that p (while not being sure that $\neg p$) or (ii) being unsure whether p (while neither being sure that p nor sure that $\neg p$), or if all three options are on the table.

¹⁶ The appeal to pills is inessential; the argument works just as well with coin flips or arbitrary decisions.

$\neg p$? Such behavior *seems* manifestly irrational (or at least pathological in some distinctly epistemic way). But if FB INTRAPERSONAL PERMISSIVISM is true, then it's not clear what if anything could be wrong with it. So FB INTRAPERSONAL PERMISSIVISM must be false.

With one qualification (to be introduced shortly), I am in agreement with White's judgments about the case: toggling of this sort is manifestly epistemically defective. But I think it's a mistake to infer from this that FB INTRAPERSONAL UNIQUENESS must be true (and so FB INTRAPERSONAL PERMISSIVISM false). This is for reasons related to the nature of weak belief. In particular, it's because (i) toggling is epistemically inappropriate for weak belief, but (ii) the weak belief-analog of FB INTRAPERSONAL PERMISSIVISM—namely,

WB INTRAPERSONAL PERMISSIVISM

It is not the case that: for any agent S, body of evidence E, and proposition p, there is exactly one weak belief attitude Φ such that: if E is S's body of evidence, then S is rationally permitted to Φ that p.

—is *much* harder to deny. And this means that in the case of weak belief, we need to account for the impermissibility of toggling in terms of something other than the idea that (intrapersonal) uniqueness is true for weak belief. But once we have such a story for weak belief, it's hard to see why we wouldn't want to say the same about full belief.

Here's my argument more explicitly:

- (P1) WB INTRAPERSONAL PERMISSIVISM is true.
- (P2) Weak belief toggling is epistemically impermissible.
- (P3) If WB INTRAPERSONAL PERMISSIVISM is true yet weak belief toggling is epistemically impermissible, then FB INTRAPERSONAL PERMISSIVISM is compatible with the impermissibility of full belief toggling.
- (C) Thus, FB INTRAPERSONAL PERMISSIVISM is compatible with the impermissibility of full belief toggling.

The case for (P1) is grounded in the examples discussed in §2. You can think a fair coin will land heads, you can think a fair coin will land tails, or you can think neither of these things in particular. None of these weak belief states is mandated by your evidence. Nor is there any reason to think that non-evidential epistemic factors (like your “epistemic standards”) settle the question in favor of one particular state. Indeed, if they *were* able to do this, they'd have to settle it in favor of agnosticism (the choice between heads and tails is a pure Buridan case, after all). But there is just no intuitive force to the thought that certain kinds of people—those whose epistemic standards allegedly disvalue fully believing falsely more than they value fully believing truly—would be making a mistake (by their own standards) in thinking that the coin will land heads rather than being agnostic on the question. What possibly could be at stake here? It's just a guess!¹⁷

¹⁷ Dorst & Mandelkern (2023) defend a theory of guessing/weak belief that could be leveraged into an objection against this line of thought. On their view, agents who form weak beliefs seek to optimize a tradeoff between informativity (the number of complete answers to the relevant question ruled out by the weak belief) and

The case for (P2) is just an appeal to brute intuition. There is something pathological seeming about a person who over the course of the flip goes,

‘I think it will land heads. Actually I think it will land tails. No, heads. No, tails. Heads. Tails. Heads. OK—Tails.’

There is something equally pathological seeming about the person who goes,

‘I think it will land tails. On second thought, I don’t have a view. On third thought, tails. Actually no, just heads or tails—that’s all I’m willing to guess. No wait, tails!’

Similar points go through with pills. Suppose I flip the coin hundreds of feet in the air. While it’s on its ascent, I ask you how you think it will land, and you say heads. I then hand you two bottles of pills, red and blue (the coin is still way up there). The red pills make you think the coin will land heads (without affecting your credence function or anything like that); the blue pills make you think the coin will land tails. It seems epistemically problematic (if not simply irrational) to start popping these pills, despite the fact that you know full well that the (weak) doxastic state you’ll be in after any given dose is one that is perfectly acceptable to be in given your evidence. *Mutatis mutandis* if we replace one of the bottles with a bottle of green pills that make you agnostic, or if we just add it as an additional option. Toggling your weak beliefs in this way is impermissible, whether by direct voluntaristic control or by medication. Hence (P2).

Before turning to (P3), an observation and a moral about weak belief toggling. The observation: although it seems pathological to toggle one’s weak beliefs *repeatedly* (whether between different incompatible opinions or between different degrees of agnosticism), it’s much less clear that it’s problematic to toggle *once*.¹⁸ To my ears, someone who (mid-flip) goes,

‘I think it will land heads. Actually, on second thought, I think it will land tails.’

is doing just fine. Likewise for someone who says something like:

accuracy (the likelihood that the weak belief is true). The optimization works in terms of a “J-value”, which represents the degree to which the agent values informativity (the higher the J-value, the more the agent prioritizes informativity). Excluding cases of ties, an agent’s J-value in combination with their evidence will determine a unique optimal answer to the question at hand. So, one might argue, weak belief is only as permissive as the setting of a J-value. If for any agent there is exactly one J-value that agent is permitted to have, then—excluding ties—for any given question there will be exactly one (possibly incomplete) answer the agent is permitted to weakly believe. It only takes a bit of squinting to see this conclusion as implying something like uniqueness for weak belief.

Two quick points in response. First, the obvious point: so long as it’s possible to weakly believe in the case of credal ties, no amount of fixity of the J-value will be enough to salvage a uniqueness principle for weak belief, since no possible J-value setting will imply that heads is a better guess than tail (or vice-versa). Second, Dorst and Mandelkern introduce the ideology of J-values in a deflationary spirit, and seem amenable to the thought that the facts that determine an agent’s J-value in a given circumstance are flexible (and perhaps even under their direct voluntaristic control; see their section 3.3). In light of this, I think it would be odd to use the ideology of J-values in an argument against permissivism about weak belief. There is little in our intuitive judgments about guessing/weak belief to suggest that the way we *happen* to make informativity/accuracy tradeoffs in certain circumstances (as revealed by the weak beliefs we happen to adopt in those circumstances) corresponds to an epistemically significant constraint on the range of such tradeoffs we permissibly *could* make.

¹⁸ Cf. Chang (1997) on decisions between incommensurate options and Moss (2014) on “credal dilemmas”.

‘I think it will land heads. Actually, never mind—I don’t have a guess one way or the other.’

Or:

‘I don’t have a particular guess. Actually, scratch that: I think it will land tails’.

Changing one’s mind *twice* is where pathology starts to creep in:

‘Heads. No, wait, tails. Actually, heads—final answer.’

Speeches like this are notably stranger than those that report a one-off switch.

As far as I can tell, then, the norm on weak belief toggling is something like this:¹⁹

THE LAW OF WEAK BELIEF TOGGLING

Toggling a weak belief once is fine. Twice is borderline. Thrice or more is unacceptable.

So really it’s not that weak belief toggling is impermissible *per se*; it’s *repeated* weak belief toggling that’s impermissible.²⁰

That’s the observation. The moral is that even if we lacked independent reason to believe WB INTRAPERSONAL PERMISSIVISM, the fact that *some* amount of toggling is permissible (even if only just a touch) makes it especially difficult to see how the impermissibility of repeated toggling could be part of an argument against WB INTRAPERSONAL PERMISSIVISM. The distinction between one-off and repeated toggling is too fine-grained to be tracked by general principles about uniqueness versus permissivism for weak belief.

Finally, (P3): if WB INTRAPERSONAL PERMISSIVISM is true yet weak belief toggling is epistemically impermissible, then FB INTRAPERSONAL PERMISSIVISM is compatible with the impermissibility of full belief toggling. I take this to be the most controversial premise of the argument. The main consideration in its favor is an abductive one: it would be surprising if whatever explains THE LAW OF WEAK BELIEF TOGGLING wouldn’t also account for the inappropriateness of full belief toggling.

Indeed, I suspect that THE LAW OF WEAK BELIEF TOGGLING follows from the following more general principle:²¹

¹⁹ A more precise statement of the principle would involve a restriction to “sufficiently short” periods of time, and would make explicit that the options one toggles between are rationally permitted. Also, as Jeremy Goodman (p.c.) has pointed out, it’s less obvious that there’s something wrong with “toggling” one’s weak beliefs by monotonically shifting one’s level of opinionatedness (i.e., starting opinionated and progressively becoming more agnostic, or starting agnostic and progressively becoming more opinionated). What doesn’t seem permissible is vacillating between levels of opinionatedness (i.e., going from opinionated to agnostic and then back to opinionated) or between ties. I’ll suppress this subtlety in what follow, but concerned readers can interpret ‘toggling’ as ‘toggling back and forth between levels of opinionatedness and/or ties’.

²⁰ There is a certain affinity to the central observation of Woodard (2020), which concerns a “puzzle about fickleness”. Woodard observes that there is an appearance of pathology in agents who repeatedly change their mind after *redilberating* (i.e., reexamining a fixed body of first-order evidence with respect to some particular question), but no such appearance of pathology in agents who change their mind *once* in this way. I lack the space to explore the relationship between these observations about toggling and redeliberation in the detail they deserve. But here I’ll emphasize (in agreement with Woodard—see her section 5) that toggling and redeliberation are fundamentally different phenomena: the former is a kind of arbitrary switching; the latter is non-arbitrary switching. Consequently, even if Woodard is right that a single change of mind in response to redeliberation is appropriate whereas multiple changes of mind are (typically) inappropriate (a supposition I’m personally inclined to resist), it’s not clear that this is for the same reason one-off toggles are appropriate while repeated toggles are inappropriate. I’d sooner call it a coincidence.

²¹ Roeber (2020, pp. 422-3) seems to be sympathetic to something like this principle.

THE FUNDAMENTAL LAW OF TOGGLING

If Φ is an (intrapersonally) permissive attitude, then toggling between Φ -states once is fine. Twice is borderline. Thrice or more is unacceptable.

This is because the judgments that motivate the instance of the principle where Φ is weak belief do not feel idiosyncratic; rather, they seem to be reflections of the fact that, as a general rule, when an agent is faced with multiple permissible options, it's reasonable for them to arbitrarily change their mind once, but not too many times.

To see this, consider another putatively permissive attitude, *preference*. Suppose S is in the midst of a Buridan case—say, with a choice between whether to accept a bet that pays \$5 if a fair coin lands heads and \$0 if tails, or vice-versa (\$5 for tails, \$0 for heads). Suppose S is asked which bet they prefer. By my lights, there is nothing infelicitous in S saying any of 'I prefer heads', 'I prefer tails', or 'I don't have a preference between them'. Taking such judgments at face value, it must be permissible for S to prefer the heads bet, to prefer the tails bet, or to have no preference between them. Can S toggle these preferences? On the one hand, it seems reasonable for S to say something like:

'I prefer heads. Actually, on second thought I prefer tails.'

On the other hand, it doesn't seem reasonable for S to say something like:

I prefer heads. Actually tails. No, wait—heads. Actually no, I don't have a preference; just give me whichever.'

So, one-off preference toggling appears mostly fine, whereas repeated preference toggling seems pathological. This mirrors the judgments in the case of weak belief. This fact outcries for explanation. THE FUNDAMENTAL LAW OF TOGGLING provides one.

Considerations of this sort suggest that we should resist arguments, like White's, that begin with the observation that full belief toggling is impermissible and end with the conclusion that FB INTRAPERSONAL UNIQUENESS is true. For permissivists about full belief have a straightforward explanation of the impermissibility of full belief toggling, namely THE FUNDAMENTAL LAW OF TOGGLING. Of course, this is not to deny that the truth of FB INTRAPERSONAL UNIQUENESS plausibly suffices to explain the impermissibility of full belief toggling. But the point is that it's not *necessary* for it, which is what is needed if the argument against FB INTRAPERSONAL PERMISSIVISM is to be dialectically effective.

That said, THE FUNDAMENTAL LAW OF TOGGLING is a bit of a double-edged sword for certain proponents of FB INTRAPERSONAL PERMISSIVISM. This is because some of these theorists have attempted to answer White's challenge by explaining the impermissibility of full belief toggling in a way that is compatible with FB INTRAPERSONAL PERMISSIVISM, but that does not involve an appeal to attitude-neutral (though admittedly yet unexplained) principles about the impermissibility of toggling, like THE FUNDAMENTAL LAW OF TOGGLING.²²

²² The only two theorists I'm familiar with who treat the impermissibility of full belief toggling as derivable from the nature of toggling itself are Roeber (2020) and Jackson (2021). As far as I can tell, however, neither has a story that is more developed than an appeal to a principle like THE FUNDAMENTAL LAW OF TOGGLING. Roeber simply insists that full belief toggling "would be nuts" (p. 423). Jackson, who goes into a bit more detail, says the following (fn. 48, p. 324):

Instead, these theorists have tried to explain it in terms of some of the specific properties of *full belief*. But given that weak belief toggling is similarly impermissible, one might start to wonder whether such explanations could really be getting to the heart of the matter.

To illustrate the worry, I'll focus on Titelbaum's (2015) account.²³ Titelbaum seeks to explain the impermissibility of full belief toggling in terms of two observations about full belief: (i) there is a synchronic norm on full belief that "requires an agent's beliefs to be consistent, and requires her to believe (at least some of) the consequences of what she believes" (p. 682); and (ii) many of our full beliefs are formed by reasoning from other full beliefs, which is a causal process that occurs over time. The thought, then, is that agents who toggle their full beliefs are likely to violate the norm described in (i). Why? Here's Titelbaum (p. 682):

Consider an agent who believes *p* and wants to determine some consequences of that belief. She begins reasoning from *p* to various conclusions that follow. As her train of reasoning continues, she considers the consequences of those conclusions, and the consequences of those consequences. She may no longer be mentally attending to *p* at all. But now suppose that as her reasoning worked through these further consequences, the agent's attitude toward *p* swung around so that she believed $\neg p$. (I realize this doesn't happen to normal human agents—the point is to see why it would be bad if it did.) An abrupt change in belief like that would vitiate the agent's reasoning from *p*; the reasoning process would no longer be one of drawing out consequences of her beliefs.

In short: if agents engaged in full belief toggling, we would expect them to end up with inconsistent beliefs (as when they toggle on the full beliefs that are the premises in the deductions they perform over some period of time). This would be bad. So it's unsurprising that full belief toggling is epistemically impermissible.

I'm skeptical.²⁴ For starters, the impermissibility of full belief toggling feels like a matter of necessity: *if* full belief toggling is in fact impermissible (as it seems to be), then it

There is also good reason to think that the fact that both attitudes are permitted for me at a time does not entail that toggling is *epistemically* permitted, either. Suppose I wake up in the morning and pour myself a bowl of cereal, and grab a spoon from the silverware drawer to eat it. It is permissible for me to eat my cereal with any of the spoons in the drawer; in fact, there is no reason to pick one spoon over another. However, this does not entail that "toggling" between spoons is permissible for me—I have already picked a spoon and once I have picked one, there is no sense switching to another. In the same way, simply because several attitudes are epistemically permitted for me does not entail that it can be rational for me to switch whenever I'd like.

I don't see this passage as doing much beyond insisting that I am permitted to take any spoon but not permitted to toggle between them. For reasons I will briefly discuss in §5, my own sense is that this is about all there is to say on the matter.

²³ Other FB INTRAPERSONAL PERMISSIVISM-friendly explanations of the impermissibility of full belief toggling include those of Weisberg (2020) and Dandeleit (2023). I suspect that the main objection I have to Titelbaum's account (see the following page) apply to these accounts as well, though I lack the space to do more than baldly state as much.

²⁴ One small point I relegate to a footnote: any explanation of the impermissibility of full belief toggling that is grounded in synchronic consistency norms won't be of much help in explaining the impermissibility of toggling between full belief and *suspension* (as opposed to toggling between two incompatible full beliefs). An agent who toggles between full belief and suspension might well find themselves fully believing the conclusions of arguments whose premises they no longer fully believe. But that doesn't entail that such an agent would be

couldn't be permissible. But (ii)—the observation that full beliefs are formed via a process of reasoning that extends over time—is only contingently true. We could have been creatures who don't need to reason (or, more conservatively: it could have been that there are creatures who form full beliefs without needing to reason). And even if we had been such creatures, it still would have been impermissible for us to toggle our full beliefs. But why should this be so, if the explanation of the impermissibility of belief toggling here in actuality is grounded in facts that don't hold in these counterfactual possibilities?

Relatedly, even if this account can explain the impermissibility of toggling a *single* full belief, it doesn't seem to do much to explain the apparent impermissibility of toggling a whole web of full beliefs. Consider a version of White's pill case where for some fully believed proposition p , the red pill does the following four things. (The blue pill works the exact same way, except with $\neg p$ in place of p .) First, the pill eliminates the full belief that p . Second, for any proposition q , if the agent fully believes q because they deduced it (in part) from p , then the pill eliminates the full belief that q and replaces it with the full belief that if p , q (if it wasn't there already). Likewise, if the agent is currently in the process of figuring out whether q follows from p , then the pill transforms the agent's attitude to the relevant premise of the deduction from a full belief to a supposition, so that the output of an affirmative verdict is the full belief that if p , q (rather than the full belief that q). Third, the pill adds the full belief that $\neg p$. Fourth, for each proposition r such that the agent fully believes that if $\neg p$, then r , the pill adds the full belief that r . An agent who pops these pills back and forth seems to be forming full beliefs in an epistemically impermissible way. But there's no threat that doing so will result in the agent violating the consistency requirements on their full beliefs, since by hypothesis the pill manages them on the agent's behalf.

These particular complaints aside, I think the most serious worry for Titelbaum's explanation is that it seems implausible when transposed into an account of the apparent impermissibility of *weak belief* toggling. I contend that the person who toggles back and forth between guessing that the coin will land heads and guessing that the coin will land tails is of no greater risk of having inconsistent weak beliefs than the person who sticks to their guns and keeps their guess as heads from the outset. I think this is intuitively clear in its own right, at least in the sense that I don't see how to fill in the details of the case such that we could reasonably expect the toggler to find themselves with inconsistent guesses (in virtue of their toggling).

But I also think there are theoretical reasons to doubt that weak belief toggling is likely to begot weak belief inconsistency. Chief among them is the fact that we don't use weak beliefs as premises in theoretical or practical reasoning. Nor do we usually reason through the consequences of our weak beliefs, unless explicitly prompted to do so. And to the extent that we *do* engage in such reasoning, it seems to be by way of full belief (though see [Pearson forthcoming](#) for an opposing perspective). If I'm in a position to say something like:

likely to have inconsistent beliefs, or would be such as to fail to believe (at least some of) the consequences of what they fully believe. So something else must be explaining our aversion to toggling in these cases.

‘I think the coin will land heads. Therefore, I think I’ll win the bet.’

—it’s only because I’m *sure* that if the coin will land heads, then I’ll win the bet.²⁵ Thus, even if certain aspects of our psychological lives involve temporally-extended epistemic processes for which it’s important that our full beliefs remain constant, it’s not clear that there are any such processes for which it’s important that our *weak beliefs* remain constant. When it comes to reasoning, weak belief is an epiphenomenon. Yet weak belief is as hostile to toggling as full belief is.

Consequently, I see little hope for an account of the impermissibility of weak belief toggling that is along the lines of Titelbaum’s account of the impermissibility of full belief toggling. And I think this transposition failure is a strong (even if not decisive) reason to find accounts like Titelbaum’s unsatisfactory. We do better avoiding parochial approaches to the problem. And an explanation of the impermissibility of full belief toggling in terms of a general principle like THE FUNDAMENTAL LAW OF TOGGING at least manages that.

Of course, this is to say nothing about *why* a principle like THE FUNDAMENTAL LAW OF TOGGING might happen to be true. I’ll return to this question briefly in §5. But for now I want to turn from toggling to moral encroachment.

4 Moral encroachment and doxastic wronging

Stated roughly, *moral encroachment* is the thesis that the epistemic status of a doxastic state can depend on its moral features (cf. Moss 2018). Somewhat more precisely, it is the thesis that if p is morally significant but q is not, then one can fail to be epistemically justified in Φ -ing that p but epistemically justified in Φ -ing that q , even when (by all traditional measures) one’s evidence for p is of the exact same quality as one’s evidence for q .²⁶ A related thesis is that of *doxastic wronging*: roughly, the thesis that if p is morally significant, then Φ -ing that p can in certain circumstances constitute a moral wrong in and of itself (i.e., independently of whatever actions one takes “downstream” of one’s Φ -ing that p).²⁷

These two theses are separable: one could accept that the epistemic status of one’s Φ -ing is sensitive to moral factors without thinking that states of Φ -ing ever wrong (in and of themselves). Or one could accept that states of Φ -ing sometimes wrong (in and of themselves) but that their epistemic status is never sensitive to moral factors. However, these

²⁵ To see this, consider a case in which I weakly rather than fully believe that I’ll win the bet if the coin lands heads. For example: suppose I’m credence .5 that the coin will land heads, credence .6 that if the coin lands heads, I’ll win the bet, and credence 1 that if the coin lands tails, I won’t win the bet. I submit that in these circumstances, competent speakers are unlikely to go around making speeches like:

‘I think the coin will land heads. And I think that if the coin lands heads, I’ll win the bet.
Therefore I think I’ll win the bet.’

And this is because the only felicitous reply to the question ‘Do you think you’ll win the bet?’ is ‘No’. As such, I think it’s far from a pre-theoretic datum that we typically reason through the consequences of our weak beliefs in anything like the way we do our full beliefs.

²⁶ For sympathetic discussion, e.g., Gardiner (2018), Moss (2018), Basu (2019), Bolinger (2020a,b), King & Babic (2020), Jackson & Fritz (2021), and Babic & King (2023). For less sympathetic discussion, see, e.g., Gendler (2011), Moss (2024), Goodsell & Hawthorne (forthcoming).

²⁷ For sympathetic discussion see, e.g., Basu (2018), Basu & Schroeder (2018), and Schroeder (2018). For less sympathetic discussion, see, e.g., Enoch & Spectre (forthcoming).

theses tend to be treated as a package view.²⁸ I'll be considering each on their own terms over the course of this section, though slightly more attention will be paid to questions of moral encroachment than doxastic wronging.

For a sense of the phenomena, consider Gendler's (2011) example of John Hope Franklin, a prominent African American historian. While in attendance at a party at the Cosmos Club in Washington DC—a party that was in fact held to honor his being awarded the Presidential Medal of Freedom—a woman attendee demanded Franklin fetch her coat, as she believed on the basis of his race that he was a member of the club's staff.

There are many who regard this as a paradigm case of doxastic wronging:²⁹ the woman did something morally wrong in virtue of holding the full belief that Franklin was a member of the staff. In particular, even if the woman had not taken any further problematic actions, like demanding that Franklin fetch her coat, the fact that she held the full belief still would have constituted a harm to him. There are also many who regard this as a paradigm case of moral encroachment: the woman's full belief was epistemically unjustified, in large part because of the moral significance of the question of whether Franklin was a member of the staff. Somewhat more precisely: despite the fact that the woman had the kind of statistical evidence that would suffice for a justified full belief on matters that are morally inert, such evidence did not suffice for a justified full belief about Franklin's occupation, since that matter is morally fraught.³⁰

Much of the existing discussion of moral encroachment and doxastic wronging has centered on *full belief*, as in the John Hope Franklin case. But comparably plausible cases of doxastic wronging and moral encroachment can be found for other doxastic attitudes, like weak belief (cf. Moss 2018, Goodsell & Hawthorne forthcoming). Suppose I know that exactly 60% of the people involved in some event are from New York. Suppose Jones is one of those people. Knowing nothing else at all about Jones, it seems reasonable to think (suspect, guess, etc.) that he's from New York, even if I may not reasonably be sure of it. There's where he's most likely to be from, after all. By contrast, if I know that exactly 60% of those people were involved in some crime—say, a Ponzi scheme—it's less clear that epistemology permits me to think (suspect, guess, etc.) that Jones is a fraudster. Moreover, it seems that Jones would be *wronged* by my thinking this of him, at least (or perhaps especially) in those cases where in fact he is among the 40% who committed no fraud whatsoever. Taken at face value, these judgments suggest that I'd need considerably higher confidence (even if not 100% confidence) to appropriately think that Jones is a fraudster. Absent such confidence, I seem to do better remaining agnostic on the matter.

For the purposes of this paper I intend to remain neutral on the question of whether these claims about full and weak belief are correct (i.e., that they can wrong and are subject

²⁸ Indeed, a standard argument for moral encroachment centers on the premise that if one's Φ -ing that p is epistemically impeccable, then one's Φ -ing that p does not constitute a moral wrong. When combined with the assumption that the facts about doxastic wronging do not supervene on the kinds of factors traditional epistemology takes to be relevant to issues of epistemic justification (e.g., one's evidence), one gets a bridge from doxastic wronging to moral encroachment. See, e.g., Basu & Schroeder (2018), Schroeder (2018), King & Babic (2020), Jackson & Fritz (2021).

²⁹ See the citations in footnote 27 above.

³⁰ As Goodsell & Hawthorne (forthcoming, fn. 3) point out, the question of whether the woman's full belief is justified by the standards of traditional epistemology is actually not entirely obvious. We'll return to this issue near the end of this section.

to moral encroachment). What I do not intend to remain neutral on, however, is a recent argument against moral encroachment due to [Goodsell & Hawthorne](#) ([forthcoming](#)). I believe their argument is simply unsound. But I also think the manner in which their argument is unsound is revealing. In particular, I believe the lesson we should take from their argument is that if morality encroaches on full belief, it is by way weak belief. And I think that this in turn raises interesting and difficult questions both about why morality encroaches, and whether and why we ought to care about doxastic wrongdoing.

Simplifying considerably, [Goodsell & Hawthorne](#)'s ([forthcoming](#)) argument against moral encroachment goes as follows. One who accepts moral encroachment for full belief ought to accept it for other doxastic attitudes, like weak belief and suspension of judgment. The intuitive motivations for positing the one look essentially the same as those for positing the other. However, it would be quite surprising if morality encroached on all these attitudes alike without there being some uniform mechanism of encroachment. And this means that certain kinds of explanations for how morality encroaches ought to be rejected, because insufficiently general. For example: since merely weakly believing a proposition does not suffice for fully believing it, an explanation of how morality encroaches on full belief cannot go by way of features that are unique to full belief, like its connections to inference and action. A "common factor" is needed. So, [Goodsell](#) and [Hawthorne](#) suggest, the general mechanism of encroachment must be *credence*, since credence is the clearest common factor to all the doxastic attitudes that are subject to encroachment. Applied to the John Hope Franklin Case, the thought would be that the woman's evidence for the proposition that Franklin is a member of the staff justifies a level of credence in that proposition that is below what is needed to be in a position to permissibly fully believe it. (By contrast, if the woman had comparably strong evidence for a morally neutral proposition—say, the proposition that a certain man across the room is drinking a martini—she might well be justified in assigning that proposition a level of credence that is compatible with full belief.) And in the case of Jones, the thought would be that one may reasonably have credence .6 that Jones is from New York, but must have credence less than .6 that he is a fraudster, even despite one's knowledge of the symmetry of the "background statistics", as it were.

[Goodsell](#) and [Hawthorne](#) then argue that the thesis that morality encroaches on credence leads to absurdity. I won't rehash the full details of their case for this particular claim (see their §6).³¹ But for a sense of the issues, note how bizarre it is to insist that, epistemically speaking, one must have lower than .6 credence that Jones is a fraudster. After all, one knows for certain that exactly 60 of the 100 people at the event of the fraudsters, that Jones is one of those 100 people, and that one has no other information relevant to the question of whether Jones is a fraudster. So the only reasonable credence one can possibly assign to the proposition that Jones is a fraudster is .6. Therefore, morality must not encroach on credence. However, if morality doesn't encroach on credence, then it can't be that what explains why one may justifiably think that Jones is from New York but may not justifiably think that he is a fraudster is some difference in the level of credence one may justifiably assign those propositions. But then what could the mechanism for encroachment

³¹ For further discussion on the issue of whether morality encroaches on credence in particular, see [Moss \(2018\)](#), [King & Babic \(2020\)](#), [Jackson & Fritz \(2021\)](#), and [Quanbeck \(2023\)](#).

be? As they write (p 17):

... given the tight connection between credence and these other doxastic attitudes, it seems impossible for the appropriateness of them to vary without corresponding variation in appropriateness of credences. If the right level of credence for you to have in two propositions is equal, then surely it is equally appropriate for you to doubt/suspect/have little credence in either of them.

So morality must not encroach on weak belief either. But if morality can't encroach on weak belief (since no plausible mechanism for encroachment), then it can't encroach on full belief. So morality does not encroach.

For the purposes of this paper I am happy to endorse just about every part of Goodsell and Hawthorne's argument. In particular, I am happy to endorse: (i) that if morality does not encroach on weak belief, then it doesn't encroach on full belief; (ii) that if morality encroaches on both full and weak belief, it is by a uniform mechanism (i.e., some factor in common to them); and (iii) that morality cannot encroach on credence.

What I deny is their claim that the mechanism for encroachment for weak belief must be credence. I think we have independent reason to reject the claim quoted just above. It is *not* impossible for the appropriateness of doxastic attitudes like *thinking* or *suspecting* to vary without corresponding variation in appropriateness of credences. As we know from the discussion of COGENCY in §2, weak belief does not supervene on credence. Two agents with identical credence functions may permissibly adopt different (weak) doxastic attitudes toward some question $Q^?$, since there are always at least two cogent answers to a question. If one has credence .6 that Jones is from New York, then relative to the question 'Is Jones from New York?' one is permitted either to weakly believe that he is or to have no opinion on the matter. In principle, then, the mere fact that one has .6 credence that Jones is from New York and .6 credence that Jones is a fraudster does *not* entail that one weakly believes that Jones is from New York (relative to the question 'Is Jones from New York?') if and only if one weakly believes that Jones is a fraudster (relative to the question 'Is Jones a fraudster?'). COGENCY allows one to weakly believe the former while failing to weakly believe the latter.

This, I suggest, is what the proponent of moral encroachment should take the mechanism of encroachment to be. Morality encroaches not by affecting one's (permissible) credences, but by ruling out otherwise cogent answers to the question at hand. In particular, one should take the lesson of the cases that motivate moral encroachment to be that COGENCY (repeated here) provides *necessary but insufficient* conditions for permissible weak belief:

COGENCY

It is rationally permissible for p to be the strongest thing S thinks relative to $Q^?$ just in case:

- (1) p is a union of complete answers to $Q^?$.
- (2) If there is a complete answer to $Q^?$, r_1 , such that r_1 doesn't entail p, then there is no other complete answer to $Q^?$, r_2 , such that: r_2 entails p, but S's evidence supports r_1 more than r_2 .

To fill the gap, we might (as a crude first pass) add an extra “Lockean” condition for morally significant propositions, along the lines of:

- (3) If someone would be wronged by S’s thinking that p, then S’s evidential support for p is sufficiently high.

We can be vague about the meaning of both ‘moral significance’ and ‘sufficient support’. All that really matters is the extra degree of freedom, since it allows us to make distinctions in the appropriateness of weak beliefs without requiring distinctions in the appropriateness of credence. In particular, it allows us to say that one may permissibly think that Jones is from New York, since *Jones is from New York* is a cogent answer to the question ‘Is Jones from New York’ and is not morally significant. But it also allows us to say that one may not permissibly think that Jones is a fraudster, since even though *Jones is a fraudster* is a cogent answer to the question ‘Is Jones a fraudster?’, it’s a morally significant proposition, and one’s evidential support for this proposition may well be insufficiently high. Perhaps one would need the kind of evidence that would license (say) .9 credence that Jones is a fraudster before one could permissibly weakly believe it. In the absence of such evidence, one is only permitted to think the trivial thing: that Jones either is or is not a fraudster.

There is also the issue of accounting for moral encroachment on *suspension* (von Klemperer 2023, Goodsell & Hawthorne forthcoming). Just as it seems possible to wrong Jones by coming to think he is a fraudster when one’s evidence for this is insufficiently strong, so too does it seem possible to wrong Jones by *failing* to come to think that he is *not* a fraudster when one’s evidence for this is sufficiently strong. For example: if we modify the Jones case so that one knows that there is only a 40% chance that Jones is fraudster, then perhaps there is only one permissible weak doxastic attitude to take to the question ‘Is Jones a fraudster?’, which is the weak belief that he is not. To account for this aspect of moral encroachment, we’ll have to add a fourth and final condition along the lines of:

- (4) For any (complete or partial answer) to Q? q such that p doesn’t entail q, if someone would be wronged by S’s failing to think that q, then S’s credence that q is sufficiently low.

So, to the extent we are convinced that morality encroaches on weak belief, we should reject COGENCY as an account of the necessary and sufficient conditions for weak belief, and instead think that the true account is given by something like the conjunction of COGENCY with (3) and (4). This account has the desirable feature of giving a uniform explanation of the mechanism of moral encroachment across all the doxastic attitudes that entail thinking. In particular, full belief entails weak belief, so in any case where weak belief is impermissible, so is the corresponding full belief. *Mutatis mutandis* for many other qualitative weak doxastic attitudes, like *being of the opinion that*, *being somewhat confident in*, and so on (cf. Hawthorne *et al.* 2016, Goodman & Holguín 2022, Goodman 2024). This suggests that Goodsell and Hawthorne are wrong in their contention that if morality encroaches at all, it encroaches on credence. One of the things that weak belief tells us about full belief is that moral encroachment is a matter of constricting the space of weakly believable answers to the question at hand.

Another nice feature of this account of the mechanism of moral encroachment is that it makes it more vivid how certain kinds of doxastic attitudes could *wrong*. This is due to the voluntaristic nature of weak belief. Suppose you and I both have credence .6 that Jones is from New York, and that I think that Jones is from New York. Does it follow that you also think Jones is from New York? No. Knowing that you and I have the same credence that Jones is from New York and that *I* think that Jones is from New York isn't enough for me to know whether *you* think that Jones is from New York. It might be enough for me to know that it's not true that you think that he's *not* from New York; but it's not enough for me to rule out the possibility that you're agnostic on the matter. To know that you're not agnostic, I need to know how you've made up your mind on the question. When it comes to morally neutral matters, like the question of whether Jones is from New York, it's not clear there's any reason to make one's mind up one way rather than the other—all the cogent answers are permissible. But for morally significant matters, like the question of whether Jones is a fraudster, not all of the cogent answers are permissible. So if you make up your mind to be of the opinion that Jones is a fraudster when you easily could have made up your mind to have no opinion on the matter, there's a real sense in which you're *responsible* for what you weakly believe on the matter. Nothing forced you into that opinion, yet you adopted it anyway. This gives Jones clearer grounds for resentment.³²

Whether this extends to *full beliefs* that putatively wrong is a more subtle matter. It does not follow from the fact that (i) full belief entails weak belief and that (ii) we can and often do form weak beliefs voluntaristically that (iii) we can and often do form full beliefs voluntaristically. Indeed, on the kind of view defended in [Holguín \(2022\)](#), full belief is non-voluntaristic, and weak belief is voluntaristic in *only* those cases where one lacks a full belief as to the true answer to the question at hand. On this view, if you fully believe that *p*, then relative to any question *Q*? to which *p* is an answer, you can't help but weakly believe relative to *Q*? that *p*.

But this view is not mandatory. In fact, I think there is reason to believe that full beliefs grounded in statistical evidence are sometimes formed voluntaristically. To give one example: suppose I grab a coin at random from a bucket filled with fair coins and double-headed coins. Without inspecting it, I flip it repeatedly, and see that it always comes up heads. On pains of skepticism, there is some least number *n* such that: after seeing *n* consecutive heads flips, I'll be in a position to rationally be sure that the coin is double-headed ([Dorr et al. 2014](#), [Goodman & Salow 2023](#), [forthcoming](#)). But does it follow that I *will* be sure that the coin is double-headed after seeing *n* heads? It's not obvious why it would. Perhaps I'll

³² This is not to say that we should only accept the claim that (weak) beliefs wrong if we're also willing to accept the claim that our (weak) beliefs are under our direct control. It is perfectly coherent to deny this. But the issues here are delicate. On the one hand, many epistemologists have made their peace with the idea that there can be norms on doxastic attitude formation even when the relevant attitudes are formed involuntarily (see [Chignell 2018](#), §3.4 and the citations therein). On the other hand, though, there seems to me an intuitive distinction between being such that *one violates a norm* by being in a certain state and being such that *one wrongs someone* by being in a certain state. In particular, the fact that it wasn't under my control whether I was in such a state doesn't do as much to undermine the feeling that my being in that state constitutes a violation of a certain norm, whereas it does go some way to undermining the feeling that I wronged someone by being in that state. The latter status seems to be more closely tied to questions of responsibility, being able to do otherwise, and the like. So, to the extent I'd want to endorse the thesis of doxastic wronging, it strikes me as desirable to explain the phenomenon in terms of attitudes over which we sometimes exert direct control. The weak belief-theoretic account is in pretty good shape in this respect.

only become sure after $n + 1$ heads, or $n + 10$, or what have you. And phenomenologically speaking, it feels like a matter of making up my mind when to start being sure.

Supposing we take these appearances at face value, we get a nice model of the John Hope Franklin case. The woman's evidence for the proposition that Franklin was a member of the staff was merely statistical in nature, grounded in her knowledge of the general distribution of club members and staff by racial identity. Many have the intuition that the woman did something wrong in fully believing Franklin was a member of the staff on this basis. But as [Goodsell & Hawthorne \(forthcoming, §1\)](#) point out, it's much less clear whether these sorts of intuitions persist in cases where the woman's full belief is grounded in non-statistical evidence. For example: suppose the woman fully believed that Franklin was a member of the staff not because of his racial identity, but because she mistook him for his identical twin, who she knew for sure was a member of the staff. Or imagine that the woman fully believed Franklin was a member of the staff because she was told this by a bunch of otherwise reliable informants. When we reimagine the etiology of the woman's full belief in these ways, it becomes considerably less clear that she wrongs Franklin in holding it. This contrast is nicely explained by the combination of (i) the view that when doxastic attitudes wrong, it's at least in part because they're formed voluntaristically; and (ii) the view that full beliefs grounded in statistical evidence are (often enough) formed voluntaristically.

To summarize, then: one of the things weak belief tells us about full belief is that *if* there is such a thing as moral encroachment, then moral encroachment on full belief plausibly goes by way of moral encroachment on weak belief, which itself need not go by way of moral encroachment on credence. This means that Goodsell and Hawthorne's general argument against moral encroachment is unsound. It also gives an intuitively satisfying account of the claim that doxastic states themselves can wrong, at least for those doxastic states that are under our voluntaristic control. Weak belief is typically like this; perhaps sometimes full belief is too.

5 Norm primitivism

I want to conclude by briefly turn to some questions about the meta-normative status of the norms associated with toggling and morality.

I'll start with the latter. Should the proponent of moral encroachment be satisfied with the previous section's answer to Goodsell and Hawthorne's challenge? I myself find the answer a bit unclear. Consider again the case of Jones, the possible New York-based fraudster. It is .6 likely on one's evidence that Jones is from New York. And it is .6 likely on one's evidence that Jones is a fraudster. But the latter proposition is unlike the former in that it is morally significant. So let us suppose, as before, that one who weakly believes that Jones is a fraudster wrongs Jones (at least when Jones is in fact innocent) and weakly believes in a manner that is epistemically impermissible.

Now imagine two further characters who are in this evidential situation: Smith and Brown. Both characters know that their evidence makes it .6 likely that Jones is a fraudster.

But only Smith chooses to weakly believe it; Brown instead merely weakly believes that either Jones is or is not a fraudster. By hypothesis, then, Smith but not Brown is in a doxastic state that wrongs Jones, and Smith but not Brown is in a doxastic state that is epistemically impermissible. The question is whether this difference between Smith and Brown should make a difference.

Consider Jones' perspective on the matter (supposing he's been informed of all this). He knows that both Smith and Brown recognize that there is a reasonable chance he is not a fraudster, but also that they (fully) believe it's more likely than not that he is. As such, Jones knows that neither will take for granted that he's a fraudster—at least not for the purposes of theoretical or practical reasoning. But he also knows that each would bet on his being fraudster (at least at even odds) if given the opportunity. More generally, Jones can be confident that there will be essentially no difference between Smith and Brown when it comes to their behavioral dispositions with respect to the proposition that he is a fraudster. This is because Jones can be confident that most all of the interesting behavioral facts concerning Smith and Brown are determined not by what they weakly believe, but by what their credences are. In fact, as far as Jones is concerned, the only real behavioral difference he can expect between Smith and Brown is the following one: that Smith, but not Brown, will answer in the affirmative if asked whether he thinks Jones is a fraudster. But is *that* really the kind of difference Jones should care about? Should he really be entertaining thoughts like, 'They're both 60% confident I'm a fraudster, but at least only Smith goes as far as to think I am one'? This would seem to unduly fetishize weak belief.

This creates a subtle dialectical situation. On the one hand, the impression that morality encroaches on weak belief (and that weak beliefs can wrong) is strong: there really does seem to be something less than fully appropriate about coming to think (say) that a certain person is a scoundrel just because one knows that some group contains 60% scoundrels and 40% saints and that that person is a member of the group. To the extent that this pre-theoretic intuition is taken as a datum that ought to be vindicated by our best theory of weak belief, there is pressure both to say that weak beliefs of this sort wrong, and perhaps also that they are epistemically impermissible. So if we're going to yield to this pressure, we may just have to learn to live with the idea that there's nothing all that deep to say about *why* such weak beliefs wrong, or *why* there are epistemic norms against holding them. The existence of such norms may well just be a brute fact of sorts.

On the other hand, one could also take these considerations to simply undermine the thesis that weak beliefs like this wrong and/or are epistemically impermissible (cf. [Jackson & Fritz 2021](#)). If *this* is as deep as the explanations of doxastic wronging and moral encroachment can go, then why even believe in such things? Perhaps it would be better just to get used to the idea that the relevant pre-theoretic intuitions are mistaken—that one who justifiably has .6 credence that Jones is a fraudster can in fact justifiably think that he is a fraudster, and that Jones need not be wronged by one's happening to think this of him.

The issue is that this kind of reaction can be difficult to quarantine. If we're willing to embrace an error-theory about moral judgments in the case of weak belief, there is pressure to say the same about the corresponding judgments in the case of other doxastic attitudes

as well. After all, if we readily hallucinate moral and epistemic asymmetries between the weak belief that Jones is from New York and the weak belief that Jones is a fraudster, then why should we be so confident that the impression of the moral asymmetries in the John Hope Franklin case are veridical? Perhaps we just need to learn to distrust the impulse to posit a moral and/or epistemic asymmetry between the status of the full belief that Franklin is a member of the staff and (say) the full belief that Franklin is from Washington, D.C., at least when both are formed on broadly statistical grounds. Otherwise we'd need some story that explains why we're more likely to be mistaken about the moral and epistemic status of weak beliefs than those of full beliefs. This is not to say that no story of this sort can be given. But I don't see an easy path to it.³³

There is a similar dialectical situation with respect to the norms on toggling. It is difficult to see what, if anything, goes badly for agents who toggle their weak beliefs. In particular, if an agent who toggles freely between weakly believing that p , weakly believing that $\neg p$, and weakly believing merely that $p \vee \neg p$ won't be any worse off than one who keeps steadfast in (say) their weak belief that p , then it's unclear what good a norm like THE LAW OF WEAK BELIEF TOGGLING is even doing for us. Yet the impression that there is something wrong with weak belief toggling is hard to avoid. This suggests that either we should be looking for an error-theory of the intuitions that suggest that weak belief toggling is permissible, or we should get used to the idea that, from a "design" perspective, there's no good reason for a norm like THE LAW OF WEAK BELIEF TOGGLING. It's just there—for better or for worse.

But whichever way we go in the case of THE LAW OF WEAK BELIEF TOGGLING, there is pressure to go the same for THE FUNDAMENTAL LAW OF TOGGLING. This pressure is not insurmountable, of course. There are enough asymmetries between the different attitudes to which THE FUNDAMENTAL LAW OF TOGGLING applies (e.g., preference, perhaps also full belief) to take seriously the hypothesis that only one of them is genuinely hostile to toggling, or the hypothesis that the explanations of the impermissibility of toggling in the two cases are fundamentally distinct. But I for one favor unified theories of what appear to be unified phenomena.

³³ One point in favor of positing a weak-versus-full belief asymmetry with respect to moral encroachment is that only full belief seems subject to the phenomenon of *pragmatic* encroachment (cf. Ross & Schroeder 2014). If I've got a lot riding on whether the bank is open at a certain time, then it's not clear I'm justified in being sure that the bank is open at that time when my only basis for that full belief is my memory of the bank's hours. Intuitively I ought to double-check to make sure that I'm remembering correctly, that the hours haven't changed, and so on. By contrast, if the question of the bank's hours are a bit of idle trivia, such double-checking isn't obviously required for rational surety. Interestingly, however, there doesn't seem to be any such contrast in the case of weak belief: it is as reasonable to think that the bank is open when I've got a lot riding on it as when it's a bit of idle trivia (at least when my only evidence for my weak belief is my apparent memory of the bank's hours). Indeed, in the high-stakes case, a speech like 'I *think* the bank is open at this time, but I should double-check to be sure' is perfectly natural. This asymmetry between full and weak belief with respect to pragmatic encroachment is mostly unsurprising: only the former has the relevant sorts of connections to practical deliberation. Such an asymmetry may be called upon in defense of the claim that morality encroaches on full belief but not on weak belief. However, as Goodsell & Hawthorne (forthcoming, §§2-3) emphasize, there are good reasons to doubt that moral encroachment on full belief can be motivated via the connections between full belief and action.

References

- Ayer, A. J. 2006. The Problem of Knowledge. *Pages 159–163 of: Honderich, Ted (ed), Ayer Writings in Philosophy: A Palgrave Macmillan Archive Collection*. Palgrave Macmillan.
- Babic, Boris, & King, Zoë Johnson. 2023. Moral Encroachment Under Moral Uncertainty. *Philosophers' Imprint*, **23**(n/a).
- Basu, Rima. 2018. The Wrongs of Racist Beliefs. *Philosophical Studies*, **176**(9), 2497–2515.
- Basu, Rima. 2019. Radical Moral Encroachment: The Moral Stakes of Racist Beliefs. *Philosophical Issues*, **29**(1), 9–23.
- Basu, Rima, & Schroeder, Mark. 2018. Doxastic Wronging. *Pages 181–205 of: Kim, Brian, & McGrath, Matthew (eds), Pragmatic Encroachment in Epistemology*. Routledge.
- Beddor, Bob. 2020. New Work For Certainty. *Philosophers' Imprint*, **20**(8), 1–25.
- Bolinger, Renée Jorgensen. 2020a. The Rational Impermissibility of Accepting (Some) Racial Generalizations. *Synthese*, **197**(6), 2415–2431.
- Bolinger, Renée Jorgensen. 2020b. Varieties of Moral Encroachment. *Philosophical Perspectives*, **34**(1), 5–26.
- Buchak, Lara. 2014. Belief, Credence, and Norms. *Philosophical Studies*, **169**(2), 285–311.
- Chang, Ruth. 1997. *Incommensurability, Incomparability, and Practical Reason*. Cambridge, MA, USA: Harvard.
- Chignell, Andrew. 2018. The Ethics of Belief. *In: Zalta, Edward N. (ed), The Stanford Encyclopedia of Philosophy*, Spring 2018 edn. Metaphysics Research Lab, Stanford University.
- Chisholm, Roderick M. 1957. *Perceiving: A Philosophical Study*. Cornell University Press.
- Clarke, Roger. 2024. Strong Belief is Ordinary. *Episteme*, **21**(3), 773–793.
- Dandelet, Sophia. 2023. Arbitrary Switching and Concern for Truth. *Synthese*, **202**(4), 1–21.
- Dorr, Cian, Goodman, Jeremy, & Hawthorne, John. 2014. Knowing Against the Odds. *Philosophical Studies*, **170**(2), 277–287.
- Dorst, Kevin. 2019. Lockeans Maximize Expected Accuracy. *Mind*, **128**(509), 175–211.
- Dorst, Kevin, & Mandelkern, Matthew. 2023. Good Guesses. *Philosophy and Phenomenological Research*, **105**(3), 581–618.
- Douven, Igor. 2009. Uniqueness Revisited. *American Philosophical Quarterly*, **46**(4), 347–361.
- Drucker, Daniel. 2020. The Attitudes We Can Have. *Philosophical Review*, **129**(4), 591–642.
- Enoch, David, & Spectre, Levi. forthcoming. There is No Such Thing as Doxastic Wrongdoing. *Philosophical Perspectives*.
- Gardiner, Georgi. 2018. Evidentialism and Moral Encroachment. *In: Kevin, McCain (ed), Believing in Accordance with the Evidence: New Essays on Evidentialism*. Springer Verlag.
- Gendler, Tamar Szabó. 2011. On the Epistemic Costs of Implicit Bias. *Philosophical Studies*, **156**(1), 33–63.
- Genin, Konstantin. 2019. Full & Partial Belief. *Pages 437–498 of: Pettigrew, Richard, & Weisberg, Jonathan (eds), The Open Handbook of Formal Epistemology*. PhilPapers Foundation.
- Gettier, Edmund. 1963. Is Justified True Belief Knowledge? *Analysis*, **23**(6), 121–123.
- Goodman, Jeremy. 2023. The Myth of Full Belief. *Philosophical Perspectives*, **37**(1), 164–171.
- Goodman, Jeremy. 2024. Degrees of confidence are not subjective probabilities. *Proceedings of Sinn und Bedeutung*, **28**(Dec.), 329–344.
- Goodman, Jeremy, & Holguín, Ben. 2022. Thinking and Being Sure. *Philosophy and Phenomenological Research*, **106**(3), 634–654.
- Goodman, Jeremy, & Salow, Bernhard. 2023. Epistemology Normalized. *Philosophical Review*, **132**(1), 89–145.
- Goodman, Jeremy, & Salow, Bernhard. forthcoming. Belief Revision Normalized. *Journal of Philosophical Logic*, 1–49.
- Goodsell, Zachary, & Hawthorne, John. forthcoming. Morality Does Not Encroach. *In: Comesana, Juan, & McGrath, Matthew (eds), Knowledge and Rationality: Essays in Honor of Stewart Cohen*. Routledge.
- Hawthorne, John, Rothschild, Daniel, & Spectre, Levi. 2016. Belief is Weak. *Philosophical Studies*, **173**(5), 1393–1404.

- Hoek, Daniel. 2022. Questions in Action. *Journal of Philosophy*, **119**(3), 113–143.
- Hoek, Daniel. 2025. Minimal Rationality and the Web of Questions. In: van Elswyk, Peter, Kindermann, Dirk, Kirk-Giannini, Cameron Domenico, & Egan, Andy (eds), *Unstructured Content*. Oxford University Press.
- Holguín, Ben. 2022. Thinking, Guessing, and Believing. *Philosophers' Imprint*, **22**, 1–34.
- Jackson, Elizabeth. 2021. A Defense of Intrapersonal Belief Permissivism. *Episteme*, **18**(2), 313–327.
- Jackson, Elizabeth, & Fritz, James. 2021. Belief, Credence, and Moral Encroachment. *Synthese*, **199**(1-2), 1387–1408.
- Kelly, Thomas. 2013. Evidence Can Be Permissive. *Page 298 of: Steup, Matthias, & Turri, John* (eds), *Contemporary Debates in Epistemology*. Blackwell.
- King, Zoe Johnson, & Babic, Boris. 2020. Moral Obligation and Epistemic Risk. *Oxford Studies in Normative Ethics*, **10**, 81–105.
- Kopec, Matthew, & Titelbaum, Michael G. 2016. The Uniqueness Thesis. *Philosophy Compass*, **11**(4), 189–200.
- Mandelkern, Matthew, & Dorst, Kevin. 2022. Assertion is Weak. *Philosophers' Imprint*, **22**(n/a).
- Moss, Sarah. 2014. Credal Dilemmas. *Noûs*, **48**(3), 665–683.
- Moss, Sarah. 2018. Moral Encroachment. *Proceedings of the Aristotelian Society*, **118**(2), 177–205.
- Moss, Sarah. 2019. Full Belief and Loose Speech. *Philosophical Review*, **128**(3), 255–291.
- Moss, Sarah. 2024. How to Do Without Encroachment. *Mind*, **133**(532), 931–971.
- Nagel, Jennifer. 2021. The Psychological Dimension of the Lottery Paradox. In: Douven, Igor (ed), *Lotteries, Knowledge, and Rational Belief: Essays on the Lottery Paradox*. Cambridge University Press.
- Pearson, Joshua Edward. forthcoming. A Puzzle About Weak Belief. *Analysis*.
- Quanbeck, Z. 2023. Belief, Blame, and Inquiry: A Defense of Doxastic Wronging. *Philosophical Studies*, **180**(10-11), 2955–2975.
- Quillien, Tadeq, Bramley, Neil R., & Lucas, Christopher G. 2023. *Guesses as Compressed Probability Distributions*.
- Roeber, Blake. 2020. Permissive Situations and Direct Doxastic Control. *Philosophy and Phenomenological Research*, **101**(2), 415–431.
- Ross, Jacob, & Schroeder, Mark. 2014. Belief, Credence, and Pragmatic Encroachment. *Philosophy and Phenomenological Research*, **88**(2), 259–288.
- Rothschild, Daniel. 2020. What It Takes to Believe. *Philosophical Studies*, **177**(5), 1345–1362.
- Schoenfield, Miriam. 2014. Permission to Believe: Why Permissivism is True and What It Tells Us About Irrelevant Influences on Belief. *Noûs*, **48**(2), 193–218.
- Schroeder, Mark. 2018. When Beliefs Wrong. *Philosophical Topics*, **46**(1), 115–127.
- Skipper, Mattias. 2023. Good Guesses as Accuracy-Specificity Tradeoffs. *Philosophical Studies*, **180**(7), 2025–2050.
- Staffel, Julia. 2016. Beliefs, Buses and Lotteries: Why Rational Belief Can't Be Stably High Credence. *Philosophical Studies*, **173**(7), 1721–1734.
- Teague, Richard. 2024. A Puzzle About Guessing and Inquiry. *Analysis*, **84**(2), 327–336.
- Titelbaum, Michael. 2015. Continuing On. *Canadian Journal of Philosophy*, **45**(5), 670–691.
- von Klemperer, Caroline. 2023. Moral Encroachment, Symmetry, and Believing Against the Evidence. *Philosophical Studies*.
- Wedgwood, Ralph. 2012. Outright Belief. *Dialectica*, **66**(3), 309–329.
- Weisberg, Jonathan. 2020. Could've Thought Otherwise. *Philosophers' Imprint*, **20**(12).
- White, Roger. 2005. Epistemic Permissiveness. *Philosophical Perspectives*, **19**(1), 445–459.
- White, Roger. 2013. Evidence Cannot Be Permissive. *Page 312 of: Steup, Matthias, & Turri, John* (eds), *Contemporary Debates in Epistemology*. Blackwell.
- Williamson, Timothy. 2000. *Knowledge and Its Limits*. Oxford University Press.
- Williamson, Timothy. forthcoming. Knowledge, Credence, and the Strength of Belief. In: Flowerree, Amy, & Reed, Baron (eds), *Towards an Expansive Epistemology: Norms, Action, and the Social World*. Routledge.

- Woodard, Elise. 2020. A Puzzle About Fickleness. *Noûs*, **56**(2), 323–342.
- Yalcin, Seth. 2018. Belief as Question-Sensitive. *Philosophy and Phenomenological Research*, **97**(1), 23–47.