

The Will in Belief

Rachel Fraser

January 3, 2022

Abstract

It's bad to hang on to a belief in the face of good evidence that it's false. That makes you *dogmatic*. Enter the classic dogmatism puzzle. The standard response? Go for *defeat*: new evidence can destroy old knowledge. But there are variants on the classic dogmatism puzzle which defeat can't help with. Worse still, these puzzles threaten to undermine the standard defeatist solution to the classic puzzle. Call these the *revenge* puzzles. Two are already noted in the literature. In this paper, I introduce a third, particularly severe revenge puzzle. I then present a unified solution to the revenge puzzles. The solution motivates a boldly revisionary conception of belief: one on which belief involves the will.

But to eliminate the will
altogether...what would that
mean but to *castrate* the
intellect?

Friedrich Nietzsche, *Genealogy of*

Morals

1 Introduction

It's bad to hang on to a belief in the face of good evidence that it's false. That makes you *dogmatic*. Enter the dogmatism puzzle (Kripke 2011; Harman 1986):

You should ignore misleading evidence. But if you know that p , you know that evidence against p is misleading. So if you know that p , you should ignore evidence against p . But that's dogmatic.¹

¹This is a little fast, but easily tightened up.

Call this the *classic* dogmatism puzzle. The standard response? Go for *defeat*: new evidence can destroy old knowledge. But there are related puzzles which are not solved by going for defeat. Call these the *revenge* puzzles. Two are already noted in the literature. In this paper, I introduce a third.

Happily, the revenge puzzles are solved by a package of views I call **Volitionism**.² Volitionism contends that belief involves the will: an agent judges that p only if she makes a *resolution*, a resolution not to (further) inquire into whether p .³ Volitionism is thus a radical view. A long tradition in philosophy holds that the faculties of will and judgement – by which, here and throughout, I mean *theoretical*, rather than *practical* judgement – are quite distinct: to make a judgement is never enough to make a decision. Volitionism, by contrast, holds that our capacity to judge depends on our ability to will.⁴ I argue for Volitionism by showing that it offers a clean, unified solution to the dogmatism puzzles.

Here is the plan. In §2, I (i) give a careful account of the classic dogmatism puzzle, (ii) introduce the two extant revenge puzzles, and (iii) introduce a third, novel revenge puzzle. We are left with a series of desiderata for an adequate anti-dogmatism. In §3, I discuss non-volitionist responses to the extant revenge puzzles, and argue that they fail. In §4, I consider non-volitionist responses to the novel revenge puzzle, and argue that they fail, too. In §5 I introduce and motivate Volitionism. In §5 section I show that Volitionism can satisfy two of the relevant desiderata. In §6 I show that it satisfies the third. The conclusion? Anti-dogmatists should be volitionists.

²Volitionism does not entail and is not entailed by doxastic voluntarism.

³What do I mean by ‘belief’ and ‘judgement’? Belief first. I’m a traditionalist: belief that p is the state expressed by an assertion that p , the state required for one to licitly reason from p , and the non-factive residue of knowledge. If you don’t want to call this state ‘belief’ (Dorst 2017; Hawthorne, Rothschild, and Spectre 2016) I won’t fight; you can *call* it whatever you want. Now judgement: a judgement that p is an act of mind which outputs a belief that p . It may or may not be the result of deliberation.

⁴The picture when it comes to practical judgement is more dappled (Fix 2018.)

2 The Classic Puzzle

Here is an appealing claim:

Ignore. Necessarily, if A knows that evidence e is misleading with respect to p , A should ignore e when assigning a doxastic attitude to p .

Ignore says that if A knows some evidence to be misleading with respect to p , A 's assignment of doxastic attitudes – beliefs, degrees of confidence, and so on – to p should proceed as though A lacks e . It is a key component of the classic dogmatism puzzle:

Classic Dogmatism. Suppose A knows that p . She then receives some new evidence e , which she recognises as evidence against p . Because A knows that p , she knows that e is evidence against a truth, and so is misleading evidence, at least with respect to p . So by **Ignore**, A 's degree of confidence that p should be as if she had never acquired e .⁵

Let's make this vivid:

Coin. Celia visits a coin factory which specialises in making fair coins: all the coins it makes are perfectly fair. Celia knows this — that's why she wants to visit the factory. Before leaving, Celia is given one of the factory's products as a souvenir. She judges, and thereby comes to know the following: *my coin is perfectly fair*. Celia then performs an experiment: she flips the coin a million times. Result: it comes up heads every time.

⁵Note that to get the classic puzzle going, we need to grant something in the ballpark of *closure*: that knowledge is closed under known implication. Yablo has recently suggested that we should prefer the weaker *immanent closure*: knowledge is (only) closed under *subject-matter preserving* known implication (Yablo 2017). This weakened principle predicts that inferences like the following fail to be knowledge preserving: ' p . So any evidence against p is misleading'. Such inferences' conclusions have a different subject matter to their premises. However, immanent closure *does* allow that inferences like those above are knowledge preserving: someone who infers that e is evidence against a truth from their knowledge that p and that e is evidence against p does not 'change the subject'.

But Celia's coin *is* fair – the experiment was a fluke. So when Celia gets the results of her experiment, she receives some misleading evidence that her coin is biased. But if Celia knows that her coin is fair at t , she may reason as follows:

My coin is fair. So evidence suggesting that my coin is not fair is misleading, at least with respect to the claim that my coin is fair. But if evidence is misleading with respect to p , my degree of confidence that for p should ignore that evidence. So it is just as probable on my my evidence now that my coin is fair as it was before.⁶

Clearly, this is a bad result. Enter **Defeat**:

Defeat. Suppose at t_1 that A knows that p . At t_2 she acquires new evidence against p . So long as that evidence is sufficiently strong, A no longer knows that p .

The standard solution endorses **Defeat**. At t , Celia acquires strong evidence that her coin is biased, so she no longer knows that her coin is fair. But if she does not know that her coin is fair, she cannot reason from the claim that her coin is fair. So the above reasoning is illicit. Good: this is the right result.

2.1 The Dogmatist's Revenge

But things are not so simple. Kripke and Lasonen-Aarnio point out revenge puzzles: **Dogmatic Intentions** and **Weak Dogmatism**, respectively. Neither is solved by **Defeat**. Time for a closer look.

⁶From here on, I'll often substitute 'probable' for 'probable on my evidence'.

2.2 Dogmatic Intentions

Dogmatic Intentions. Suppose A knows that p at t_1 . She realises that she might get some evidence against p in the future. But she knows that any such evidence will be misleading. So she forms an intention: ignore any evidence she gets against p (Kripke 2011; Beddor 2019).

If A knows that future evidence will be misleading, it should be fine for her to intend to ignore it. And if A knows p , she knows evidence against p is misleading. So it should be fine for A to intend to ignore future evidence against any p she currently knows. But forming such an intention looks dogmatic.⁷

Let's make this vivid. Suppose Celia knows her souvenir coin is fair. On her walk home from the factory, she idly imagines a scenario in which she gets evidence that her coin is biased – perhaps that her coin comes up heads a million times in a row. She resolves to maintain her high confidence that her coin is fair regardless of whether she receives such evidence. That seems dogmatic. But here, **Defeat** is no help: when Celia forms her intention, her knowledge that p has not yet been defeated.⁸ We have our first desideratum:

The Intention Desideratum Solve the puzzle of dogmatic intentions.

2.3 Weak Dogmatism

Suppose classic dogmatism draws you to **Defeat**. Then you probably think that there's an evidential threshold for knowledge, and that it's (sometimes) below 1.⁹ Then an agent A might receive evidence which makes p less likely without

⁷Though perhaps not in every case (Kripke 2011; Fantl 2018). I discuss such exceptions at length in the final section.

⁸One might try: you can intend to ignore evidence against p only if you don't know that you won't get any future evidence against p . But you only know that p if you know you won't get any future evidence against p (Malcolm 1952). This, though, generates scepticism (Unger 1978).

⁹Argument: **Defeat** is awkward for infallibilists (Baker-Hytch and Benton 2015; Lasonen-Aarnio 2014).

defeating A 's knowledge that p (Lasonen-Aarnio 2014). More vividly:

Trial After receiving her fair coin, judging, and thereby coming to know that her coin is fair, Celia decides to flip her coin a few times. She flips it 20 times. It comes up heads 20 times.

This sequence is clearly evidence against the coin being fair: such a sequence is much more likely given a biased than a fair coin. But it doesn't seem strong enough to defeat Celia's knowledge that her coin is fair.¹⁰ So she knows the coin is fair even after the twenty flips. So she can reason thus:

My coin is fair. So evidence suggesting that my coin is not fair is misleading, at least with respect to the claim that my coin is fair. But my probability for p should ignore misleading evidence. So the probability that my coin is fair is just as high as before.

This is another bad result that **Defeat** can't help with. *Ex hypothesi*, Celia's evidence does not defeat her knowledge, so she can reason from the claim that her coin is fair. But if Celia reasons as above, she's dogmatic (Lasonen-Aarnio 2014). We have our second desideratum:

The Weak Desideratum Solve the weak dogmatism puzzle.

Note that both weak dogmatism and dogmatic intentions really are *revenge* – rather than merely residual – puzzles. Why? Because they undermine the motivation for **Defeat**:

[A]rguing for the defeasibility of knowledge from the absurdity of the dogmatist conclusion [in the classic dogmatism puzzle] is inherently problematic, for...[t]he defeasibility of knowledge alone does not shield one from a problematic dogmatism. To avoid that kind of

¹⁰If you disagree, just work with a shorter sequence of heads.

dogmatism, one must deny [Ignore]. But of course, denying [Ignore] alone would have blocked the reasoning in the original dogmatism puzzle (Lasonen-Aarnio 2014).

In other words: these new puzzles pressure anti-dogmatists to reject **Ignore** – the weak puzzle because it directly relies on **Ignore**; the intentions puzzle because, if **Ignore** is false, it's implausible that agents should intend to ignore evidence they know to be misleading. But once we drop **Ignore**, we don't need **Defeat** to avoid dogmatism.

2.4 A New Revenge Puzzle

Time for a new revenge puzzle. Suppose that A has a body of evidence e . e is made up of two subsets of evidence g and m . A realises that g strongly supports p , whereas m is very weak evidence against p . A considers e carefully, taking both the evidence in favour of p and the evidence against p into account. After careful weighing of the competing considerations, A judges, and comes to know that p . A then reasons as follows:

In fact p . So all the evidence m that I just considered against p was misleading. But if m is misleading with respect to p , my probability for p should ignore m . So p is more probable than I initially thought.

This look like terrible reasoning. But it is difficult to see where A has gone wrong. She does know that p : *ex hypothesisi*, A 's evidence, though mixed, does support knowledge that p . The only way to avoid this possibility involves infallibilism, which, for familiar reasons, risks robbing us of the standard solution to the classic problem (Baker-Hytch and Benton 2015; Lasonen-Aarnio 2014). So A does know that m is misleading evidence. And it is hard to see how A might count as having a credence for p that ignores m unless she raises her

degree of confidence that p . After all, at present, A 's degree of confidence that p is informed by m . Had she not been exposed to m , but only to g , she would have assigned a higher probability to p in the first place.

This is the **Dogmatic Bootstrapping** puzzle. The **Dogmatic Bootstrapping** puzzle is significantly more severe than either of the revenge problems studied thus far, because even those willing to countenance the indefeasibility of knowledge should be reluctant to admit this bootstrappy reasoning as licit. Biting the dogmatic bullet *might* be an option with the first two revenge puzzles. It's not an option here. We have our third desideratum:

The Bootstrapping Desideratum Solve the puzzle of dogmatic bootstrapping.

3 Against the Revenge Puzzles

In this section I examine some extant responses to the revenge puzzles. They don't work.

3.1 Responding to Weak Dogmatism

According to Lasonen-Aarnio, our options for responding to **Weak Dogmatism** look like this:

Option 1. Reject Ignore.

Option 2. Accept weak dogmatism: Celia is rational to reason that her probability should not change.

Lasonen-Aarnio admits that this is a nasty choice. It is not appealing to give up **Ignore**. As Lasonen-Aarnio points out, locutions like 'e is misleading

evidence against p , but it should make me less confident in p ' have a Moore-paradoxical feel. Any adequate solution to the dogmatism puzzles, then, must retain **Ignore**.

This leaves Option 2. Option 2 is also very bad. Suppose Celia finds herself at the mercy of a mercurial psychopath, who says he will kill her if the coin he flips comes up tails, and give her a million pounds if the coin comes up heads. She can choose between her coin – which, recall, has come up heads twenty times in a row – or another randomly selected coin. Celia should obviously prefer her coin. But the weak dogmatist cannot accommodate this: Celia should be just as confident of each coin that it is fair. Upshot: neither option is palatable.

3.2 Responding to Dogmatic Intentions

Beddor suggests a solution to **Dogmatic-Intentions**. It turns on:

Intention. It is irrational for A to intend at t_1 to $[\phi$ at $t_2]$ if, at t_1 , A is in a position to know that it will be irrational for A to $[\phi$ at $t_2]$.

Beddor motivates **Intention** by appealing to Kavka's toxin puzzle (Kavka 1983). You are offered a deal by an eccentric billionaire, Elon. If you drink his vial of toxin, you will be unwell for a day, but suffer no long term ill-effects. The toxin will be delivered to you at 9am tomorrow. Elon will deposit 1 million pounds in your bank account at 1am tomorrow just so long as you intend, at midnight tonight, to drink the toxin at 9am tomorrow. Judgement: you cannot, at midnight, rationally intend to drink the toxin tomorrow morning. Why not?

Answer: you can foresee at midnight that by 9am, the money will either be in your bank account or not. Nothing you do at 9am will determine whether you get the money or not. If you don't have the money already, drinking the toxin will only make you ill for a day. If you already have the money, you don't need to drink the toxin to acquire it. So you can foresee at midnight that it

will be irrational for you to drink the toxin tomorrow morning (Beddor 2019). **Intention** is motivated thus: if we plug it in to the above line of reasoning, we can retrieve our initial judgement that you cannot, at midnight, rationally intend to drink the toxin the next morning.

Beddor argues that **Intention** plus **Defeat** solves **Dogmatic Intentions**. Consider Celia. Right now, she knows her coin is fair. She cannot rule out that she will tomorrow receive strong but misleading evidence e that her coin is biased. She can foresee that if she receives e , her knowledge that p will be defeated. If her knowledge that p is defeated, it will be irrational for her to ignore e . Add **Intention** and we get: Celia cannot rationally intend to ignore e . Problem solved.

3.3 Discussion

Except: no it isn't. There are two problems with Beddor's account. First, it can't handle an important class of dogmatic intentions. As Lasonen-Aarnio points out, not all evidence is defeating evidence. But it can be dogmatic to intend to ignore evidence whether or not that evidence is defeating evidence. Consider:

Paleontologist It has long been widely accepted in Pauline's field that the *diplodocus* only appeared towards the end of the Kimmeridgian age. Pauline herself has known 'the D-hypothesis' for years. A few years ago, Pauline developed a refinement of radiometric dating which yielded novel more evidence for the D-hypothesis. Everyone in the field increased their confidence in the D-hypothesis accordingly. Today, Pauline discovers, quite unexpectedly, that a new paper examining her radiometric dating has just been published. She forms an intention to maintain her current degree of confidence in the D-hypothesis even if she reads the paper and finds it

casts doubt on her radiometric dating.

Pauline is clearly irrational. But Pauline is not intending to ignore *defeating* evidence for the D-hypothesis: even if her data were completely discredited, she would still be in a position to know the D-hypothesis – it has been known for years – and so in a position to know that the evidence provided by the hostile paper is misleading. To properly satisfy the intention desideratum, we need to be able to explain why it can be irrational to intend to disregard even ‘weak’ – that is, non-defeating – evidence against some presently known p .

Second, **Intention** is either false or unmotivated. Consider the following:

Support. It is irrational for A to ϕ at t if A 's deliberation-available reasons at t require A 's to refrain from ϕ -ing. (That is, if any agent with A 's reasons who considers at t whether to ϕ is rationally required to refrain from ϕ -ing.)

Support is either false or true. If **Support** is false, **Intention** is unmotivated. If **Support** is true, **Intention** is false. That's bad news for **Intention**.

Let's take the horns in turn. First suppose: **Support** is false. We lose our argument for **Intention**. The argument for **Intention** depends on the claim that at 9am tomorrow, it will be irrational for you to drink the toxin. But that argument presupposed **Support**. We argued like thus:

Premise At 9am, your deliberation-available reasons do not favour drinking.

Conclusion So at 9am it is irrational for you drink.

The move from premise to conclusion relies on **Support**. But without the conclusion we have no argument for **Intention**. So if **Support** is false, **Intention** is poorly motivated.

Second, suppose: **Support** is true. Together, **Support** and **Intention** generate bad results. Consider:

Reluctant Jogger Homer has not been exercising much recently. He is beginning to feel the effects. On Wednesday, he judges – and indeed, the balance of reasons supports his judgement – that he really should begin to exercise more regularly. He resolves to go for a daily run, starting on Saturday. However, come Saturday, as his alarm starts to go off, his thoughts start to change. His bed is very warm and comfy, and he has had a hard week. He could start his running next week. The balance of Homer’s reasons have shifted: on Wednesday, he weighted his future health more heavily than his Saturday morning comfort. Now, Homer weighs his present comfort more heavily than his future health. The balance of Homer’s reasons now favours his staying in bed. (Holton 2004).

Suppose Wednesday-Homer can foresee that Saturday-Homer’s reasons will favour staying in bed. (All too plausible.) Is it therefore irrational for Wednesday-Homer to resolve to go for a run come Saturday? Surely not. It is under precisely such circumstances – circumstances in which we can foresee and wish to guard against future temptation – that we need and rely on resolutions. But **Support** and **Intention** together predict that it is irrational for Wednesday-Homer to resolve to go for a Saturday run. Given **Support** it will, on Saturday be irrational for Homer to go for a run, as *ex hypothesi*, Saturday-Homer’s reasons support his staying in bed rather than going for a run. And Wednesday-Homer can foresee this. So assuming Wednesday-Homer also knows **Support**, he knows it will be irrational for Saturday-Homer to head off on his morning run. But then, by **Intention** it is irrational for Wednesday-Homer to resolve to go for a run on Saturday. That’s a bad result. Something has gone wrong. It can’t be **Support**; *ex hypothesi* **Support** is true. It must be **Intention**. So if **Support** is true, **Intention** is false.

One could try a dodge: perhaps **Reluctant Jogger** misdescribes things.

Had Wednesday-Homer made no resolution, his Saturday-reasons would favour staying in bed. But Wednesday-Homer *has* made a resolution. That gives Saturday-Homer a special reason to go for a run. But the dodge comes at too high a price: it is unattractive to think that we can give ourselves extra reasons to ϕ simply by resolving to ϕ (Holton 2004).¹¹ The bad result stands. We should reject **Intention**.

3.3.1 Conditional Complications

But then, maybe we don't need **Intention**. Suppose Celia leaves it open that she will receive evidence e against p tomorrow. We want the result that it is irrational for her to intend to ignore such evidence.

Here's a choice-point. Either Celia *knows* she will receive e tomorrow, or she does not. If she *knows* today that she will receive evidence against p tomorrow, then, arguably – given some sort of reflection principle (Van Fraassen 1984) – she has a defeater right now for p , and so fails to know that p right now. But if this is so, no appeal to **Intention** is needed; **Intention** is otiose.

Now suppose Celia does not know she will receive e tomorrow; she just can't rule e 's receipt out. Then, arguably, Celia cannot foresee that it will be irrational for her to ignore e tomorrow. She can foresee only something weaker: *if* she gets e *then* it will be irrational for her to ignore it. But then **Intention** is silent on whether Celia may intend to ignore e . To get the desired result we need:

Intention-Conditional It is irrational for A to form at t_1 the conditional intention [ϕ if F comes to pass] if, at t_1 , A is in a position to know that

¹¹What if I really value keeping resolutions? Then the fact that I have resolved to ϕ may be a reason to ϕ when combined with the preference not to break resolutions. But this will only apply to those who value not breaking resolutions. But not everyone values keeping resolutions – some of us value our flighty spontaneity! A different tack, then: perhaps, knowing that I had resolved, previously, to go for a run might give me a reason to go for a run insofar as it suggests to me that deep down I really do care about my health. Sure, but the resolution is not doing any genuine work here; it is only a proxy for information about my real values (Holton 2004).

it will be irrational, if F comes to pass, for A to $[\phi]$.¹²

Intention-Conditional can be motivated with a a tweaked version of the toxin puzzle. Suppose Elon says ‘I will put the money in your bank account iff you form the following conditional intention: if you know the money is in your bank account, drink the toxin.’ Judgement: you cannot rationally form said conditional intention. Why not? *If* you know the money is in your bank account, it is irrational to drink the toxin: you already have the money. This is perfectly foreseeable. If we plug **Intention-Conditional** to this line of thought, we recover our judgement that it is impossible rationally form Elon’s proposed intention, so **Intention-Conditional** is well motivated.

Unfortunately, **Intention-Conditional** faces the same problems as **Intention**: it is either unmotivated (if **Support** is false) or false (if **Support** is true). If **Support** is false, we lose our argument for **Intention-Conditional** in much the same way that we lose the motivation for **Intention**. What if **Support** is true? Well, suppose Homer regularly eats to excess. Right now he wants to stop overeating, but can foresee that once he begins eating, things will look different: slimming down will seem less important and eating more pleasurable. A natural strategy? Form a conditional intention, say: stop eating if I need to undo my belt. **Intention-Conditional** and **Support** together say this resolution is irrational. So both cannot be true.

¹²Alternatively:

Intention-Strong. It is irrational for A to intend at t_1 to $[\phi$ at $t_2]$ if, at t_1 , A leaves it open that it will be irrational for A to $[\phi$ at $t_2]$.

But **Intention-Strong** is too strong: I may rationally resolve to buy a cheap car whilst leaving it open that I win the lottery. But I know that if I won the lottery, it would be irrational for me to buy a cheap car.

3.4 Onwards

Intention and its variants fail because they cannot accommodate the role played by resolutions. Holton suggests an alternative:

Resolve. It is irrational for A to resolve at t_1 to $[\phi$ at $t_2]$ if, at t_1 , A is in a position to know that, she should, at t_2 , reconsider whether to ϕ .

Resolve makes space for rational resolutions. That Wednesday-Homer resolved to go for a Saturday run does not give Saturday-Homer a reason to for a run, but it does give Saturday-Homer a strong reason not to reconsider whether to go for a run. Homer may be resolute or irresolute. Suppose he is resolute: he does not reconsider whether to go for a run. Then it will be rational for Saturday-Homer to go for a run: so long as one does not reconsider, it is rational to do what one has rationally resolved to do. Suppose he is irresolute: come Saturday, he acts contrary to his reasons and reconsiders whether to go for a run. He will have acted irrationally in reconsidering, but *once he reconsiders*, the rational thing for him to do really will be to stay in bed. His having had, prior to reconsideration, a reason not to reconsider, is not itself a reason to go for a run.

We need a final tweak. At present **Resolve** only delivers results for agents who can know quite a lot about what t_2 will be like. I might know that if I win the lottery tomorrow, then I should reconsider whether to buy the cheap but ugly car that I currently plan on buying. But I can know *this* without knowing that tomorrow I should reconsider whether to buy the cheap car, because I don't know whether I will win the lottery. Nonetheless, if I know I should reconsider whether to buy the cheap car if I win the lottery, it seems irrational for me to resolve to buy the cheap car even if I win the lottery:

Resolve-Conditional. An agent may not intend to $[\phi$ at t_2 even if $F]$ if she

is in a position to know that, if F is the case, then at t_2 , she should reconsider whether to ϕ .

3.5 Back to Beddor

We can now distinguish two ways I might intend to ignore evidence. I might form a *disregarding* intention:

Disregard! If I reconsider how likely it is that p I will assign p a probability that ignores e .

Alternatively, I might form a *blocking* intention:

Block! I will not reconsider how likely it is that p even if I get e .

Both look dogmatic. So the intention desideratum has two parts. Part 1: explain why disregarding intentions are irrational. Part 2: explain why blocking intentions are irrational.

Disregarding intentions are relatively easy so long as I intend to disregard defeating evidence. For such intentions, Beddor's story largely goes through once we substitute **Resolve-Conditional** for Intention. **Resolve-Conditional** predicts that my disregarding intentions are irrational if I can foresee the following: if I get new evidence e against some presently known p and reconsider whether p , I must reconsider whether to assign p a probability which ignores e . Given **Defeat**, it seems I can foresee this. I know that if I get defeating evidence against p , I will stop knowing that p . And if stop knowing that p I lose an important reason for thinking e is misleading. So I must reconsider whether I ought to assign p a probability which ignores it.

Unfortunately, this solution does not extend to cases in which I intend to disregard only 'weak' – that is, non-defeating – counter-evidence. And blocking intentions are harder across the board. **Resolve-Conditional** deems my

blocking intentions irrational only if I can foresee: if I receive e at t_2 , I will be rationally required to reconsider whether p at t_2 . For that we need a principle like:

Reconsider! Suppose A has judged that p at t_1 and at t_2 receives evidence against p such that were A to reconsider at t_2 whether p she would assign a lower credence to p . Then A must at t_2 re-open inquiry into whether p .

Is **Reconsider** plausible? It is not *obviously* false. But there's a decent case for its falsity. It's very demanding claim. Most of the new evidence I receive is such that it would be a lot of work to figure out which of my beliefs that p are such that, were I to reconsider them in light of this new evidence, I would lower my credence. To ensure conformity with the principle, I'd have to reconsider *a lot* of my beliefs every – or almost every – time I got new evidence. I don't have time for that, and nor do you. But if I don't do this, there's no way I will avoid violating **Reconsider!** without getting wildly lucky.

Of course, not everyone is moved by demandingness considerations – life is hard. But there's a deeper worry. It's plausible that the point of having beliefs (rather than or in addition to credences) is that beliefs simplify our reasoning, deliberation, and updating procedures (Holton 2014; Ross and Schroeder 2014; Staffel 2019). **Reconsider!** is in tension with this picture: forming a belief: given **Reconsider**, forming a belief does nothing to make updating on new evidence simpler. But in that case, why bother with a belief at all? The upshot? We need a different story about dogmatic intentions.

4 First Pass Responses To Bootstrapping Dogmatism

In this section, I'll canvas three potential responses to the bootstrapping version of the dogmatism puzzle. They don't work.

4.1 Attempt One: Restrict Ignore

Maybe we can restrict **Ignore**:

Ignore-Weak. Like **Ignore**, except in cases in which A has initially come to know that p on the basis of evidence which includes the misleading evidence; in such cases the misleading evidence must not be ignored.

This is patently *ad hoc*. And it has horrible implications. Suppose Sally and Celia are shown a coin. Sally is told on Monday that it comes from the fair coin factory, and then on Tuesday that it came up heads twenty times in a row. Celia is told on Tuesday though it came up twenty times in a row, it comes from the fair coin factory. Sally must, given **Ignore-Weak**, be more confident than Celia that the coin is fair. That's bad.

Worse still, there are close variants of the above bootstrapping puzzle which **Disregard-Weak** can't help with. Call these the *forward-looking* variants of the bootstrapping puzzle. Suppose A knows that p , and that she will tomorrow attend a lecture by an expert testifier who will be talking about whether p . She can reason like this:

p . And tomorrow an expert will tell me whether p . But given that p and given that they are an expert, the expert will tell me the truth about whether p , which is that p . And that an expert will tell me that p is good evidence that p – evidence I didn't have a minute ago! So the probability of p is even higher than I initially thought.

4.2 Attempt Two: No Double Counting

It's bad to 'double-count' evidence. Maybe our dogmatic bootstrapper illicitly 'counts' the evidence for p twice over.

Let's start by getting clearer on the idea of double counting.

[Double counting is] the fallacy of using the same evidence twice, to attempt to get an extra, albeit illicit, 'boost' to the posterior probability of an hypothesis...Suppose that a prosecutor introduces evidence that the defendant's fingerprints are on the murder weapon in order to raise the probability that the jury rationally assigns to the defendant's guilt. Suppose that they raise their collective probability to 0.5 on the basis of evidence introduced up to and including that stage. It would be fallacious for someone to then argue, 'Well, we all agree that the probability of the defendant's guilt is 0.5. Now, consider the fact that his fingerprints were found on the murder weapon. This fact should boost our estimate of his guilt to 0.75. The reason is that that very evidence has already been used in getting the posterior to 0.5 and cannot be reused in that way to further boost the posterior of the same hypothesis (Juhl 2007).

We can extract a working characterisation of double counting from these remarks: an agent double-counts some evidence when they use some evidence e to raise the probability of some proposition p , where their probability for p already takes e into account.

Someone who reasons like our dogmatic bootstrapper does not seem to double-count in anything like this sense. Our dogmatic bootstrapper raises their probability that p in light of learning that their evidence against p – that is, evidence m – was misleading. But the initial probability they assigned to p did not take into account that m was misleading. So it does not seem that

anything has been double-counted. One might insist: an agent who reasons from some evidence e to p , and then from p to some further thesis q ‘double counts’ the evidence for p . But this won’t do: it’s perfectly innocuous to reason thus. Consider Holmes, who reasons thus: ‘There was a cigarette butt left at the scene, so the murderer is a smoker. So the murderer can’t have been Jane – she doesn’t smoke’. Holmes does not double count the evidence of the cigarette butt.

4.3 Attempt Three: Doxastic Conservatism

Maybe doxastic conservatism is the answer. Doxastic conservatives think adopting an attitude gives you a reason to stick with it: ‘an agent is in some measure justified in maintaining a belief simply in virtue of the fact that the agent has that belief’ (Christensen 1994). It can be given a credal twist:

Credal Conservatism If A rationally assigns a degree of confidence n to p , that A has assigned that degree of confidence to p gives her, in the absence of special considerations, a strong reason to maintain that degree of confidence that p .

How might this help with dogmatic bootstrapping? Recall the reasoning:

In fact p . So all the evidence m that I just considered against p was misleading. But if m is misleading with respect to p , my probability for p should ignore m . So p is more probable than I initially thought.

The reasoner, we might think, acts irrationally because she does not act in accordance with her reason to stick with the probability she initially assigned to p .

But there are big problems with the conservative response. First, it is mysterious why assigning a credence might give you a reason to maintain it (Chris-

tensen 1994). Second, it's unclear that 'special circumstances' can be interpreted widely enough to disallow the very dogmatic reasoning we want to disallow. Third, and most significantly, **Credal Conservatism** appears to license its own species of bootstrappy reasoning, so any victory against bootstrappy dogmatism that relies on it will be pyrrhic. Suppose *A* has some complex, difficult-to-interpret evidence. She is unsure what credence she should have that *p*. Eventually, she guesses, and assigns a probability of 0.7. If **Credal Conservatism** is true, she can now be highly confident that this is the credence in *p* she should have. This seems like a bad result. So we should reject **Credal Conservatism**.

5 Introducing Volitionism

We have some puzzles. Their solution? **Volitionism**. Volitionism's central commitment is that it is a core function of belief to close inquiry (Friedman 2019).¹³ This core commitment motivates three further interlocking claims:

Openness. Let *p* be some complete answer to some question *Q*.¹⁴ Anyone who inquires into *Q* at *t* is such that they allow the chance that *p* to bear on their behaviour at *t*.

Don't-Believe-and-Inquire (DBI). One ought not to inquire into *Q* at *t* and believe *p* at *t*, where *p* is a complete answer to *Q* (Friedman 2019).

Judgements-Are-Resolutions (JAR). If *A* judges that *p*, she resolves not to (further) inquire into whether *p*.¹⁵

¹³For related lines of thought, see also Lawlor 2014; Staffel 2019; Holton 2014.

¹⁴The ideology of a 'complete answer' is borrowed from Friedman. To illustrate: to the question 'Is anyone going to the party?', 'Sarah is' is an incomplete answer. 'Sarah and Simon are going, but no one else' is a complete answer.

¹⁵The idea of **Judgements-Are-Resolutions** is not that a belief that *p* is nothing more than a resolution not to reconsider whether *p*. (That would yield the absurd result that a belief that *p* is a belief in *p*'s negation.) Nor am I committed to a constitution claim; viz., that

In this section, I explain and briefly motivate these commitments. The full argument for the view will emerge in the next two sections, as I show that volitionism can solve the three revenge puzzles.

5.1 Motivating Volitionism

I understand inquiry in much the same way as Friedman (Friedman 2019). For Friedman, when we inquire, we try to settle a question. (‘Where are my keys?’, or ‘How many planets are there?’) Sharpened up: inquiry requires a question-directed attitude. For A to genuinely inquire into some question Q , A must have an interrogative attitude to Q ; that is, question Q must be ‘open’ in A ’s thought. But if inquiry is a matter of *trying* to settle questions, forming beliefs is how we achieve what inquiry aims at. It’s how we settle questions. By contrast, I do not settle the question of Q by believing that it is probable that p , or assigning a high credence to the claim that p .

Friedman’s picture motivates **Openness**. Suppose that Q is the question ‘Is it the case that p ?’ Both p and not- p are complete answers to this question. Someone who inquires into Q clearly allows the chance that p , as well as the chance that not- p to bear on their behaviour. For someone who did not allow the chance that p to bear on their behaviour could not engage in an activity organised around an interrogative attitude to the question ‘is it the case that p ?’ But inquiry into Q is just such an activity.

Friedman’s picture also motivates **DBI**. It seems irrational to simultaneously ϕ and try to ϕ . But if inquiry into Q is an attempt to settle Q , and believing p means having settled p , one ought not believe p and inquire into whether p .

Friedman shows that **DBI** can be further motivated with (i) cases and (ii) linguistic data. Cases first. Friedman contrasts two inspectors, Morse and Morse*.

beliefs are partly constituted by resolutions; **Judgements-Are-Resolutions** is compatible with the weaker claim that judgements must be accompanied by resolutions.

Both are called to the scene of a murdered doctor. Both search the scene, talk to witnesses, and so on. Both, in short, engage in activities characteristic of inquiry. But there's a twist: whilst Morse has no idea who killed the doctor, Morse* knows full well that *he* is the murderer. Intuition: only Morse is genuinely inquiring into who killed the doctor. This motivates DBI. Now for the linguistic data. The following are infelicitous:

1 # *p* but I'm still inquiring into whether *p*.

2 # *p* but I'll look into whether *p* further tomorrow (Friedman 2019).

Here I add a forth argument for DBI. BDI follows from **Openness** once we grant an auxiliary assumption:

Deliberation. It is rationally incoherent to do both of the following at once: (i)

believe that *p*, (ii) allow the chance that not-*p* to bear on one's behaviour.

(Fantl 2018; Fantl and McGrath 2012).

Why accept **Deliberation**? Fantl and McGrath argue for **Deliberation** by contrasting deliberation that seems licit with deliberation that seems incoherent:

(a) On the one hand, the ice will hold my weight, and walking across the lake is a shortcut. On the other hand, I might slip and fall. Which is more important: that the ice is a shortcut that will hold me, or that I might slip and fall? ((Adapted from Fantl and McGrath 2012).

(b) # If I walk across the lake, I'll get there faster. And the ice will hold my weight. On the other hand, the ice might not hold me. Which is more important, that the ice will hold my weight, or that it might not? (Adapted from Fantl and McGrath 2012.)

As Fantl and McGrath point out, the first bit of deliberation is easy to make sense of: the deliberating agent is weighing up reasons which pull in different directions. There does not seem to be any vacillation: the agent has simultaneous deliberative access to all these competing reasons. In other words, the agent may rationally allow all of these competing reasons to bear on his behaviour at once.

The second bit of deliberation is not like this at all. If agents could have simultaneous deliberative access to ‘ p ’ and ‘it might not be the case that p ’, we would not expect such a sharp contrast between (a) and (b). In (b), the agent seems instead to be vacillating between two different deliberative perspectives. From one of these perspectives, ‘ p ’ is a deliberatively available reason but ‘it might not be that p ’ is not. From the other, ‘it might not be that p ’ is available but ‘ p ’ is not. The lesson? It is incoherent to believe p – that is, to adopt a perspective from which p is deliberatively available – whilst also treating ‘it might not be that p ’ as a deliberatively available reason (Fantl and McGrath 2012). **Deliberation** is the natural upshot. And when combined with **Openness**, **Deliberation** gets us DBI.

At this point, perhaps you worry: can a fallibilist accept **Deliberation**? If not, then it is not clear that **Deliberation** can play any meaningful role in responding to the dogmatism puzzle – a puzzle that is standardly taken to arise only for fallibilists.

Happily, despite a superficially infallibilist flavour, **Deliberation** is perfectly compatible with fallibilism (Fantl 2018). If **Deliberation** said that no coherent agent can both know that p and take it that there is a chance that not- p , then it might pose a problem for fallibilism. But **Deliberation** does not say that. Rather, it says that an agent cannot simultaneously treat both of these considerations *as reasons*. **Deliberation** posits a ‘deliberative blind-spot’, not

a requirement that beliefs come with certainty.¹⁶

5.2 Judgements are Resolutions

Central to the volitionist picture is the idea that a core function of belief that p is to close inquiry into whether p (Friedman 2019). But an act of mind can only *effectively* close inquiry into whether p if it gives the agent a reason to refrain from re-opening inquiry into whether p . The hiring committee have only effectively closed their inquiry as to who they ought to hire if they have some reason not to re-open the question as to who they should hire; if they have no reason not to immediately start reconsidering who they ought to hire, then their inquiry is not effectively closed. So a judgement that p can only effectively close inquiry into whether p – and so play its core functional role – if it gives an agent a reason not to (re-)open inquiry into whether p . We then have a puzzle: how might judging that p give an agent a reason not to re-open inquiry into whether p ? I contend that the best answer is that is **Judgements-Are-Resolutions** is true.

Without **Judgements-Are-Resolutions** how might making a judgement that p at t_1 give an agent reason not to reconsider, at t_2 , whether p ? Here are three alternatives:

- 1 When A judges that p at t_1 on the strength of sufficiently decisive evidence that p , A 's strong evidence at t_1 gives her at t_2 a reason not to reconsider whether p at t_2 .
- 2 When A judges that p at t_1 on the strength of sufficiently decisive evidence that p , and recognises, at t_1 , the strength of her evidence, the strong evidence at t_1 , together with her recognition of it at t_1 , give her, at t_2 , a reason not to reconsider whether p at

¹⁶See D'Cruz 2013 for discussion of deliberative blind-spots.

3 When A judges that p at t_1 on the strength of sufficiently strong evidence that p and knows that her evidential situation with respect to p will not change (much), her knowledge that she will get no new evidence gives her, at t_2 , a reason not to reconsider.

(1) looks implausible. How can the mere fact of my having had decisive evidence for p *yesterday* give me any reason to do anything today? What matters for what I should do today are the reasons I have today. (2) is similarly unpromising. How could the fact that I yesterday recognised anything about my evidence with respect to p at t_1 matter to what I should do today?

(3) looks better. If I knew yesterday that q would be true today, then today it is true that q , and so q could be a reason for me today. However. For (3) to furnish an explanation of how rational belief *per se* can effectively close inquiry, it would need to be the case that all rational judgement that p is accompanied by knowledge that the degree to which one's evidence supports p will not change, or, at the very least, that will not change much. Both Leitgeb and Lawlor defend something like this view; for Leitgeb, belief is stably high subjective probability: one believes p so long as, for all q compatible with p , $\text{Cr}(p|q)$ is above some threshold r (Leitgeb 2014). For Lawlor:

Belief's function is to permit the anticipation of constancy of the truth of one's belief, such that one can act without further inquiry. For belief to be rational, this anticipation itself must be rational (Lawlor 2014).

But the requirement that all rational judgement that p be accompanied by knowledge that one's evidential support for p will remain (roughly) constant is too strong. Consider:

Charlatan. Right now, Kelly knows that p . But she knows that tomorrow

she'll encounter a fabulously persuasive and apparently reliable orator who will tell her that p is false. Right now she knows this orator is a charlatan; but she also knows anyone who hears him speak forgets that he is a charlatan. So Kelly knows that tomorrow her evidence will support a belief that $\neg p$.¹⁷

Nonetheless, it is clearly permissible for Kelly to today judge that p . So (3) won't do: one may judge that p even whilst knowing one's future evidence will not support p .

Upshot: if (1) to (3) exhausted the candidates, we would be in trouble. Happily, **Judgements-are-Resolutions** provides a solution. Recall, first, how resolutions work. Wednesday-Homer's resolution to go for a Saturday morning run gives Saturday-Homer a reason *not to reconsider whether* to go for a run. And when an agent has rationally resolved to ϕ and rationally does not reconsider whether to ϕ , it is rational for them to ϕ . **Judgements are Resolutions**, then, suggests the following treatment of **Charlatan**. Wednesday-Kelly's judgement-generated resolution not to (re-)open inquiry into whether p gives Thursday-Kelly a reason not to reconsider whether to (re-)open inquiry into p . And just as Wednesday-Homer's run makes it rational for Saturday-Homer to go for a run *so long as he does not reconsider whether to do so*, Wednesday-Kelly's resolution makes it rational for her to keep inquiry into whether p closed, *so long as she does not reconsider whether to keep the inquiry closed*. And that's how beliefs close inquiry: they give agents reasons not to reconsider whether to inquire, thus making it rational to keep the inquiry closed.

That's the rough motivation for Volitionism. Now let's see the work it can do. We have three desiderata in need of satisfaction: intention, weak, and

¹⁷This case shares suggestive features with some classic counterexamples to strong reflection principles – see (Briggs 2009). If I know that I'll get genuine expert testimony tomorrow that p is false, then I surely don't now count as knowing that p . By if I know that I will receive similar testimony from a fraud, this does not undermine my present knowledge.

bootstrapping. Let's go.

6 Putting Volitionism to Work – Part 1

In this section I solve the weak and the bootstrapping dogmatism puzzles.

6.1 The Weak Puzzle, Solved

A refresher. At t Celia knows her coin is fair. Then she flips it; it comes up heads twenty times in a row. The evidence is not strong enough to defeat her knowledge. So she can reason thus:

My coin is fair. So evidence suggesting that my coin is not fair is misleading, at least with respect to the claim that my coin is fair. But if evidence is misleading with respect to p , my degree of confidence that p should ignore that evidence. So it is just as probable now that my coin is fair as it was before.

We can now articulate a solution. Let p be the claim that the coin is fair. The solution starts with the following banal claim: Celia's reasoning is licit only if she is permitted to reason from the claim that p . Hence:

1 Celia's reasoning is licit only if she has an outright belief that p .

But at the same time, Celia is doing the following: assessing how her evidence bears on p , with a view to letting that assessment guide her in assigning attitudes to p . It's hard, then, to make sense of what is going on here unless we take Celia to be *inquiring* into whether p . Hence:

2 Celia is inquiring into the question of whether p .

And DBI tells us:

3 If Celia has an outright belief that p and is inquiring into whether p , then she then she is doing something illicit.

From (2) and (3), we get:

4 If Celia has an outright belief that p , she is doing something illicit.

But combined with (1), (4) gets us:

Conclusion Celia is doing something illicit.

Good: this is the right result. And it's a result for volitionism, because the result relies on DBI.

6.2 Objections and Replies

Hang on, though. Is (2) true? Perhaps one thinks something like this:

In the reasoning above, we need not interpret Celia as inquiring into whether p . Rather, Celia may be interpreted as inquiring into the following question: *how probable is it that p* . So we should reject (2).

This is a good objection. Happily, the volitionist can respond. Recall the principle **Openness** introduced above: Let p be some complete answer to Q . Anyone who inquires into Q is such that, they allow the chance that p to bear on their behaviour. From **Openness**, it follows that anyone who inquires into how probable it is that r allows the chance that q to bear on their behaviour, where q is a complete answer to the question of how probable it is that r .

Now, it seems to me that the following is true: not- r is a complete answer to the question 'How probable is it that r ?'¹⁸ Consider the following exchanges,

¹⁸I am bracketing cases in which the question is interpreted as concerning objective probability.

where A knows that the results for Michigan and Georgia have been called, but doesn't realise that B has already checked the news:

A How likely is it that Trump won Georgia?

B Trump lost Georgia!

C How likely is it that Biden won Michigan?

D Biden *did* win Michigan!

Or consider a case in which B but not A is looking out of the window:

A How likely is it to rain today?

B It's already raining.

It follows, together with **Openness**, that anyone who inquires into the question 'How probable is it that p ' allows the chance that not- p to bear on their behaviour. But in that case, it is just as irrational to inquire into the question 'How probable is it that p ?' whilst believing p as it is to inquire into the question 'Is it the case that p ?' whilst believing p .

6.3 Tweaked Dogmatism

Still, you might have reservations. The first reservation is linguistic: it insists that p is not a complete answer to the question 'How probable is p ?'. The second worry imagines that Celia tweaks her reasoning slightly:

My coin is fair. So evidence suggesting that my coin is not fair is misleading, at least with respect to the claim that my coin is fair. But if evidence is misleading with respect to p , my degree of

confidence that p should ignore that evidence. ~~So it is just as probable now that my coin is fair as was before.~~ So my evidence supports p just as much as it did before.

Celia still seems guilty of dogmatic reasoning. But she no longer seems to be inquiring into either the question of *whether her coin is fair* or into the question of *how probable it is that her coin is fair*. Rather, she seems to be inquiring into something like the following question: ‘(to what degree) does my evidence support the claim that my coin is fair?’

The arguments given above cannot account for what is troubling about these tweaked bits of dogmatic reasoning. But if volitionism can’t account for the badness of tweaked dogmatic reasoning, it can’t account for the badness of dogmatic reasoning.

One might try to repeat the move we trialed above, and insist that ‘ p ’ is a complete answer to the question, ‘To what degree does my evidence support p ?’ But that move won’t work for a second time: it isn’t plausible that someone who asks whether their evidence supports p receives an answer of any kind when they are told that p :

A To what degree does my evidence support Trump having won Georgia?

B # Trump lost Georgia!

So we need a different response. Happily can respond to the linguistic and the ‘tweaked reasoning’ worry in one go. Here’s the strategy. As well as **Deliberation**, we endorse:

Deliberation*. It is rationally incoherent to both: believe that p , allow the chance that one’s evidence doesn’t fully support p to bear on one’s behaviour.

Openness and **Deliberation*** together predict that the tweaked dogmatic reasoning is illicit. **Openness** tells us that Celia, by inquiring into the degree to which the evidence supports p , allows the chance that her evidence does not fully support p to bear on her behaviour. But she also relies, in said inquiry, on the outright belief that p . Hence she violates **Deliberation***. So her reasoning is illicit. Further, this account of why Celia's reasoning is defective does not depend on taking ' p ' to be a complete answer to the question of how probable it is that p . So the account also addresses the linguistic worry. Once again, volitionism gets the right result. Clearly, though, this account relies on **Deliberation***. Happily, we can argue for **Deliberation*** in the same way that we argued for **Deliberation**. We argued for the latter by noting a marked contrast between ordinary felicitous reasoning in which an agent weighs competing considerations, and cases in which an agent weighs p against the chance that not- p . Now consider a variant:

(a) On the one hand, the ice will hold my weight, and walking across the lake is a shortcut. On the other hand, I might slip and fall. Which is more important: that the ice is a shortcut that will hold me, or that I might slip and fall?

(c) # On the one hand, the ice will hold my weight, and walking across the lake is a shortcut. But if my evidence doesn't guarantee that the ice will hold me, I shouldn't walk across the lake. And my evidence might leave it open that the ice will break. Which is more important, that the ice will hold my weight, or that my evidence might leave it open that it won't? ¹⁹

(c) has the same paradoxical feel as (b). Thus, insofar as (b) motivates **Deliberation**, (c) motivates **Deliberation***.

¹⁹Note that such an agent will not violate standard akratic norms.

Let us take stock. Agents who inquire into some question Q allow, for any given complete answer to Q , the chance that said answer is true to bear on their behaviour. But for some answers – ‘not- p ’, ‘my evidence doesn’t support p ’ – there may be a rational tension between (i) allowing the chance of those answers’ truth to bear on one’s behaviour, and (ii) believing that p . It is of precisely this sort of tension of which weakly dogmatic agents are guilty.

Now for the bootstrapping puzzle.

6.4 The Bootstrapping Puzzle, Solved

A refresher. A judges and comes to know that p on the basis of a body of mixed evidence. Given **Ignore** it seems A may reason as follows:

In fact p . So the evidence m I just considered against p was misleading. But misleading evidence with respect to p is irrelevant to how probable it is that p . So m is irrelevant to the probability of p . So p is more probable than I initially thought.

More vividly, suppose Jeri reasons like this:

In fact the drugs do work. So the study in *The Lancet*, which suggested that they don’t, is misleading. But misleading evidence with respect to p is irrelevant to how probable it is that p . So the study in *The Lancet* is irrelevant to the probability that the drugs work. So that the drugs work is more likely than I initially thought.

Ugh. Happily, the solution to **Weak Dogmatism** generalises:

By inquiring into Q , Jeri allows the chance that p is false to bear on her behaviour, because ‘not- p ’ is a complete answer to the question ‘How probable is p ?’ (See above.) Hence, she either (i) both believes

p and allows the chance that not- p to bear on her behaviour (bad) or (ii) does not believe p . But if she does not believe p , then she reasons from a claim – viz., that the drugs do work – that she does not know (also bad). Upshot: volitionism gives a unified solution to both the weak and the bootstrapping puzzle.

You might object: Jeri does not allow the chance that p is false to bear on her behaviour. For her inquiry takes it for granted that her evidence either supports p to a high or to a *very* high degree, and that doesn't seem readily compatible with her allowing the chance that not- p to bear on her inquiry.

This is a good objection. Nonetheless, the volitionist solution still generalises. Suppose we take Jeri to be inquiring into something like the following question: does my evidence support p to degree x or to degree y ? Even if both x and y are very close to 1 (indeed, even if one of them is equal to 1), so long as x and y take different values, someone who inquires into this question must leave it open that their evidence supports p to a degree less than 1, and so that their evidence does not fully support p . And so someone who inquires into this question allows the chance that their evidence does not fully support p to bear on their behaviour. But then, if Jeri is allowing the chance that her evidence does not fully support p to bear on her behaviour, she cannot also reason from p , without violating **Deliberation***. The upshot? Volitionism gives a unified solution to both the weak and the bootstrapping puzzle. Hurrah for volitionism.

6.5 A Test

Any adequate solution to the dogmatism puzzles must retain **Ignore**. So the volitionist solution had better be compatible with **Ignore**. Happily, it is. The best way to see this is by considering an argument to the contrary:

Consider Jeri. Volitionism predicts that she must maintain a proba-

bility for the claim that the drugs work which is sensitive to, rather than one which ignores, the evidence in *The Lancet*. But **Ignore** says that she must do the opposite. So volitionism is incompatible with **Ignore**.²⁰

First, it's wrong to think there's a special problem for volitionism here. Volitionism predicts the (strongly!) intuitive result that Jeri should not raise her confidence that the drugs work on coming to know that the drugs work. If this result is incompatible with **Ignore**, then **Ignore** is incompatible with an intuitive data point, not a distinctively volitionist commitment. If there's a problem here for volitionism, it's because there's a problem for everyone.

Happily, though, there is no problem for volitionists here. That's because **Ignore** does not entail what the argument above takes it to. The argument above presumes that Jeri violates **Ignore** simply by doing both of the following at once: (i) knowing that the drugs work, (ii) having a probability for the claim that the drugs work which is sensitive to *The Lancet's* evidence that they don't. But this is incorrect. To violate **Ignore** Jeri would need to be doing a third thing in addition to (i) and (ii). She would need, in addition, to be:

(iii) in the process of assigning a doxastic attitude to the claim that the drugs work.

But Jeri need not be doing this third thing. She may have already assigned her attitudes, and moved on with her life. In that case, **Ignore** should have no quarrel with Jeri.

Suppose, on the other hand, Jeri *is* in the process of assigning a probability to the claim that the drugs work. At this point, the volitionist solution kicks in. Here's how. First, if Jeri knows that *The Lancet* study is misleading, it's because she knows that the drugs work. So **Ignore** only tells her to ignore *The*

²⁰We could run an argument *mutatis mutandis* concerning the weak puzzle.

Lancet's study whilst she's assigning a probability if she knows that the drugs work whilst she's assigning a probability. Second, if Jeri is in the process of assigning a probability to the claim that p , Jeri is allowing the that chance not- p to bear on her behaviour. But **Deliberation** tells us that no rational agent can (i) allow the the chance that not- p to bear on her behaviour whilst also (ii) believing that p . Upshot: if Jeri is rational, she won't believe – and so will not know – that the drugs work whilst in the process of assigning a probability to the claim that the drugs work. So if Jeri is rational, Jeri won't know that *The Lancet*'s study is misleading whilst she's assigning her probability to the claim that the drugs work. And so **Ignore** won't tell her to ignore *The Lancet* study whilst she's assigning this probability.

Not only, then, is volitionism *compatible* with **Ignore**. It also presents a complete account of how and why **Ignore** is compatible with our strongly held anti-dogmatic intuitions as to the unacceptability of dogmatic bootstrapping. So there's no problem for volitionists here. But there *is* still a problem for everyone else.

7 Putting Volitionism to Work – Part 2

Thus far, the anti-dogmatic work has been done by **Openness** and **DBI**. **JAR** – the most radical of the volitionist commitments – has been silent. Thus, even those who reject *full* volitionism can help themselves to the solutions developed above. This section offers no such succour. When it comes to solving the puzzle of dogmatic intentions, **JAR** is central.

At first sight, **JAR** might seem an unlikely component in an anti-dogmatic epistemology. **JAR**, recall, is the view that an agent who judges that p resolves to close inquiry that p . Combined with Holton's view of resolutions, it tells us that, for an agent who has rationally judged that p , it is rational to keep inquiry

into whether p closed. But that sounds very...dogmatic.

This worry rests on a mistake about resolutions. A rational resolution gives the resolving agent a strong *pro tanto* reason not to reconsider. But even strong *pro tanto* reasons can be defeated (Holton 2004). Consider the following variants of **Reluctant Jogger**:

Lava Homer has not been exercising much recently, and he is beginning to feel the effects. On Wednesday, he judges – and indeed, the balance of reasons supports his judgement – that he really should begin to exercise more regularly. He resolves to go for a daily run, starting on Saturday. However, come Saturday, he hears on the radio that his running route has become engulfed by lava.

Drizzle. Homer has not been exercising much recently, and he is beginning to feel the effects. On Wednesday, he judges – and indeed, the balance of reasons supports his judgement – that he really should begin to exercise more regularly. He resolves to go for a daily run, starting on Saturday. However, come Saturday, he hears on the radio that it is forecast to drizzle.

In **Lava**, Homer clearly ought to reconsider whether he ought to go for a run. The lava is an excellent reason to reconsider, as well as an excellent reason not to go for a run. The drizzle, by contrast, is not a good reason to reconsider: it may be a reason not to go for a run, but it is not also a good reason to reconsider whether to go for a run. Only the lava, not the drizzle defeats the *pro tanto* reason provided by the resolution. We can capture the difference thus: the lava but not the drizzle is *unsettling* with respect to Homer's resolution to go for a run.

So: when are the *pro tanto* reasons provided by rational resolutions defeated? It would be a mistake to expect a tidy account: practical rationality is messy. But, as Holton remarks, we can give some plausible rules of thumb:

it is irrational to reconsider if one is faced with the very temptations that the resolution was designed to overcome, or if one's judgment will be worse than it was when the resolution was formed. It will be rational to reconsider if the reasons for forming the resolution no longer obtain, if the agent discovers that her judgement was seriously impaired during her original deliberation, if the stakes are much higher than one initially assumed, or if circumstances otherwise turn out to be importantly different from those anticipated – for example, if one finds oneself in a scenario which you did not, when forming the resolution, treat as a live possibility (Holton 2004).²¹

You might worry that there's something paradoxical about this picture. Suppose I have resolved to ϕ . Then I get some new information. If the new information is *not* unsettling, I should not reconsider whether to ϕ . On the other hand, if the new information *is* unsettling, I should reconsider whether to ϕ . So to work out whether to reconsider whether to ϕ , I need to figure out whether the new information is unsettling. But to figure this out, I need to know whether I should still ϕ . And to figure that out, I need to reconsider whether to ϕ . But that means I need to reconsider whether to ϕ before I can figure out whether I ought to reconsider whether ϕ . But that's impossible! Oh dear.

Lucky for me, this worry rests on two mistakes. First, it is false that I can never figure out whether some new information is unsettling with respect to my resolution to ϕ without reconsidering whether to ϕ . It's easy to figure out that learning I was drunk when I decided to ϕ , or that the stakes are much higher now than they were when I made my resolution, are unsettling with respect to my resolution to ϕ . I can work that out without thinking at all about the first-order question of whether I ought to ϕ . Second, the worry above assumes that the only way for me to be guided by reasons is for me to deliberate on and reason from those reasons. But this is false. I can be guided by reasons dispositionally.

²¹See Holton 2014 for discussion of 'live possibilities'.

A reason R can guide me to θ in virtue of my having a disposition to θ when R -type circumstances obtains (Holton 2004).

We can get a better feel for the proposed rules of thumb with some cases:

Hypochondriac Louis suffers from hypochondria: whenever he notices a twinge or a click, he starts to calculate the odds of it being a symptom or an indication of some cancerous growth. And indeed, many of the clicks and twinges he experiences are mild – very mild! – evidence for the claim that he has cancer. But right now, Louis knows that whenever he experiences the twinges, he – by his present lights – vastly overestimates the extent to which they support the cancer hypothesis. In an attempt to avoid hypochondriac rumination, Louis yesterday judged that he does not have cancer, and resolved to stop thinking about it. Today, he experiences a twinge, and does not reconsider whether he has cancer.

Husband Sarah's husband occasionally works late. For a while, this made Sarah very nervous: whenever he worked late, she would anxiously weigh up the evidence as to whether or not he was cheating on her. This made her very unhappy. After one careful survey of the evidence, she judges that her husband is faithful; she thereby resolves not to start, the next time he works late, to begin weighing up the evidence as to whether he is cheating on her. When he next works late, she does as she intends.

Lizards. Right now, I know that the Queen is not a lizard. But I know I may well encounter some at least apparent evidence for the claim that she is a lizard, because I know I may well encounter testimony to the effect that she is a lizard from a conspiracy theorist. I think seriously considering this evidence would be a waste of time, so I resolve to ignore any such testimony. Suppose, on the other hand, I were to learn tomorrow that Prince Charles is a lizard. Learning that would make it rational for me

to re-open inquiry into whether the Queen is a lizard, because my initial judgement was not made whilst taking seriously the possibility of receiving such evidence.

Policeman. Lilly has found fairly decisive evidence that her boyfriend, Talar, is cheating on her: she has found that he owns a secret phone, on which he has been exchanging sexually explicit messages with another woman. She judges and comes to know that he is cheating on her. She resolves to confront him. She expects a tearful denial that the phone belongs to him, or something similar. Lilly also suspects that once she is actually speaking to him, she will believe any denial he issues because of desperate wishful thinking. So she resolves to ignore his denials. Once she confronts him, though, he tells her something completely unexpected: he is an undercover policeman, and the messages are part of a sting. Given how different Talar's response is to anything she anticipated, and that Lilly has never considered whether she would consider this sort of scenario to count as cheating, it is rational for her to consider re-opening inquiry.

Holton's remarks pattern well with my judgements about these cases: in the former two, reconsideration does not seem called for, in the latter two, it does; this is just what Holton's rules of thumb predict. I now turn to the intention puzzle proper.

7.1 The Intention Puzzle, Solved

A refresher. Suppose you now know that p . If you know that p , then you know that if you get evidence against p tomorrow, it will be misleading. And if you know that e will be misleading, it seems rational to form blocking intention: to intend not to reconsider whether p if one gets e . If e is not defeating evidence, you may form a disregarding intention instead: reconsider whether p if one

receives e , but assign p a probability which ignores e . Both seem dogmatic.

Disregarding intentions – even disregarding intentions focused on weak evidence – are easy for the volitionist. Given DBI, it is foreseeable that if I (permissibly) re-open inquiry into whether p , it will no longer be rational for me to believe p , and so I will lose my reason for thinking e misleading. So it is foreseeable that if I reconsider whether p , I must reconsider whether to assign p a credence which ignores e . Thus by **Resolve-Conditional** it is irrational for me to intend: if I receive e and reconsider whether p , assign p a credence which ignores e . Good: the problem of disregarding intentions is solved.²²

We still need to explain the irrationality of blocking intentions. Above, I considered an explanation that rested on **Reconsider**. Maybe I was too quick to dismiss **Reconsider**. Maybe it's just obvious that an agent is irrational if they have a degree of confidence that p at t which is higher than the degree of confidence they would have were they to rationally reflect on the evidence they actually have at t . So **Reconsider** is true.

It's not. **Reconsider** gives the wrong result in three key cases: **Charlatan**, **Hypochondriac** and **Husband**. In these cases, agents rationally maintain at t a degree of confidence which outstrips the degree of confidence they would have were they to rationally reflect on the evidence they have at t .

We now have a complicated spread of data. It seems that whilst it is sometimes (perhaps often) irrational to intend to block defeating evidence for p , it can be rational to intend to block defeating evidence (see: **Charlatan**). And whilst it can be rationally permissible to intend to block non-defeating evidence against p (see: **Husband**, **Hypochondriac**), it is sometimes impermissible to do so (see: **Paleontologist**). To properly satisfy the intention desideratum, we must be able to explain this range of data points. We need a story that does

²²I will be using '(re)consider whether p ' and '(re-)open inquiry into whether p ' interchangeably.

justice to the complexity of the data.

There are two options. We can tell a backwards-reaching story or a synchronic story. A backwards-reaching story explains which intentions I can rationally form *now* in terms of which actions I might permissibly perform *later*. A synchronic approach explains which intentions an agent may form at t in terms of what the agent's available-to-deliberation reasons support at t . Below, I consider and dismiss some synchronic constraints.

7.2 Synchronic Constraints

To go synchronic, we need to explain why it is irrational for A at t to form an intention to ignore future evidence in terms of the agent's situation at t . Here goes.

Suggestion 1: you should intend to block evidence if you can foresee that you will fail to conditionalize on it. This gets the right results in **Charlatan** and **Hypochondriac**, though not **Husband**. More generally, it will radically over-generate rational blocking intentions. Not many of us are perfect Bayseian updaters – we forget things! And most of us know that, and so can foresee failures to conditionalise.

A weaker constraint with a similar spirit might do better: perhaps I can form a rational blocking intention if I can foresee that I won't respond to future evidence as I currently think that I should. Again, this doesn't handle **Husband**. More generally, it will over-permit blocking intentions. Maybe I have good evidence that as people age they find arguments for modal realism more plausible. It does not seem rational for me to intend to block new arguments for modal realism that I might hear in twenty years time, despite having good evidence that I will find them more persuasive than I would if I were to hear them today.

Suggestion two: it is permissible to form a blocking intention if you know that your epistemic position will be better if you block than if you don't. Read one way, this hopeless. This will make it permissible for anyone who knows p to intend to block all defeating evidence for p : if I know that p then I know that I am in a better epistemic position if I continue to believe that p than if I do not.

A different gloss might work. Perhaps it's permissible for me to intend to block evidence e if I am rationally confident that the credence in p I will end up with if I block e is closer to $\text{Cr}(p|e)$ than the credence in p I will end up with if I take e into account. But this won't do either. Why not? Consider our jejune modal anti-realist again. Suppose that, unbeknownst to him, in twenty years, he'll be given a new argument – call it the novel argument – for modal realism, which he has never encountered before. For him to assess whether blocking e_a – the evidence provided by the novel argument – will get him a credence closer to $\text{Cr}(\text{modal realism is false}|e_a)$ than not blocking will, he would need to have some sense of the value of $\text{Cr}(\text{modal realism}|e_a)$. But he has no clue of that, because, ex hypothesi, right now he has no clue what the novel argument is.

Suggestion 3: You may intend to block e for p if e is 'low impact' evidence against p but not if it is 'high impact' (Dallmann 2017). But on any reasonable construal of high vs low impact evidence, this predicts that Pauline the paleontologist is rational in forming a dogmatic intention and Kelly, anticipating the charlatan's testimony, is not. So this won't do.

Suggestion 4: You may at t intend to block e for p if at t you know: (i) p (ii) the conditional 'if e then p ', and (iii) your knowledge of the conditional does not depend on your having low confidence that e (Sorensen 1988).²³ This gives the right result in **Charlatan**. Before Kelly goes to see the charlatan, Kelly knows p . She also knows that: even if the charlatan tells her that p is false, p . And she

²³Thanks to Bernhard Salow for making me take this suggestion seriously.

is highly confident in the antecedent of the conditional. So her knowledge of the conditional doesn't depend on low confidence in the antecedent. Compare this with the situation of Celia in **Coin**. Before she flips her coin a million times, she knows both that the coin is fair. Maybe she also knows the conditional: if the coin comes up heads a million times in a row, it's nonetheless fair. But if she does know this conditional, it's knowledge which depends on her being highly confident that the antecedent is false. So she may not permissibly intend to block the evidence of a million flips coming up heads.

Unfortunately, this suggestion cannot handle intentions to ignore non-defeating evidence. Pauline the dogmatic paleontologist knows both the D-hypothesis and that even if her dating is unsound, the D-hypothesis is true. And her knowledge of the conditional does not depend on low confidence in the antecedent, because the antecedent is not a defeater for the consequent.²⁴ So by the above suggestion she may intend to block the evidence that her dating is unsound. But that's the wrong result.

What if we tweak the proposal: you may intend to block e for p if (i) you know p , (ii) have high confidence in the conditional 'even if e , p ' and (iii) your exact degree of rational confidence in the conditional does not depend on your low confidence that e . This handles the problem of Pauline the paleontologist: upon learning e her confidence in the conditional should decrease. But where the initial proposal was too concessive, the new proposal is not concessive enough: it gives the wrong results in **Husband** and **Hypochondriac**.

Upshot: the synchronic strategy does not look promising.

²⁴Proof: because the antecedent is not a defeater for the consequent, she can know the conjunction of the antecedent and the consequent. But anyone who knows [p and q] can know [if p then q].

7.3 Volitionism to the rescue

Suppose I have resolved to buy a cheap car, taking my lack of funds to be a very important consideration, but neglecting to take seriously the possibility in which I win the lottery. I may not, without further deliberation, resolve to buy the car even if I win the lottery. Why not? Because it is foreseeable that I should reconsider whether to buy the cheap car if I were I to win the lottery. More generally, an agent who has carefully deliberated and resolved to ϕ may not, without further deliberation, form the intention to [ϕ even if F] if it is foreseeable if F happens, they will be required to reconsider whether to ϕ . The volitionist can use this to build an account of when we may intend to block evidence. The volitionist, recall, thinks that judging that p involves resolving to close inquiry into whether p . The volitionist also has a story about when the reasons provided by resolutions are defeated: in general, an agent who has resolved at t to close inquiry into whether p must consider whether to re-open inquiry if circumstances are very different to those anticipated when making the resolution, or they learn that their judgement at t was importantly compromised, and so on. The circumstances under which an agent will be required to reconsider whether to keep inquiry closed, or to re-open it, are ϕ are *unsettling circumstances*. The volitionist story about when it is permissible to intend to ignore evidence will go like this:

Suppose an agent who has judged that p at t and so resolved at t to close inquiry into whether p . They form an intention to ignore evidence e if they get it. This is permissible only if it is not foreseeable that receipt of that evidence would not unsettle the resolution to close inquiry into p that was made at t .²⁵

²⁵Given the messiness of practical rationality, this will not amount to a reductive analysis of which disregarding intentions are permissible: our judgements as to which intentions are permissible and which situations would count as unsettling are probably too intermingled for that. But it provides an illuminating structural analysis nonetheless.

This gives good results. It predicts that it is permissible for the agents in **Charlatan**, **Husband** and **Hypochondriac** to intend to block the charlatan's testimony, bodily twinges, and late night working respectively. Why? Because the agent's initial – rational – judgements were made whilst taking seriously the prospect of receiving this evidence. Thus receipt of that evidence is not an unsettling circumstance for their resolution to close inquiry into p . By contrast, it predicts that it is impermissible for Pauline the paleontologist to intend to block evidence that her radiometric dating is unreliable, and impermissible for the agents in **Lizards** and **Policemen** to intend to block discovery that Prince Charles is a lizard, or the boyfriend's testimony that he is an undercover policeman. Why? Because receipt of those bits of evidence marks their situations as importantly different from those anticipated when the agent's initial judgements were made. Thus the volitionist can satisfy the intention desideratum. And they can do so with the flexibility and nuance required by our spread of examples.

Importantly, volitionism does not commit us to nasty results like the following. Suppose that Wanda takes very seriously the possibility that she won't win lottery: she knows the lottery is fair, and that she has only one of millions of tickets. Nonetheless, she judges, on the basis of wishful thinking, that she will win the lottery. For the volitionist, she thereby resolves to close inquiry into whether she will win the lottery. One might worry as follows:

Volitionism predicts that Wanda may ignore the newspaper reports that say her numbers did not come up. It makes this prediction because she made her resolution to close inquiry into whether she will win the lottery whilst taking seriously the possibility that she would not win the lottery, and so taking seriously the possibility that she would read such reports in the newspaper. Hence such reports

are not unsettling with respect to Wanda's resolution, and, once she reads them, Wanda may permissibly maintain her judgement that she is a winner.

Clearly, that would be a bad result. But it's not a result to which volitionism is committed. Volitionism would generate such a result only if paired with the claim that Wanda's initial resolution to close inquiry was a rational one when made. But volitionists are committed to no such claim.

Let us take stock. Rational resolutions prevent certain novel information – 'it's drizzling', for example – from impacting on what we ought to do. They quarantine this information from the question of how to act. Rational beliefs act similarly. They prevent certain novel information – 'another conspiracy theorist says that the Queen is a lizard', perhaps – from having any impact on the doxastic attitudes we ought to have. They quarantine this information from the question of what to believe. They do this because they are tightly connected with a particular kind of resolution: the resolution to close inquiry. This picture can explain both why it's generally irrational to intend to ignore evidence against some known p , and why such intentions are sometimes permissible, because it allows us to see the questions raised by such intentions as one species of a broader genus. That genus? The question of which conditional resolutions – ϕ even if F an agent who has rationally resolved to ϕ might make.

8 Conclusion

We started with three desiderata for an adequate response to the dogmatism puzzle. Volitionism satisfies all three. This strongly motivates volitionism. I close with one final point in volitionism's favour. Volitionism does not *just* solve the revenge puzzles. It also provides a deeper explanation of why the standard solution to the classic dogmatism puzzle – viz., **Defeat** – gets things right.

Why is it, we might ask, that agents must give up their beliefs in standard defeat cases? Volitionism provides an answer. In standard defeat cases, I contend, two things happen. First, the new evidence is *unsettles* the agent's prior resolution to close inquiry into whether p . And once the agent is required to re-open inquiry into whether p , the agent is required – by DBI – to give up their belief that p during this inquiry. Second, because the agent's total evidence no longer supports a belief that p ; they will not be able to close this newly opened inquiry by judging that p . Hence they must remain without a belief that p .

Volitionism, then, not only provides a solution to the 'revenge' puzzles which is compatible with, rather than undermining of, the standard solution to the classic puzzle. By locating **Defeat** within a rich explanatory network of claims concerning the dynamics of rational belief, it transforms would-be-avengers into sympathetic allies. That's no mean feat.

So: it turns out that theoretical judgement requires an act of will. That's surprising. You might try to blunt the interest of the thesis with an appeal to bounded rationality. Here's the idea. It is rational for us to close inquiry because we are limited: cognitively, computationally, and so on. That is why we need to form beliefs. But unbound creatures would never rationally close inquiry, or at the very least, would never rationally close inquiry until they reached absolute certainty. So the fusing of theoretical judgement and the will is an artefact of our cognitive frailty.²⁶ I'm not sure that makes the fusing less interesting. Nonetheless, I think the thought is mistaken in a deeper way, which is brought out by **Husband**. Few long term projects are such that the agents who embark on them will always be such that their evidence supports their continued endeavour (Buchak 2017; Morton and Paul 2019). Nonetheless, it can be rational to stick with such projects. In many cases – a doctorate, for example – it will be enough for eventual success to have resolved in advance to

²⁶For this sort of thought, see Holton 2014.

a course of action, and so have resolved to keep acting *as if* the endeavour is in good standing. But for some long-term projects – marriage, friendship – this is not enough. Sarah’s marriage is not a success if she continually leaves it open that her husband is unfaithful, however well she acts. For her marriage to be a success, she must really believe that he is faithful. What **Husband** reveals, then, is that the capacity to close certain kinds of inquiry into those we love is required for us to enter into trusting interpersonal relationships. And the value of such relationships has little or nothing to do with our bounded cognitive capacities. Even computationally unlimited creatures would need a capacity to judge, if they were to have friends and partners in whom they placed their trust.²⁷

References

- Malcolm, Norman (1952). “Knowledge and belief”. In: *Mind* 61.242, pp. 178–189.
- Unger, Peter (1978). *Ignorance: A case for scepticism*. OUP Oxford.
- Kavka, Gregory S (1983). “The toxin puzzle”. In: *Analysis* 43.1, pp. 33–36.
- Van Fraassen, C (1984). “Belief and the Will”. In: *The Journal of Philosophy* 81.5, pp. 235–256.
- Harman, Gilbert (1986). *Change in view: Principles of reasoning*. The MIT Press.
- Sorensen, Roy A (1988). “Dogmatism, junk knowledge, and conditionals”. In: *The Philosophical Quarterly (1950-)* 38.153, pp. 433–454.
- Christensen, David (1994). “Conservatism in epistemology”. In: *Nous* 28.1, pp. 69–89.

²⁷Acknowledgements: Thanks to Allan Hazlitt, Nick Hughes, Tristram McPherson, Beau Madison Mount, David Plunkett, and Bernhard Salow for helpful comments. Thanks also to members of the Oxford Epistemology Group and Washington Ethics and Epistemology Group for helpful discussion.

- Holton, Richard (2004). “Rational resolve”. In: *The Philosophical Review* 113.4, pp. 507–535.
- Juhl, Cory (2007). “Fine-tuning and old evidence”. In: *Noûs* 41.3, pp. 550–558.
- Briggs, R. (2009). “Distorted reflection”. In: *Philosophical Review* 118.1, pp. 59–85.
- Kripke, Saul A (2011). *Philosophical troubles: Collected papers*. Vol. 1. OUP USA.
- Fantl, Jeremy and Matthew McGrath (2012). “Pragmatic encroachment: It’s not just about knowledge”. In: *Episteme* 9.1, pp. 27–42.
- D’Cruz, Jason (2013). “Volatile reasons”. In: *Australasian Journal of Philosophy* 91.1, pp. 31–40.
- Holton, Richard (2014). “Intention as a Model for Belief”. In:
- Lasonen-Aarnio, Maria (2014). “The dogmatism puzzle”. In: *Australasian Journal of Philosophy* 92.3, pp. 417–432.
- Lawlor, Krista (2014). “Exploring the Stability of Belief: Resiliency and Temptation”. In: *Inquiry* 57.1, pp. 1–27.
- Leitgeb, Hannes (2014). “The Stability Theory of BeliefThe Stability Theory of BeliefHannes Leitgeb”. In: *The Philosophical Review* 123.2, pp. 131–171.
- Ross, Jacob and Mark Schroeder (2014). “Belief, Credence, and Pragmatic Encroachment 1”. In: *Philosophy and Phenomenological Research* 88.2, pp. 259–288.
- Baker-Hyatt, Max and Matthew A Benton (2015). “Defeatism defeated”. In:
- Hawthorne, John, Daniel Rothschild, and Levi Spectre (2016). “Belief is weak”. In: *Philosophical Studies* 173.5, pp. 1393–1404.
- Buchak, Lara (2017). “Faith and steadfastness in the face of counter-evidence”. In: *International journal for philosophy of religion* 81.1-2, pp. 113–133.

- Dallmann, Justin (2017). *When obstinacy is a better (cognitive) policy*. Ann Arbor, MI: Michigan Publishing, University of Michigan Library.
- Dorst, Kevin (2017). “Lockeans maximize expected accuracy”. In: *Mind* 128.509, pp. 175–211.
- Yablo, Stephen (2017). “Open knowledge and changing the subject”. In: *Philosophical Studies* 174.4, pp. 1047–1071.
- Fantl, Jeremy (2018). *The limitations of the open mind*. Oxford University Press.
- Fix, Jeremy David (2018). “Intellectual Isolation”. In: *Mind* 127.506, pp. 491–520.
- Beddor, Bob (2019). “The Toxin and the Dogmatist”. In: *Australasian Journal of Philosophy*, pp. 1–14.
- Friedman, Jane (2019). “Inquiry and belief”. In: *Noûs* 53.2, pp. 296–315.
- Morton, Jennifer M and Sarah K Paul (2019). “Grit”. In: *Ethics* 129.2, pp. 175–203.
- Staffel, Julia (2019). “How do beliefs simplify reasoning?” In: *Noûs* 53.4, pp. 937–962.