



Fallibility and Dogmatism

Bernhard Salow

To cite this article: Bernhard Salow (2025) Fallibility and Dogmatism, Australasian Journal of Philosophy, 103:1, 23-38, DOI: [10.1080/00048402.2024.2381692](https://doi.org/10.1080/00048402.2024.2381692)

To link to this article: <https://doi.org/10.1080/00048402.2024.2381692>



© 2024 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group



Published online: 09 Aug 2024.



Submit your article to this journal [↗](#)



Article views: 1205



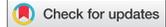
View related articles [↗](#)



View Crossmark data [↗](#)



Citing articles: 1 View citing articles [↗](#)



Fallibility and Dogmatism

Bernhard Salow 

Magdalen College, Oxford

ABSTRACT

The strongest version of the dogmatism puzzle argues that, when we know something, we should resolve to ignore or avoid evidence against it. The best existing responses are fallibilist, and hold that decisions should be governed by underlying probabilities rather than our knowledge. I argue that this is an overreaction: by paying close attention to the principles governing belief-revision, and to subtly different ways in which knowledge can govern decision-making, we can dissolve the puzzle without the need for controversial theoretical commitments. The resulting solution demonstrates fruitful and underexplored points of interaction between ‘traditional’ epistemology and ‘formal’ theories of belief-revision, and clears the ground for more systematic theorizing about how and when we should be open to changing our minds.

ARTICLE HISTORY Received 17 August 2022; Revised 25 September 2023

KEYWORDS dogmatism paradox; fallibilism; knowledge-action principles; belief revision

1. Introduction

If p is true, evidence against p is misleading. So if you know p , you can deduce that evidence against p is misleading. But if you know that some evidence is misleading, you should avoid engaging with it, by ensuring you never encounter it or ignoring it when you do. Yet it’s also clear that you should remain open to changing your mind even about things which you, in fact, know. That is the dogmatism puzzle, discovered by Kripke, and introduced to the literature by Harman (1973).

Harman (1973) also offers the canonical response. Knowledge is defeasible: it can be lost when you encounter opposing evidence. So, once you receive the evidence against p , you will no longer know p , and hence no longer know that the evidence is misleading. So, you will not be entitled to ignore it after all.

Unfortunately, as Hawthorne (2004: 181) and Kripke (2011: appendix C) argue, this leaves a residual puzzle. Harman’s response explains why you shouldn’t ignore the evidence *once you receive it*—you will no longer know that it is misleading then. But it doesn’t explain why you shouldn’t take precautions *beforehand*, when you still know that it is misleading. So it doesn’t explain why you shouldn’t take a pill that prevents you from changing your mind about p ; nor why you shouldn’t avoid situations that might yield evidence against p , say by avoiding anyone who disagrees.¹

¹ Lasonen-Aarnio (2014) raises another objection to Harman’s response, which I discuss in §5.

One response is to adopt fallibilism, the view that almost everything we know has a non-zero probability of being false. For when you then reason with these probabilities, the benefits of engaging with the evidence in the unlikely case that it's not misleading can outweigh the disadvantages of doing so in the much likelier case that it is.² But solving the puzzle this way means either denying that knowledge plays a significant role in practical reasoning (if your probability-based decisions to engage with the evidence are rational even though you know that the evidence is misleading), or else accepting widespread 'pragmatic encroachment' (if your knowledge that the evidence is misleading disappears whenever you face such decisions).³

This is unnecessary. I will show that we can solve the residual puzzle without fallibilism. In particular, we can solve it while both granting knowledge a significant role in practical reasoning and rejecting pragmatic encroachment. If such a solution is possible, it is preferable. For the fallibilist solution carries many theoretical commitments: that enough of our knowledge has non-maximal probability, but also that the relevant probabilities exist in the first place. Perhaps these commitments are correct. But they go beyond what is necessary to resolve the puzzle, and thus generate a less illuminating explanation of why the dogmatic argument fails; much like, when a marble ends up at rest at the bottom of a basin, an explanation that appeals to its exact starting location is less illuminating than one which makes do with the more minimal fact that it was dropped somewhere over the basin.⁴

I will argue that the key to resolving the dogmatism puzzle is not to accept fallibilism, but to take a consistent stance regarding a controversial principle of belief revision. Using 'discover that q ' to mean that q is added to one's evidence, the principle is:

Preservation if q is consistent with everything you are entitled to believe, then you are entitled to retain all of your beliefs upon discovering that q .

In particular, I will argue that if Preservation is true, the dogmatic argument cannot be run (§3). And if Preservation is false, we can distinguish two principles about knowledge and action that are otherwise equivalent. Only one licenses the dogmatic argument; we should opt for the other, and can do so without diluting the role of knowledge in practical reasoning (§4). So, whether Preservation is true or false, the dogmatic argument collapses.

But before arguing this, I will reconstruct the dogmatic argument in more detail.

2. Reconstructing the Puzzle

The version of the dogmatism puzzle set out above closely follows the literature. It is not, however, the most perspicuous version. One issue is that it focuses on *knowing* that p . But the puzzle arises for any attitude which is (a) defeasible, in that further evidence can destroy it, even though (b) we are entitled to 'reason from' propositions to

² Good (1967) and Oddie (1997) argue that gathering and accommodating cost-free evidence always maximizes expected utility and expected accuracy. Arguably, some of their assumptions are too strong, but the results still show that the probabilities rarely support dogmatic decisions.

³ Ye (2016) and Beddor (2019) defend the former. Hawthorne (2004: 181) suggests the latter: given Good's (1967) result, this naturally falls out of the proposal—developed by Fantl and McGrath (2002) and Weatherston (2005)—that pragmatic encroachment occurs whenever the difference between maximal and actual probability affects your preferences over available options.

⁴ The example is from Strevens (2008: 434–35).

which we bear this attitude when deciding what to do. Knowledge is one candidate for such a relation—but other plausible ones are justified outright belief or a justified state of being sure. To postpone some subtleties discussed in §5, I will initially focus on (justified, outright) belief.⁵

Another issue is the talk of evidence being *misleading*. I'm not sure what 'misleading evidence' means in ordinary English, but I doubt it means *evidence against the truth*, as the argument assumes. For, in the ordinary sense, two people can agree that some observation or statistic used to defend p is misleading, even while disagreeing about whether p is true. Better, then, to drop the term, to avoid equivocating.

My version of the puzzle thus begins as follows. Suppose you justifiably believe that p , and know that you do. Since you believe yourself to be rational, you believe that you would give up your belief that p if you were to encounter and accommodate sufficiently strong evidence against p .⁶ Since you believe that p , you believe that your belief that p is true. Putting the two together, you can deduce that you would lose a true belief if you were to encounter and accommodate strong evidence against p . So you believe that, as far as your accuracy about p is concerned, it would be better to avoid or ignore strong evidence against p .⁷

To complete the puzzle, we need an argument from here to the rationality of some dogmatic decision. This would be straightforward if you could justifiably believe that one course of action A —opening a book, talking to a certain person—would result in you receiving strong evidence against p , while another course of action B —returning the book unopened, avoiding the person—would not. For if justified belief plays any role in practical reasoning, then presumably the following principle is true:

Strong Dominance If you justifiably believe that B 's consequences would be better than A 's, you should choose B over A .

Since you justifiably believe that A would yield strong evidence against p , and that B would not, and since you also justifiably believe that (as far as your accuracy about p is concerned), it is better not to encounter strong evidence against p , this principle recommends B : returning the book unopened.

It is hard to see, however, how the required combinations of beliefs could be rational. For suppose you are justified in believing that A would lead you to learn something that would defeat your justification for believing p . Then that is surely enough to already defeat your justification for believing p . Even if you never open the book, you shouldn't believe p if you believe that it compellingly calls p into question. But if you lack justification to believe p , the dogmatic reasoning can't get started.⁸

⁵ I will often drop 'justified' and 'outright'. I will also use 'rational' and 'reasonable' interchangeably with 'justified' and 'is entitled to believe' interchangeably with 'has justification to believe'.

⁶ I assume throughout that you believe both that you will accommodate evidence rationally and that you won't lose relevant evidence you currently have. When these conditions aren't met, it needn't be dogmatic to ignore or avoid evidence.

⁷ Kripke (2011: 43–44) also focuses on your concern for truth. Of course, sensible people don't care only about their accuracy about p ; and the other things they care about (being responsible believers, being accurate on topics other than p , etc) may still support engaging with additional evidence. But this doesn't dissolve the puzzle: if you are willing to consider arguments against p because, even though you think this unhelpful to your accuracy about p , it promotes other goals you have, this still seems problematically dogmatic. So I will follow the dogmatist in ignoring all effects aside from your accuracy about p .

⁸ Similar reasoning is why Sorensen (2018: §3) discusses cases where you may get no evidence. The reasoning here appeals to something like van Fraassen's (1984) 'reflection principle', especially its variant for outright belief discussed by Bovens (1995).

We can't sidestep this concern by focusing on other decisions: taking a belief-preservation-pill, forming a generic resolution to ignore any evidence against p (or about p) you come across, forming a conditional resolution to ignore the book if it contains such-and-such material. For Strong Dominance won't support these decisions unless you believe that they would definitely leave you better off than the alternative, and this won't be true unless you believe that the alternative would definitely involve losing your belief. Yet, since you're not worried about becoming irrational or simply forgetting, it is hard to see how you could be justified in believing this while also believing p .⁹

The dogmatist should respond by appealing to a stronger principle of practical reasoning. Perhaps being justified in believing p is incompatible with positively believing that certain courses of action would yield strong evidence against p . But surely, this response claims, you could suspend judgement about whether they would. When you either believe that p is true or suspend judgement on whether it is, I will say that *you believe that p might be true*;¹⁰ the claim is then that you can be justified in believing p while also believing that certain courses of action might yield strong evidence against p . Moreover, it seems at least somewhat plausible that, even if accommodating evidence that is weak or further supports p wouldn't reduce your accuracy about p , it also wouldn't increase it. How could it, when you already believe the truth?¹¹ But then we can complete the puzzle with

Weak Dominance If you justifiably believe that B's consequences would be at least as good as A's, and that they might be better, you should choose B over A.

For you can believe that an action—reading that book, say—would either make no difference to whether you believe p (if it yields further evidence supporting p , or weak evidence against p), or else result in you not believing p (if it is strong evidence

Some might worry that the reasoning has sceptical consequences. Couldn't you believe p , while also believing that some evidence out there would rationally convince you to drop that belief? Couldn't you, for example, believe that Arsenal won last night's match based on one newspaper report, while also believing that—given the possibility of error and the sheer number of rival papers—some paper somewhere will have misreported the result? You could, but that's not enough. To generate a problem, we need to add that you also believe that, when you encounter the evidence, it will still have this evidential force. And we can't do that. If we imagine that you always read all the papers, and are thus bound to come across the erroneous report eventually, then it's surely possible that you'll first come across enough papers that agree with you to reasonably maintain your belief about the results when you encounter the disagreeing one. And if we imagine that the paper was pre-selected to disagree with you, it seems reasonable to maintain your belief when you read it: for this paper, but not your original source, was chosen in a way that yields false reports most of the time. See Salow 2018 for related discussion.

Goodman and Salow (*forthcoming*) develop reasons to think that there may, after all, be cases in which one can believe both p and that gathering additional evidence will defeat one's justification to believe p . These cases are highly unusual—I set them aside here.

⁹ Kripke's formalization of the argument doesn't explicitly invoke Strong Dominance. Instead, the central premise is 'If A knows that taking an action of type T leads to consequence C , and A wishes above all else to avoid C [...], then A should resolve now not to take any action of type T '; we then substitute 'looking at alleged evidence against p ' for type T (2011: 43–44). But this faces the same problem: since one plausibly can't know both p and that looking at alleged evidence against p will definitely lead one to (rationally) stop believing p (the undesirable consequences C), this principle doesn't support the conclusion.

¹⁰ Should fallibilists balk at this usage? After all, they hold that one can fully believe or know something while acknowledging a small chance of error. But they can read my usage of 'might p ' as reporting not merely some chance that p , but of reporting a sufficiently high chance that p to render a contrary belief unjustified. (Such a usage, moreover, wouldn't be unfamiliar—it is presumably the usage that makes ' p , but it might be that not p ' sound incoherent).

¹¹ You could believe p more firmly, so that you would continue to believe it even if you were to encounter evidence that would previously have defeated your belief. But that won't save us from the conclusion that, as far as your short-term accuracy about p is concerned, you should avoid the evidence. This is bad enough.

against p). And while you're not entitled to believe that the second *would* happen, you are entitled to believe that it *might* happen. Moreover, you believe that the first possibility is an equally good outcome to that achieved by not reading the book, while the second is worse. So, by Weak Dominance, you shouldn't read the book.

This is the version of the dogmatic argument I will focus on. In the next section, I will show that it presupposes that we reject Preservation—for if Preservation is true, there are no cases in which you believe both p and that you might receive strong evidence against p . I will then show that without Preservation, Weak Dominance diverges from the following, very similar-looking principle:

Conditional Dominance If you justifiably believe that B's consequences would be better than A's if they aren't equally good, you should choose B over A.

Conditional Dominance captures everything appealing about Weak Dominance, without supporting dogmatic reasoning. If Preservation is false, then, we should respond to the dogmatic argument by rejecting Weak Dominance for Conditional Dominance. We thus get a satisfying solution to the puzzle, whether Preservation is true or not.

3. Preservation and its Limits

The dogmatic argument based on Weak Dominance requires cases in which you believe p while also believing that you might receive sufficiently strong evidence against p to defeat your justification for believing p . Can there be such cases? Not if Preservation is true. For to think that some proposition q might become part of your evidence, you have to think that q might be true.¹² And if you believe that q might be true, you can't believe anything inconsistent with q .¹³ So, if you think that q might become part of your evidence, then q is consistent with everything you believe. But then, by Preservation, you would remain entitled to all of your beliefs upon discovering that q is true, that is, upon q being added to your evidence. So you would remain justified in believing p upon discovering q , meaning that q is not sufficiently strong evidence against p . So no evidence which you believe you might receive can be strong evidence against something you currently believe.

Preservation, moreover, is not an *ad hoc* principle. It is part of AGM, the standard theory of belief revision.¹⁴ It follows from Smith's (2016) normic theory of justification.¹⁵ And it follows from Bayesian conditionalization if we identify outright belief with credence 1, as Clarke (2013), Greco (2015), and Dorst (2019) do—even if we employ primitive conditional probabilities, so that new evidence that received a prior credence of 0 can undo a prior credence of 1.¹⁶

¹² This follows from the claim, defended by Williamson (2000), that all evidence is true. Even if we think that evidence can be false, I don't think we should allow that one can anticipate this. But if I'm wrong, and one can anticipate receiving false evidence, there may be nothing dogmatic about avoiding it. Note that Good's (1967) and Oddie's (1997) arguments that we should always welcome additional evidence also presuppose certainty that no falsehood will become part of one's evidence. See Das (2023) for discussion.

¹³ Recall that I use 'believes that it might be the case that q ' to mean 'believes that q or suspends judgement about whether q '; so this is simply the claim that if p is inconsistent with q , one can't rationally believe p while also believing or suspending judgement about q .

¹⁴ AGM was formulated by Alchourrón, Gärdenfors, and Makinson (1985). Lin (2019) is a recent introduction, focused on Preservation.

¹⁵ See Smith 2016: §4.3, §7.3.

¹⁶ Dorst rejects the identification of credence 1 with belief, but identifies it with the attitude of being sure, which plays the theoretical role I'm taking belief to play.

Notably, many of the examples featuring in the dogmatism literature don't clearly violate *Preservation*, and thus don't clearly exhibit the structure required for the dogmatic argument. The following case, from Hawthorne (2004: 71), is representative:

Newspapers

You believe that Arsenal won last night's match against Chelsea, having read this in the *Times*. You also know that reading the *Guardian* report anything other than a win would defeat your justification to believe that Arsenal won.

This is supposed to generate pressure to avoid reading the *Guardian*, or to ensure you ignore it if you do. This implicitly assumes that a further condition is satisfied: you believe that the *Guardian* might not report a win. For if you don't believe this, why should you worry about reading it?

Yet, on the most natural versions of the case, this condition isn't satisfied. When I read the *Times* report an outcome of a match, I believe that the *Guardian* reports that outcome too. After all, the *Guardian* is a decent paper: when the outcome was such-and-such, this is what the *Guardian* will report. Perhaps this doesn't generalize to *all* newspapers; there are some I think quite unreliable, who might well have gotten it wrong. But I would continue to believe the *Times* even if I discovered that those papers disagreed—after all, I think that the *Times* is reliable and that they are not.

Something similar is true of other standard cases. Sorensen (1988) cites his belief that his car is parked where he left it, which he would give up if his friend Doug told him that it's not. So far, that seems plausible; but now we need to add that he also thinks his friend Doug might actually tell him this. Plausibly, the case collapses at that point. For this requires Sorensen to think that Doug might tell him that his car isn't where he left it *even though it is*—but if he thinks that, why would he trust Doug's testimony?¹⁷

Both of these examples involve rebutting defeaters. But the situation is no different for undercutting ones. The wall in front of you looks red, and so you justifiably believe that it is red. You know that, if you were to notice red lighting when you look at the ceiling, this would defeat your justification to believe that the wall is red. Can we add to the story that you also justifiably believe that, when you look up, you might find such lighting? Plausibly, we cannot: if you had reason to believe that there might really be trick lighting around, you would not have been justified in taking the wall's appearance at face value.

So *Preservation* isn't obviously false. And endorsing it does make for a very satisfying response to the dogmatic argument. The response is essentially that you should see no need to avoid or ignore evidence, because you should believe that any evidence you might actually receive won't threaten your current view. Much like the proverbial schoolyard bully, the dogmatist behaves poorly not because he is too self-assured, but because he is, at heart, not self-assured enough. If only he stood properly behind his belief that p —by maintaining that q , being strong evidence against p , is false and thus not something he might discover—he would not act as he does.

That being said, I do think there are cases that put pressure on *Preservation*. Here is one:

¹⁷ Kripke's (2011: 35, 45) main example is similar, so the same point applies. Sorensen (2018) gives a better example; but I don't think it's obvious that, in his case, the discovery defeats your justification, given what you know about how the counterevidence was compiled.

Coin

A coin that you know to be fair was flipped until it landed heads.

Arguably, you are entitled to have some beliefs about how often the coin was flipped: for example that it wasn't flipped 1,000,000 times.¹⁸ But you should not be maximally opinionated: you should allow that it might have been flipped at least three times. But now take a number towards the top end of what you should leave open: maybe it's 8. Since 8 is at the top end, you don't believe that the coin might have been flipped a fair bit more often than that—12 times, say. But you do think that it might have been flipped 8 times, hence that it might have been flipped at least 8 times. And if you were to discover that it was flipped at least 8 times, you'd have to adjust what you believe. In particular, you would then no longer be entitled to believe that the coin wasn't flipped 12 times. This makes you a counterexample to Preservation.

We can't escape this counterexample by maintaining that you should have no beliefs about how often the coin was flipped. For, surely, if you were to discover that it was flipped over 1,000,000 times, this would defeat your knowledge that it was fair. So there'd still be some p (namely: that the coin is fair) which you believe, and some q (namely: that it was flipped at least 1,000,000 times) consistent with all your beliefs, such that you would be required to abandon p upon discovering q —and so Preservation would still fail.

Similar examples arise without explicitly probabilistic structure. Consider

Suitcase

You pick up your suitcase, and it strikes you as weighing about 15 kg.

Plausibly, for some $x > 0$, you should believe all and only that the suitcase weighs somewhere between $15-x$ and $15+x$ kg. After all, your estimates are neither perfectly reliable nor entirely uninformative. For concreteness, suppose x is 3. Now suppose you discover something you currently think might be true, namely that the suitcase weighs more than 17.5 kg. Then you should re-evaluate your earlier conviction that it weighs less than 18kg: if your estimate was definitely out by more than 2.5 kg, it may well have been out by as much as 3 kg. (If 17.5 kg isn't enough to trigger this intuition, imagine discovering that it weighs more than 17.99 kg.) Preservation fails.¹⁹

Some will resist these counterexamples, perhaps suspecting that they rely on subtle context-shifts and/or illegitimate assumptions about vagueness. I needn't disagree, since I have already shown that the dogmatic reasoning breaks down if Preservation is true. But the examples are compelling enough that we should consider how to respond to dogmatism if Preservation is false.

¹⁸ Compare Hall (1999), Dorr, Goodman, and Hawthorne (2014), and Goodman and Salow (2018).

¹⁹ While Suitcase resembles Williamson's (2000) examples of inexact knowledge, its status as a counterexample to Preservation is independent of Williamson's controversial claims about margins for error or the KK principle. See Goodman and Salow 2023, forthcoming, for discussion.

A similar counterexample to Preservation arises in the Surprise Exam paradox, where you are initially justified in believing that there will be a surprise exam in a certain period, leave open that it will be towards the end of that period, but will lose your justification to believe that there will be a surprise exam if you discover that it wasn't towards the beginning or middle of the period. Kripke (2011) discusses this example at length in the paper formulating the dogmatism paradox, but doesn't make explicit this connection between them.

4. Conditional Dominance

Suppose, then, that Suitcase is a counterexample to Preservation. And suppose that there is no pragmatic encroachment, so that it remains such a counterexample even when you are offered the opportunity to discover whether the suitcase weighs more than 17.5 kg. Then Weak Dominance supports dogmatism. For you can reason as follows: ‘If I were to find out whether the suitcase weighs more than 17.5 kg, I would either discover that it doesn’t weigh more than 17.5 kg, or that it does—either might happen. In the first case, I would keep believing that it weighs less than 18 kg, so I’d neither gain nor lose accuracy on that issue. In the second case, I would give up that belief. But that belief is true, so this is a loss in accuracy. So finding out definitely wouldn’t improve matters, and might make things worse. So I should avoid finding out whether the suitcase weighs more than 17.5 kg and/or make sure I don’t change my mind when I do’. But the dogmatic conclusion seems no more acceptable in this (admittedly contrived) example than in Newspapers.

Without pragmatic encroachment (to which I return in §6), then, there remains only one option: Weak Dominance will have to go. However, we must tread carefully here. The simplest fallibilist response also rejects Weak Dominance, concluding that outright belief plays no important role in practical reasoning. If our response is to be more minimal, more tailored to dogmatism in particular, we must replace Weak Dominance with a principle that preserves such a role.

Fortunately, Conditional Dominance provides just this kind of replacement. Weak Dominance, recall, says that you should choose *B* over *A* whenever you believe that *B* would be either equally good or better. Conditional Dominance says that you should choose *B* over *A* whenever you believe that *B* would be better if it isn’t equally good. At first sight, the two seem interchangeable.

But they are not. A minor divergence concerns cases where you believe that *A* and *B* would be equally good, but also believe that *B* would be better if they’re not. Weak Dominance goes silent here, while Conditional Dominance tells you to choose *B*. This consequence of Conditional Dominance seems welcome. You might outright believe that Hume died before 1800, and hence that a \$100 bet that Hume died before 1800 will net the same result as an unconditional pay-out of \$100; still, you should choose the latter.²⁰

More surprisingly, the two principles also diverge in our counterexamples to Preservation. For in these cases there is an important difference between believing a disjunction (*q* or *p*) and believing the corresponding conditional (if not *q*, then *p*), even when you believe that *q* might be false. In Suitcase, for example, you believe that the suitcase weighs either no more than 17.5 kg or no more than 18 kg; but you do not believe that, if it weighs more than 17.5 kg, it still weighs less than 18 kg. This judgement is intuitive, given the Preservation-violating construal of this case, and it is backed up by the plausible principle that you’re justified in believing the

²⁰ Roeber (2016: 177–78) discusses a similar example. One might worry that, when combined with a probability 1 view of belief, this conflicts with Expected Utility Theory, since the two options have the same expected utility. However, the best versions of Expected Utility Theory allow that options with equal expected utilities needn’t be equally permissible. For even if we assign probability 1 only when mathematically required, there will be some options that are clearly better than others while having 0 probability of resulting in a different outcome: for example, a sure pay-out of \$100 and one that requires a fair coin flipped infinitely often to land heads at least once. See Easwaran 2014: 14–15 and Briggs 2019: §3.2.3.

indicative conditional ‘if not q , then p ’ only if discovering not q would leave intact your justification for believing p .²¹

Consider, then, how Conditional Dominance applies in Suitcase. You’re wondering whether to find out whether the suitcase weighs more than 17.5 kg. Which do you think is better if they don’t have equally good outcomes? Well, not finding out will leave you believing that the suitcase weighs at least 18 kg. Finding out will leave you believing this only if it weighs at most 17.5 kg. So the hypothesis that the two don’t have equally good outcomes just is the hypothesis that the suitcase weighs more than 17.5 kg. So our question becomes: which outcome, continuing to believe that it weighs less than 18 kg or dropping that belief, do you believe is better if it weighs more than 17.5kg?

It is tempting to think that the answer is: you believe that continuing to believe this is better, if that is so. After all, you believe that this belief is true, and truth is all you care about here. But that reasoning is fallacious. What matters is not what you believe is better *simpliciter*, but what you believe is better *if the suitcase weighs more than 17.5 kg*. And, as we just saw, you do not believe that, if it weighs more than 17.5 kg, it still weighs less than 18 kg. So you also do not believe that your current belief that it weighs less than 18 kg is definitely true, if it weighs more than 17.5 kg. So you do not believe that retaining this belief is better, accuracy-wise, if the suitcase weighs more than 17.5 kg. So you do not believe that the consequences of not finding out are definitely better if the suitcase weighs more than 17.5 kg. Since that hypothesis just is the hypothesis that the two options don’t lead to equally good outcomes, you do not believe that not finding out is better if the outcomes are different. Conditional Dominance thus doesn’t support the dogmatic conclusion that you should avoid finding out.

So Conditional and Weak Dominance give different verdicts in Suitcase; and Conditional Dominance escapes the dogmatic verdict that you should avoid finding out. The difference arises because, when you believe that your action might have a certain outcome (such as you losing your belief that the suitcase weighs less than 18 kg), Weak Dominance looks at your unconditional evaluation of that outcome, while Conditional Dominance looks to your evaluation of that outcome *on the hypothesis that this action causes it*. Since you can believe unconditionally that an outcome is bad, without believing that it is bad on the hypothesis that it comes about as a result of this action, the two come apart.

Conditional and Weak Dominance are, nevertheless, extremely close. For suppose the following instance of Preservation holds: If it is consistent with everything you are entitled to believe that A and B ’s outcomes are equally good, then you are entitled to retain all of your beliefs upon discovering that A and B ’s outcomes are equally good. Then we can recover Weak Dominance from Conditional Dominance. For suppose you meet the conditions of Weak Dominance: you believe that B ’s outcome is at least as good as A ’s and that it might be better. Since you believe that B ’s outcome might be better, it is consistent with everything you believe that A and B ’s outcome are not equally good. It follows from our instance of Preservation that you would remain entitled to believe that B ’s outcome is at least as good as A ’s upon discovering that they aren’t equally good. Since B ’s outcome being at least as good but not equally good as A ’s is the same as B ’s outcome being better, you would also be entitled to

²¹ I address concerns about this principle in §5.

believe that *B*'s outcome is better than *A*'s upon discovering that they aren't equally good. But then you are already entitled to believe that *B*'s outcome is better than *A*'s if they aren't equally good. And so Conditional Dominance says that you should choose *B*, which is exactly the verdict Weak Dominance would call for. So Weak Dominance follows from Conditional Dominance when the relevant instance of Preservation holds.²²

This observation undercuts most of the appeal of endorsing Weak Dominance instead of, or in addition to, Conditional Dominance. For even if Preservation can fail, §3 showed that it has intuitive appeal in a wide range of cases. We can thus explain both why Weak Dominance seems attractive, and why it is often a perfectly acceptable heuristic for rational decision making, while endorsing only Conditional Dominance. The observation also shows that the retreat to Conditional Dominance is in an important sense minimal for solving the dogmatism puzzle. For it leads us to reject Weak Dominance only in cases where Preservation fails—and these are exactly the cases in which Weak Dominance licences dogmatic reasoning.

This completes my argument for replacing Weak Dominance with Conditional Dominance, if Preservation can fail. Should there be more? The only reason I have given to reject Weak Dominance for Conditional Dominance is that Weak Dominance leads to dogmatism in counterexamples to Preservation; this shouldn't convince a committed dogmatist. But it doesn't need to. The dogmatist is arguing for an implausible conclusion: they are worth taking seriously only because their argument for this conclusion looks compelling. It thus suffices to point out that one of their premises is not, in fact, well-motivated. Highlighting the alternative of endorsing only Conditional Dominance, and showing that this preserves any appeal Weak Dominance might have had, achieves this—no independent argument against Weak Dominance is needed.

5. Objections and Replies

I have now laid out my solution to the puzzle. We can maintain that outright belief plays a significant role in practical reasoning via Conditional Dominance, without licensing dogmatism. If Preservation is false, Conditional Dominance fails to entail Weak Dominance, thus invalidating the dogmatic argument. If Preservation is true, Conditional Dominance entails Weak Dominance, but Weak Dominance doesn't licence dogmatism, because we shouldn't believe that we might encounter evidence that would change our mind.

Before closing, I will discuss four objections to my solution.

Objection 1. *Your solution is formulated in terms of belief; the original puzzle concerned knowledge.*

Both solutions generalize to plausible knowledge-theoretic analogues of our decision principles. Consider first:

Weak Knowledge Dominance If you know that *B*'s consequences would be at least as good as *A*'s, and that they might be better, you should choose *B* over *A*.

Given Preservation, this principle never leads to dogmatism. For to know that the consequences of forming a belief-protecting resolution might be better than those of not

²² Similar reasoning also establishes the converse entailment when you believe that *A* and *B* might not be equally good.

forming it, you need to justifiably believe that they might be better. But then you need to justifiably believe that you might discover counterevidence strong enough to defeat the belief. But then, by Preservation, the belief isn't justified, and thus isn't knowledge.

Without Preservation, this principle generates dogmatism. But we can replace it with:

Conditional Knowledge Dominance If you know that B's consequences would be better than A's if they're not equally good, you should choose B over A.

We then observe that, since discovering that the suitcase weighs more than 17.5 kg would defeat your knowledge that it weighs less than 18 kg, you do not know that, if it weighs more than 17.5 kg, it still weighs less than 18 kg. So Conditional Knowledge Dominance supports no dogmatic resolution.

There is, however, a minor wrinkle. Instead of Weak Knowledge Dominance, one might suggest that the proper analogue of Weak Dominance is:

Weak Knowledge Dominance* If you know that B's consequences would be at least as good as A's, and for all you know they would be better, you should choose B over A.

Weak Knowledge Dominance* can support dogmatism even given Preservation. For suppose that, while you believe that the additional evidence will further support p , it in fact tells strongly against p . Preservation then allows you to believe and know p , but—in fact, and thus for all you know—you will lose your true belief by engaging with the evidence.

However, Weak Knowledge Dominance* is implausible for closely related reasons. Suppose that you have independent reports that Chelsea and Everton each won their most recent match. In both cases, the report would ordinarily suffice for knowledge, and Chelsea generally win more often, so your overall reasons for believing that Chelsea won are slightly stronger. This time, however, Chelsea lost their match and your report about them is mistaken; so while you know that Everton won their match, you do not know that Chelsea won theirs. In this case, Weak Knowledge Dominance* says that, when choosing between a bet that Everton won and an analogous bet that Chelsea won, you have to choose the former. This strikes me as implausible. Weak Knowledge Dominance does not share this prediction, and so is independently more attractive.

Objection 2. In discussing Conditional Dominance, you assume that you justifiably believe the indicative conditional 'if p then q ' only if discovering p would not defeat your justification for believing q . But this is controversial—for example, it is false if indicative conditionals are just material conditionals.

We can separate the assumption into two parts. The first is that you justifiably believe the conditional 'if p then q ' only if you justifiably believe q on the hypothesis that p , where the latter attributes a conditional or suppositional belief rather than a belief with a conditional as its content. The second is that you are entitled to believe q on the hypothesis that p only if discovering p would not require you to stop believing q .

Defending the first assumption would require engaging with the vast literature on the semantics of indicative conditionals, which I cannot do here. However, we can avoid this assumption simply by rephrasing everything in terms of conditional belief throughout, including in Conditional Dominance. The resulting discussion is harder to parse; but it is, I think, no less compelling.

By contrast, the second assumption is essential. Fortunately, it is easier to defend. The standard worries about it are based on examples like the following: on the hypothesis that your sister is a master spy, you believe that you will never find out; but if you were to discover that she is a master spy, you would not conclude that you will never find out. (See van Fraassen (1980), who attributes these to Thomason.) But these are cases where you don't *just* discover the antecedent, but also something else, in this case, that you did find out. So they cast no doubt on a more careful version of the needed claim: that if discovering p and *nothing else relevant to q* would require you not to believe q , you're not entitled to believe q on the hypothesis that p .

Objection 3. As far as your accuracy about p is concerned, willingness to consider further evidence is not just permissible but required. Yet the solution discussed here can't deliver this stronger verdict. For example, if you believe that the evidence out there wouldn't change your mind, you will have no reason to gather it.

There is a weak and a strong version of this complaint. The strong version says that the solution is incompatible with the claim that willingness to consider further evidence is always required. The weak version says only that the solution doesn't predict this claim.

The strong version, implicit in the final sentence above, is mistaken. It assumes that we should accept something like

Arbitrary If you believe that A and B would lead to equally good outcomes, you may choose either.

But Arbitrary does not follow from Weak Dominance or Conditional Dominance. In fact, it conflicts with Conditional Dominance in cases described in §4, in which you believe that two options would have equally good outcomes but nonetheless have views about which would be better if they are not equally good. Conditional Dominance gets those cases right; so Arbitrary should be rejected.

To assess the weak version of the complaint, distinguish two cases. One concerns your dealings with evidence that will definitely be conclusive either way. Conditional Dominance will always recommend consulting such evidence. For you will believe both (a) that if the evidence makes a difference, it will lead you to believe not- p and (b) that if the evidence makes a difference (that is, if it is conclusive evidence against p), p is definitely false. So you will believe that if the evidence makes a difference to your accuracy about p , it will improve it, by replacing a false belief with a true one. So even the weak version of the complaint does not apply in these cases.

The weak version does, however, apply when you think that, if the evidence makes a difference, it will make a moderate one, leaving you believing neither p nor not- p . For then you will not believe that p is definitely false, conditional on the evidence making a difference—rather, you will be unsure. And so you will be equally unsure about whether giving up your belief that p will make things better, accuracy-wise.

However, this is just a general problem of how all-out beliefs can feature in decision-making when we lack outright beliefs about outcomes—how they can explain, for example, why you should bet on the horse that is most likely to win when you don't outright believe that it will. Any solution to that problem will generalize to this one: for surely, if you are uncertain about p conditional on some hypothesis, you should also think that if that hypothesis is true, being uncertain about p is, on

balance, a better bet—accuracy-wise—than believing p .²³ So while the response by itself doesn't predict that openness to further evidence is always required, a natural development of it will.

Objection 4. Like Harman's solution, which it extends, this solution applies only to evidence against p that is strong enough to defeat your justification for believing p . But won't there still be a problematic argument for ignoring or avoiding weaker evidence against p ?²⁴

To give an argument for ignoring or avoiding such evidence, we need to identify an adverse consequence to accommodating it when p is true. This can't be that you'll give up your belief in p . But it might be that you (a) lower your credence in p or (b) lower the resilience of your belief that p to further counterevidence.

For (a) to be an option, we need to assume that credences exist, and that outright belief is compatible with less than maximal credence. That view faces a prior, independent problem. For p entails that a non-maximal credence in p is less accurate than a maximal one. So when you believe p while having a non-maximal credence in it, you believe that your current doxastic state is less accurate than some identifiable alternative. How could such a self-undermining state be rational? Natural responses to this problem—such as denying that higher credences in truths are always more accurate, or maintaining that we mustn't use outright beliefs when reasoning about our credences—will immediately block any dogmatic argument based on (a) as well. So the argument based on (a) poses no *new* puzzle. So I will set it aside.²⁵

Appealing to (b) does make for a new puzzle. Fortunately, our solution extends to it. For lowering the resilience of one's belief to counterevidence does not, by itself, change one's accuracy about p ; it only does so if and when one gets the relevant counterevidence in the future. So the hypothesis that it makes a difference whether one accommodates the evidence isn't just the hypothesis that the evidence is weak evidence against p , but rather the hypothesis that the evidence is weak evidence against p that is later followed by further evidence against p , which together with the first evidence is strong enough to defeat one's justification. But then we can re-run our objection to the dogmatic reasoning: if Preservation is true, you should believe that this won't happen; and at any rate, you also believe that, if it does happen, p might well not be true, so that losing your belief that p needn't make you less accurate.

6. Conclusion

I have argued that we can resolve the dogmatism puzzle by focusing on Preservation and Conditional Dominance. Since this solution makes no appeal to the chance of being wrong, it is available to fallibilists and infallibilists alike.

This solution offers a more minimal, and thus more illuminating, diagnosis of where the dogmatic argument goes wrong than its more theoretically loaded

²³ This is an all-out belief version of the kind of 'Immodesty' principle discussed in the accuracy literature. For defence, see e.g. Oddie 1997, Greaves and Wallace 2006, and Joyce 2009.

²⁴ Lasonen-Aarnio (2014) develops this objection to Harman's solution.

²⁵ See also Fraser *forthcoming*. In fairness, Lasonen-Aarnio highlights that those who endorse knowledge defeat usually assume that knowledge and outright belief are consistent with non-maximal probability: this means that they can't simply dismiss this as someone else's problem. However, probabilistic models are not the only models of defeat. For example, the normality model developed by Smith (2016) and Goodman and Salow (2018) strikes me as equally promising, and is consistent with rejecting credences or requiring credence 1 for knowledge or belief.

competitors. This is clearest when comparing to the classic fallibilist, who simply rejects Weak Dominance in cases like Suitcase while offering nothing to put in its place—a response which, whatever the larger merits of the view, is much more radical than what’s required.

The same applies when comparing to the fallibilist who accepts pragmatic encroachment. This theorist can, admittedly, uphold Weak Dominance while rejecting Preservation: for they can hold that, in Suitcase, your justification to believe that the suitcase weighs less than 18 kg disappears when you get the option to check whether it weighs more than 17.5 kg. One worry is that this isn’t a particularly natural verdict even about this example. But a more important concern is that it isn’t clear what contains the occurrence of pragmatic encroachment to these relatively rare counterexamples to Preservation. For even in cases like Newspapers, the small probability of being wrong may significantly affect your preferences over various actions (such as reading another paper, or steeling yourself against changing your mind). On standard theories of encroachment, that will be enough to make you lose your knowledge and belief in these cases too. Clearly, the result is significantly more revisionary than replacing Weak Dominance with a principle that shares its non-dogmatic predictions and looks, at first sight, completely equivalent.

The solution set out here is, therefore, the most targeted one available. Abstracting from the details, it suggests that remaining open to evidence isn’t wise because you should always think there is a *tiny* chance you’re wrong. Rather, it’s wise because you should always think that, if the evidence you would discover tells strongly against your view, there is a *good* chance you’re wrong.²⁶ Emphasizing that Conditional Dominance doesn’t license dogmatism is the minimal way of making this point: it thus provides the best explanation of where the dogmatic reasoning goes wrong.

Acknowledgments

For helpful feedback and discussion, I thank Bob Beddor, Kevin Dorst, Rachel Fraser, Ginger Schultheis, Paulina Sliwa, Jack Spencer, Bob Stalnaker, and Roger White, the editor and referees, and audiences at MIT, Princeton, and the Universities of Cambridge, Glasgow, Kent, and Oxford.

Disclosure Statement

No potential conflict of interest was reported by the author(s).

ORCID

Bernhard Salow  <http://orcid.org/0000-0002-4437-8211>

References

Alchourrón, Carlos, Peter Gärdenfors, and David Makinson (1985) ‘On the Logic of Theory Change: Partial Meets Contraction and Revision Functions’, *Journal of Symbolic Logic* 50: 510–30. doi:10.2307/2274239.

²⁶ In fact, I think that this is also the thought behind the fallibilist arguments of Good (1967) and Oddie (1997); see Salow [forthcoming](#).

- Beddor, Bob (2019) 'The Toxin and the Dogmatist', *Australasian Journal of Philosophy* **97**: 727–40. doi:10.1080/00048402.2018.1556309.
- Bovens, Luc (1995) "'P and I will believe that not-P": Diachronic Constraints on Rational Belief', *Mind* **104**: 737–60. doi:10.1093/mind/104.416.737.
- Briggs, R A (2019) 'Normative Theories of Rational Choice: Expected Utility', in Edward Zalta, ed., *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University. <https://plato.stanford.edu/archives/fall2019/entries/rationality-normative-utility>.
- Clarke, Roger (2013) 'Belief is Credence 1 (in Context)', *Philosopher's Imprint* **13**: 1–18.
- Das, Nilanjan (2023) 'The Value of Biased Information', *British Journal for the Philosophy of Science* **74**: 25–55. doi: 10.1093/bjps/axaa003.
- Dorr, Cian, Jeremy Goodman, and John Hawthorne (2014) 'Knowing Against the Odds', *Philosophical Studies* **170**: 277–87. doi:10.1007/s11098-013-0212-9.
- Dorst, Kevin (2019) 'Lockeans Maximize Expected Accuracy', *Mind* **128**: 175–211. doi:10.1093/mind/fzx028.
- Easwaran, Kenny (2014) 'Regularity and Hyperreal Credences', *Philosophical Review* **123**: 1–41. doi:10.1215/00318108-2366479.
- Fantl, Jeremy and McGrath, Matthew (2002) 'Evidence, Pragmatics, and Justification', *Philosophical Review* **111**: 67–94. doi:10.1215/00318108-111-1-67.
- Fraser, Rachel (forthcoming) 'The Will in Belief', *Oxford Studies in Epistemology*. Oxford University Press.
- Good, I J (1967) 'On the Principle of Total Evidence', *British Journal for the Philosophy of Science* **17**: 319–21. doi:10.1093/bjps/17.4.319.
- Goodman, Jeremy and Bernhard Salow (2018) 'Taking a Chance on KK', *Philosophical Studies* **175**: 183–96. doi:10.1007/s11098-017-0861-1.
- Goodman, Jeremy and Bernhard Salow (2023) 'Epistemology Normalized', *Philosophical Review* **132**: 89–145. doi:10.1215/00318108-10123787.
- Goodman, Jeremy and Bernhard Salow (forthcoming) 'Belief Revision Normalized', *Journal of Philosophical Logic*.
- Greaves, Hilary and David Wallace (2006) 'Justifying Conditionalization: Conditionalization Maximizes Expected Epistemic Utility', *Mind* **115**: 607–32. doi:10.1093/mind/fzl.
- Greco, Daniel (2015) 'How I Learned to Stop Worrying and Love Probability 1', *Philosophical Perspectives* **29**: 179–201. doi:10.1111/phpe.12059.
- Hall, Ned (1999) 'How to Set a Surprise Exam', *Mind* **108**: 647–703. doi:10.1093/mind/108.432.647.
- Harman, Gilbert (1973) *Thought*. Princeton University Press.
- Hawthorne, John (2004) *Knowledge and Lotteries*. Oxford University Press.
- Joyce, James (2009) 'Accuracy and Coherence: Prospects for an Alethic Epistemology of Partial Beliefs', in Franz Huber and Christoph Schmidt-Petri, eds., *Degrees of Belief*: 263–97. Springer.
- Kripke, Saul (2011) 'Two Paradoxes of Knowledge', in *Philosophical Troubles*: 27–51. Oxford University Press.
- Lasonen-Aarnio, Maria (2014) 'The Dogmatism Paradox', *Australasian Journal of Philosophy* **92**: 417–32. doi:10.1080/00048402.2013.834949.
- Lin, Hanti (2019) 'Belief Revision Theory', in Richard Pettigrew and Jonathan Weisberg, eds., *The Open Handbook of Formal Epistemology*: 349–96. The PhilPapers Foundation. <https://philarchive.org/rec/PETTOH-2>.
- Oddie, Graham (1997) 'Conditionalization, Cogency, and Cognitive value', *British Journal for the Philosophy of Science* **48**: 533–41. doi:10.1093/bjps/48.4.533.
- Roeber, Blake (2016) 'The Pragmatic Encroachment Debate', *Nous* **52**: 171–95. doi:10.1111/nous.12156.
- Smith, Martin (2016) *Between Probability and Certainty*. Oxford University Press.
- Salow, Bernhard (2018) 'The Externalist's Guide to Fishing for Compliments', *Mind* **127**: 691–728. doi:10.1093/mind/fzw029.
- Salow, Bernhard (forthcoming) 'The Value of Evidence', in Maria Lasonen-Aarnio and Clayton Littlejohn, eds., *The Routledge Handbook of the Philosophy of Evidence*. Routledge.
- Sorensen, Roy (1988) 'Dogmatism, Junk Knowledge, and Conditionals', *Philosophical Quarterly* **38**: 433–35. doi:10.2307/2219708.
- Sorensen, Roy (2018) 'Stealing Harman's Thought: Knowledge Saboteurs and Dogmatists', *Synthese* **198**: 1787–99. doi:10.1007/s11229-018-01945-5.

- Strevens, Michael (2008) *Depth: An Account of Scientific Explanation*. Harvard University Press.
- van Fraassen, Bas (1980) 'Review of Brian Ellis, Rational Belief Systems', *Canadian Journal of Philosophy* **10**: 497–511. doi:10.1080/00455091.1980.10715741.
- van Fraassen, Bas (1984) 'Belief and the Will', *Journal of Philosophy* **81**: 235–56. doi:10.2307/2026388.
- Weatherson, Brian (2005) 'Can we do Without Pragmatic Encroachment?', *Philosophical Perspectives* **19**: 417–43. doi:10.1111/j.1520-8583.2005.00068.x.
- Williamson, Timothy (2000) *Knowledge and its Limits*. Oxford University Press.
- Ye, Ru (2016) 'Misleading Evidence and the Dogmatism Puzzle', *Australasian Journal of Philosophy* **94**: 563–75. doi:10.1080/00048402.2015.1091485.