Belief Revision Revised

Abstract. I outline a novel counterexample to the plausible principle of belief revision, Anticipation: if you would not be justified in believing *p* were you to learn *e*, and you would not be justified in believing *p* were you to learn not-*e*, you cannot now be justified in believing *p*. If I'm right, not only is the leading theory of belief revision false (Alchourrón, Gärdenfors, and Makinson 1985), so are various proposed weakenings, such as those in (Leitgeb 2017), (Goodman and Salow 2018) and (Lin and Kelly 2012). While Goodman and Salow (2021) (ms) have more recently defended a theory which predicts different counterexamples to Anticipation, their theory fails to predict the counterexample I outline. I defend an alternative theory that makes all of the correct predictions on this issue. Endorsing this theory—or any other theory that rejects Anticipation—has radical consequences; for instance, it means either that it can be rational to refuse free evidence before making a decision, or else that justified beliefs do not play an important role in rational decision making.

Belief revision theory concerns when and how one's beliefs should change upon learning new information. One extremely plausible idea is as follows.

> **Broken Brakes.** A cycling trip is planned for tomorrow, but you've just discovered the brakes on your bike are broken. On considering whether to take it to the local repair shop, you realise the following. Were you to learn that the repair shop is closed, you'd be justifiably unsure whether you'd be able to make the trip—while you might be able to borrow a bike from a friend, you also might not. At the same time, were you to learn that that the local repair shop is *open*, you'd again be justifiably unsure whether you'd be able to make the trip—there's no way to tell whether the shop will have the parts necessary for a quick fix.

Can you, nevertheless, *now* be justified in believing that you'll make the trip? It doesn't seem like it. There is a proposition *e*—*the repair shop is open*—such that no matter whether you were to learn it or its negation, you'd fail to be justified in believing you'll make the trip. Plausibly, this equivalent effect of learning *e* or of learning not-*e* ought to be anticipated, preventing you from justifiably believing you'll make the trip.

This idea is codified by the widely endorsed principle of belief revision, Anticipation. Anticipation is entailed by the dominant theory of belief revision, AGM (Alchourrón, Gärdenfors, and Makinson 1985), as well as by various proposed weakenings of AGM, such as those in (Leitgeb 2017), (Goodman and Salow 2018), and (Lin and Kelly 2012). This wide endorsement of Anticipation is unsurprising. It makes very plausible predictions in simple cases. Moreover, as we'll see, implausible consequences follow if Anticipation is false: bizarre assertions involving indicative conditionals are licensed, and counterexamples arise for the platitude that one should always accept free evidence before making a decision. Despite all this, my goal here is to argue that Anticipation is, in fact, *false*, and to outline a theory of belief revision that properly respects this fact.

Here's the plan. §1 outlines and motivates Anticipation in more detail. §2 presents my central case against it. §2.1 outlines a known and widely accepted counterexample to the logically stronger principle Preservation. §2.2 argues that anyone who finds this counterexample to Preservation persuasive should also find the counterexample to Anticipation I outline persuasive. §3 considers how recent theories of belief revision interact with this counterexample to Anticipation. §3.1 shows that while Lin and Kelly's (2012) approach and Goodman and Salow's (2018) earlier approach can successfully predict the counterexample to Preservation, they cannot predict my counterexample to Anticipation. §3.2 considers Goodman and Salow's (2021) (ms) more recent approach and shows that while it can predict *different* counterexamples to Anticipation, it nevertheless fails to accommodate my counterexample. §4 defends an alternative theory of belief revision that makes all of the correct predictions on this issue.

In keeping with the literature, I'll be making two assumptions. First, one's justified beliefs play a role in rational decision making that cannot be reduced to the role played by credences; for example, as premises in practical reasoning— if one justifiably believes that P then one is justified in reasoning from the premise that P.[1] Second, justified beliefs are closed under deduction: if one has justification to believe $P_1,...P_n$, and $P_1,...P_n$ jointly entail Q, then one has justification to believe Q.[2] These assumptions rule out simple Lockeanism, the view that you have justification to believe P iff the probability that P conditional on your evidence is sufficiently high.[3] You might read this paper as a reductio of my assumptions. This is a debate for another time. But note that rejecting my assumptions would itself be a significant result.

[1](Leitgeb 2017), (Kelly and Lin 2021), (Fantl and McGrath 2009), (Comesaña 2020), (Hawthorne and Stanley 2008).

[2](Alchourrón, Gärdenfors, and Makinson 1985), (Goodman and Salow 2018), (Goodman and Salow 2023), (Leitgeb 2017), (Lin and Kelly 2012).

[3]Simple Lockeanism nevertheless occupies some interesting dialectical space in this paper; see §2.2, objection 4.

1 Anticipation

Here's Anticipation:[4]

Anticipation
If one would not be justified in believing p were one to learn that e as total information, and one would not be justified in believing p were one to learn not-e as total information, one cannot *now* be justified in believing p.

Three clarifications. First, I follow the literature on belief revision in stating these principles in the subjunctive mood: they concern what one *would* believe *were* one to learn new information.[5] This is arguably not quite right, as it invites counterexamples not usually of interest to those studying belief revision. If Dr. Evil ensures that, regardless of whether I'll learn e or learn not-e, I'll have my memories which justify my belief that p erased, then we have a counterexample to Anticipation as stated: I'm justified in believing p, but wouldn't be were I to learn e or were I to learn not-e. For that reason, it's better to understand such principles in terms of *conditional belief*: whether, on hypothetically adding e to your body of evidence, you'd be justified in believing p.[6] However, since this technical notion of conditional belief is less familiar, and since the subjunctive gloss will serve well for our purposes, I shall stick with it, setting problem cases of this kind aside.

Second, although I will often informally drop it, the qualification that the principle only concerns propositions that are learned "as total information" is crucial for avoiding further counterexamples. Consider:

> **Marmite.** You justifiably believe you'll never learn whether you like Marmite. However, were you to learn that you like Marmite (say, by tasting it), you'd be justified in believing that you've learned whether you like Marmite, and were you to learn that you *don't* like Marmite, you'd again be justified in believing that you've learned whether you like Marmite.

There is a proposition e—*you like Marmite*—such that no matter whether you learn it or its negation, you'd fail to be justified in believing proposition p–*you'll never learn whether you like Marmite*. Since you are *now* nevertheless justified in believing p, don't we have a counterexample to Anticipation? No. The case in which you learn you like Marmite

[4] I take the name and principle from a previous draft of (Goodman and Salow 2023). Goodman and Salow (ms) later discuss a generalisation of this principle under the name '$\Pi-$'. (Kraus, Lehmann, and Magidor 1990) and (Freund and Lehmann 1996) discuss the principle under the name 'Negation Rationality'.

[5] E.g: (Gärdenfors 1986), (Huber 2013), (Leitgeb 2017), (Lin 2019) and (Goodman and Salow ms).

[6] Cf. (Ramsey 1926, p. 247) who famously uses this notion of conditional belief in his discussion of indicative conditionals.

by tasting it is a case in which you learn more than just *e*, you also learn the stronger proposition that *you've learned whether e*. Let's suppose your total information in this case is characterised by proposition e'. To properly assess whether we have a counterexample to ANTICIPATION, the further case to consider is not one in which you learn not-*e*, but rather a case in which you learn as total information not-e'. Since e' entails that you've learned whether you like Marmite, not-e' is compatible with scenarios in which you *haven't* learned whether you like Marmite. So it is not at all clear that learning not-e' gives you reason to give up your belief in *p*, meaning we have no counterexample to ANTICIPATION.

Third, the notion of justification at issue is what is referred to as 'propositional' justification, rather than 'doxastic'.[7] However, I often use the locution 'justified in believing' rather than 'have justification to believe' due to naturalness. Accordingly, I will be assuming that the epistemic agents at issue form all and only the beliefs they have justification to, in a way that is sufficient for those beliefs to be justified.

Why accept ANTICIPATION? Beyond its considerable intuitive appeal, we can give two further motivations. First, failures of ANTICIPATION license bizarre assertions, given two plausible ideas: (i) one who would not be justified in believing *p* were one to learn *e* (as total information) is accordingly not justified in believing, and rather should doubt, the conditional 'If *e*, *p*';[8] and (ii) one is epistemically permitted to assert those propositions one is justified in believing.[9] If both ideas are right, then failures of ANTICIPATION license bizarre assertions. Suppose that ANTICIPATION *fails* in **Broken Brakes**: you are *now* justified in believing you'll make the trip tomorrow, even though you wouldn't be either on learning that the repair shop is open or upon learning that it is closed. Then you'll be in a position to assert the highly infelicitous: "I'm not sure whether I'll make the trip if the bike shop is closed. I'm also not sure whether I'll make the trip if the bike shop is open. Nevertheless, I'll make the cycling trip!"

Second, if ANTICIPATION is false, serious doubts emerge concerning the thesis that one should, if given the opportunity, always look at free evidence before making a decision. Although this idea has not gone unquestioned, counterexamples to it have so far required agents that are risk-averse, as in (Buchak 2010), or agents that fail to know what their evidence is, as in (Salow and Ahmed 2019). The falsity of ANTICIPATION puts pressure

[7] See (Silva and Oliveira forthcoming) for recent discussion.

[8] We may be tempted by an even stronger principle: one has justification to believe the conditional 'If *e*, *p*' iff one would have justification to believe *p* were they to learn *e* as total information. I won't make this stronger assumption here; doing so requires care concerning triviality results (Gärdenfors 1986). I am sympathetic to contextualist replies to these triviality concerns, see (Lindström 1996) and (Mandelkern and Khoo 2019).

[9] A widely endorsed thesis, sometimes referred to as 'Entitlement Equality'. Philosophers who disagree on the norm of assertion, such as (Williamson 2000) and (Lackey 2008), tend to nevertheless agree on Entitlement Equality. But see (Hawthorne, Rothschild, and Spectre 2016) for objections to it.

on this claim even without assuming that rational agents can be risk-averse or ignorant of their own evidence. Suppose you justifiably believe that a yet undetonated wartime bomb found at the local park won't go off; at the same time, there exists some proposition *e*, such that, were you to learn *e*, or were to learn not-*e*, you wouldn't be justified in believing that bomb won't detonate. Should you take up a free opportunity to learn whether *e*? It's hard to see why. Learning whether *e* will will (assuming you're rational) have significant practical consequences, such as evacuating the local area. Yet from your current perspective, such actions are completely unnecessary; after all, the bomb will not go off (or so you think).[10]

2 Against Anticipation

Anticipation is a highly plausible, well-motivated principle. Despite this, I will now argue that Anticipation is false. I begin in §2.1 by outlining known counterexamples to a strictly stronger principle of belief revision, Preservation. In §2.2 I argue that proper appreciation of the intuitions grounding these counterexamples to Preservation forces us to accept counterexamples to Anticipation, too.

2.1 Preservation Failure

Let's begin, then, with Preservation—a consequence of the dominant theory of belief revision AGM (1985) and centrepiece of Leitgeb's (2014; 2017) recent theory:

> **Preservation**
> If one is justified in believing *p* and one is justified in leaving *e* open, then one would still be justified in believing *p* were one to learn that *e* as total information.[11]

Preservation possesses some intuitive plausibility. For instance, if learning that the repair shop is closed would defeat your justification for believing you'll make the trip tomorrow, then, plausibly, you can only be justified in believing you'll make the trip tomorrow if you are also justified in believing that the repair shop isn't closed.

[10]This challenge is similar to Kripke's second dogmatism puzzle (Kripke 2011). However, the problem presented by Anticipation failure is harder than Kripke's puzzle—solutions to the latter can't obviously be applied to the former. For example, (Carter and Hawthorne forthcoming) propose to solve Kripke's puzzle by noting that (substituting "knows" for "justifiably believes"), although looking at the evidence may risk losing a justified belief that *p*, it may also result in obtaining a higher-order justified belief that one justifiably believes *p*. However, if Anticipation fails, one is *guaranteed* to lose the justified belief in question if one looks at the evidence, meaning Carter & Hawthorne's solution cannot be applied here.

[11]One leaves *e* open just in case one does not believe not-*e*.

Nevertheless, Preservation faces decisive counterexamples. I will mainly focus on the following, simple case. It was originally deployed in (Dorr, Goodman, and Hawthorne 2014) to argue against the KK principle, but has since been used to argue against Preservation in (Stalnaker 2019, ch. 8), (Goodman and Salow 2021) and (Goodman and Salow 2023):[12]

> **Flipping for Heads.** In front of us is a coin you know to be fair. I am going to flip it until it lands heads or has otherwise been flipped 1,000 times. You know all of this. Once I am done, I will have produced a sequence of tails, followed by a heads, or a sequence of 1,000 tails.

Plausibly, you are justified in believing that the coin won't be flipped all 1,000 times; equivalently, that the coin will not land tails 999 times in a row. At the same time, you are clearly *not* in a position to rule out the coin landing tails at least once.[13] Indeed, you are likewise not in a position to rule out the coin landing tails twice, three times, and so on. Since you *are* justified in ruling out the coin landing tails 999 times, we will eventually reach some number of tails, k such that whilst you are not justified in ruling out the coin landing tails k times, you are justified in ruling out the coin landing tails $k + 1$ times. For concreteness, let's suppose that k is equal to 20, meaning you are in the following situation:[14]

(i) You are not justified in believing that the coin will not land on tails 20 times.

(ii) You are justified believing that the coin will not land on tails 21 times.

Trouble for Preservation arises on considering what you should believe upon learning that the first flip lands on tails. If Preservation is right, since you left it open that the first flip would land tails, learning this has happened should not change any of your beliefs; in particular, you should still believe that the coin will not land on tails 21 times. This is counter-intuitive. It means that while your beliefs about how many tails there will be *in total* has not changed, there has been a shift in your *de se* beliefs about how many *more* tails there will be from your present moment: at the start of the experiment you thought there would be at most 20 *more* tails, yet after seeing the first flip land tails, since you still think there will be at most 20 tails *in total*, you now think that there will be at most 19 *more* tails. This is objectionable: assuming your knowledge that the coin is fair

[12] Stalnaker (2019, ch. 8) remains sympathetic to Preservation, but does not offer a full account of how to deal with counterexamples of this kind.

[13] One is justified in ruling out that p just in case one is justified in believing not-p.

[14] You may suspect that the value of k is vague. In §2.2, I argue that this observation cannot be leveraged in support of Preservation/Anticipation (see "Objection 2".

remains upon seeing the first coin land tails, you have just as much reason to think there will be 20 more tails at the start of the experiment as you do after the first flip lands on tails. Analogously, if instead of seeing the first flip, you are instead told that the coin was flipped once *yesterday* and landed tails, this would give you no reason to change your beliefs about what sequence the coin might produce from *now*. To reason as PRESERVATION recommends therefore smacks of the gambler's fallacy: after seeing the first flip land tails, you now believe that the first heads is guaranteed to come sooner than you did before—as if a heads is now somehow more "overdue"—despite the fact that your beliefs about the bias of the coin have not changed.

More specifically, PRESERVATION goes wrong in violating the following principle:

> **NO GAMBLER'S FALLACY**
> If agent A_1 knows at t_1 that coin C_1 is fair, and agent A_2 knows at t_2 that coin C_2 is fair, then A_1 is justified in believing that C_1 will not produce sequence X from being flipped after t_1 iff A_2 is justified in believing that C_2 will not produce sequence X from being flipped after t_2.[15]

PRESERVATION is inconsistent with NO GAMBLER'S FALLACY: treating you before and after the first coin flip as distinct agents, PRESERVATION says that you-before the-first-flip cannot believe the coin won't land tails 20 more times whereas you-after-the-first-flip *can*. Contra PRESERVATION, then, it seems that upon seeing the first flip lands tails in **Flipping for Heads**, your beliefs about how many tails there might be should shift by one from (i) and (ii): you should now leave open the coin landing tails 21 times, and rule out it landing tails 22 times.

2.2 ANTICIPATION Failure

If we found the above argument against PRESERVATION convincing, the following example should also convince us that ANTICIPATION is false.

> **Flipping for Both.** In front of us is a coin you know to be fair. I am going to flip it until it lands on heads at least once and on tails at least once, or until I have otherwise flipped it 1,000 times. You know all of this. Once I am done, I will have produced either a sequence of heads followed by a tails, a sequence

[15] The principle will do as stated for our purposes, although it is imperfect: an agent who observes close to all 1,000 possible flips in **Flipping for Heads** can justifiably rule out particularly long sequences being produced that an agent who has only observed a couple of flips can't. It's tricky to state exactly what the correct restricted principle should be, but it would likely consider only those beliefs about possible future sequences that agents have formed on a purely inductive basis from facts about the bias of the coin.

of tails followed by a heads, a sequence of 1,000 repeating heads, or a sequence of 1,000 repeating tails.

Plausibly, the same two facts apply in this case for the same reason they did in **Flipping for Heads**:

(i) You are not justified in believing that the coin will not land on tails 20 times.

(ii) You are justified believing that the coin will not land on tails 21 times.

Symmetric considerations apply concerning the number of *heads* you are justified in believing may occur, giving us:

(iii) You are not justified in believing that the coin will not land on heads 20 times.

(iv) You are justified believing that the coin will not land on heads 21 times.

Given that justified beliefs are closed under deduction, (ii) and (iv) further entail that:

(v) You are justified in believing that the coin will not land the same way 21 times.

We now have all the required materials to argue against Anticipation.

First, consider what would happen were you to learn that the first flip has landed tails. This case is no different to what would happen in an analogous scenario in **Flipping for Heads**: to maintain your belief outlined in (ii), that the coin will not land tails 21 times, would violate No Gambler's Fallacy. Instead, you should now leave it open that the coin will lands tails 21 times. This, in turn, means revising your belief outlined in (v) that the coin will not land the same way 21 times.

Second, consider what would happen were you to learn that the first flip has landed *heads*. Since this case is symmetric to the case in which the first coin land tails, symmetric conclusions apply: learning that the first flip has landed heads should result in you leaving open that the coin will land heads 21 times, which means revising your belief outlined in (v): that the coin will not land the same way 21 times.

And now we have a counterexample to Anticipation. Initially, you are justified in believing that the coin will not land the same way 21 times in a row. But you would not be justified in believing this were you to learn that the first flip has landed tails, and you would not be justified in believing this were you to learn that the first flip has landed heads. So Anticipation is false.

This argument ought to be completely convincing to anyone who is convinced by the argument against Preservation in §2.1. My argument against Anticipation does not bring with it new substantive commitments not already present in the argument against Preservation.

Before moving on, I'll reply to various objections to my arguments so far.

Objection 1: Lotteries

Objection: "Your arguments go wrong at the very first step: one is not permitted to believe the coin will eventually land heads. To do so is to form a belief analogous to a belief that your lottery ticket will be a loser."

Reply: As Hawthorne et al. (2014) argue, this reply is skeptical. Many of our ordinary beliefs that we take to be justified arguably have a structure sufficiently similar to the above coin-flipping cases to generate problems for Preservation and Anticipation. Consider your beliefs about the weather over the next month. Plausibly, you are justified in believing it will rain at some point. Perhaps your strongest justified belief is that it will rain at some point over the next two weeks. However, upon learning tomorrow that it has not rained, we can construct a similar argument against Preservation as above: if tomorrow you ought to have the same *de se* beliefs as today—there will at most be no more rain for two weeks—then tomorrow there should be a revision in your beliefs: you should no longer believe there will not be rain for, in total, 15 days in a row. And it is not too hard to see how to extend this into an argument against Anticipation: just consider in addition your strongest justified belief about how many days *with* rain there might be in a row; say it's 7 days. If you learn tomorrow that it has rained, you arguably ought to extend this prediction to 8 days. So, no matter what you learn about the weather tomorrow, you'll have to give up your belief that there will be at most 14 days without rain from now and at most 7 days of consecutive rain from now.

There is a general formula here: we are often justified in believing some process will eventually produce a certain output O. Yet, as time proceeds, we can remain equally justified in our *de se* beliefs regarding how quickly O will occur *from our present moment*. It is exactly cases with this structure that causes trouble for Preservation and Anticipation, in the ways outlined above. But since these beliefs are common place—you believe that not every paper you grade in this next batch will be a C; that your partner will be home from work at some point over the next two hours; that at least one grape on this bunch will be delicious; etc—denying that they are justified is skeptical.

Objection 2: Vagueness

Objection: "The arguments objectionably exploit assumptions that are not plausible once we accept there's vagueness about the boundaries of one's beliefs. For instance, in **Flipping for Heads**, there *is no precise* k for which k is the smallest number that you are justified in believing there will not be k heads in a row. The boundary is imprecise."

Reply: Perhaps that's right. But I have a hard time seeing how to leverage this observation in support of Preservation and Anticipation. The fact that k is vague gets the result that the above counterexamples are harder to *identify* than I have claimed. But there's a gulf from that conclusion to the further conclusion that Preservation and Anticipation are nevertheless *true*. Consider the dominant approach to vagueness, Supervaluationism (Fine 1975), on which a claim is true iff it is true on every single admissible precisification of its vague terms. In order for Preservation (for example) to be true, it will have to hold under every single precisification of the vague term 'believes'. But the idea that one's beliefs should not violate No Gambler's Fallacy in a case like **Flipping for Heads** remains just as compelling for precise belief states as it does for our own, fuzzy belief states.

Maybe the objection is rather that, due to vagueness, we should reject classical logic, which my arguments have implicitly assumed. This would be a radical option. While I have no new objections to it, it's surely at least worth looking at alternatives, such as endorsing a theory of belief revision that instead rejects Preservation and Anticipation.

Objection 3: A New Argument for Preservation?

Objection: "I agree with your conditional claim: if we accept the argument against Preservation, we should accept your argument against Anticipation. But I apply *modus tollens* where you apply *modus ponens*: we should reject the argument against Preservation"

Reply: I have some sympathy here. After all, it is not completely obvious that we should prioritise accepting No Gambler's Fallacy over Anticipation. Both have considerable intuitive appeal. Perhaps it's No Gambler's Fallacy rather than Preservation and Anticipation that should go.

The problem is that a full endorsement of Preservation has further intolerable consequences. Suppose that, instead of seeing just the first flip of the coin land tails, you see the first *20* flips all land tails. Assuming, as we did above, that (i) and (ii) hold—you must initially leave open 20 tails in a row but can rule out 21 tails in a row—Preservation delivers an extremely worrying consequence. Since you left it open that the coin will land tails 20 times in a row, Preservation says you need not revise your beliefs, including your belief that the coin will not land tails 21 times in a row. That means Preservation predicts you can justifiably believe the next flip will land heads. That cannot be right. Absent inadmissible information (Lewis 1980), one cannot be justified in believing that a coin is fair and will land heads when next flipped.

Objection 4: A New Argument Against Closure?

Objection: "Your argument against Anticipation crucially assumes Closure: that if one justifiably believes premises $p_1, ..., p_n$, which jointly entail q, then one is in position to justifiably believe q by competently deducing it from those premises. If we deny Closure, we can agree with (ii) and (iv)—you are justified in believing there won't be 21 tails, and you are justified in believing there won't be 21 heads—yet disagree with (v)—you are justified in believing the coin will not land the same way 21 times. What we have is a new argument against Closure, not an argument against Anticipation.

Reply: I don't any new objections to views which deny Closure, such as simple Lockeanism, mentioned in the introduction.[16] But I deny that my arguments can be construed as a *new* argument against Closure.

The reason is this. Denying Closure offers us no help in responding to the first argument against Preservation; moreover, denying Closure offers us no help in defending the following, weaker principle:

Probabilistic Preservation
If one is justified in believing p, but one would not be justified in believing p were one to learn e as total information, then one is justified in taking e to be sufficiently unlikely.

No matter where we set the threshold for "sufficiently unlikely", we can construct a **Flipping-for-Heads**-type counterexample to Probabilistic Preservation. For instance, suppose we set "sufficiently unlikely" in Probabilistic Preservation to (the very lax) "less than 90% likely". Then the following case serves as a counterexample:

> **Rolling Ten Sides for One.** In front of us is a fair ten-sided die. The die is blue on one side and red elsewhere. I am going to roll until it lands on blue, or until I have otherwise rolled it 1,000 times.

Presumably, you are again justified in believing that I will not roll the die all 1,000 times, and there is thereby some smallest k such that you believe I will roll red at most k times. Using the same arguments as above, we can come to the conclusion that on rolling the die once and it landing red, you should now leave open the possibility of, in total, $k + 1$ rolls landing on red. And that means that we'll have a counterexample to Probabilistic Preservation even if "sufficiently unlikely" is set at the extremely permissive "less than 90% likely". It's not too hard to see how to create a similar argument no matter how lax we set the threshold for "sufficiently unlikely".

Here's the point. When the threshold for "sufficiently unlikely" is set low enough, anyone who is motivated to accept Anticipation should likewise be motivated to accept Probabilistic Preservation. For near-identical arguments in favour of Anticipation in

[16] See (Goodman and Salow ms, §9.1) for criticisms of simple Lockean approaches to coin-flipping cases.

§1 can be given in favour of sufficiently weak variants of Probabilistic Preservation. Consider the infelicitous assertions violations of Probabilistic Preservation would license, such as "I'm not sure if I'll make the trip if the repair shop is closed. And it is almost certainly closed. Nevertheless, I'll make the cycling trip!" Moreover, failures of lax variants Probabilistic Preservation will also motivate evidence avoidance: if you would not be justified in believing the bomb will remain undetonated were you to learn that *e*, and you are almost certain that *e*, then you again may well have good reason to refuse a free opportunity to learn whether *e*. Given that Anticipation but not Probabilistic Preservation have the same motivations, it's hard to motivate a view that gives us one but not the other. Therefore, since denying Closure will only help preserve Anticipation, but not Probabilistic Preservation, the above arguments cannot be construed as a new argument against Closure.

3 Theories of Belief Revision

I have argued that anyone who rejects Preservation due to counterexamples like **Flipping for Heads** ought also reject Anticipation due to counterexamples like **Flipping for Both**. This section discusses what recent approaches in the theory of belief revision say about these cases, and argues they are wanting.

The dominant approach to belief revision, AGM (1985), as well as Leitgeb's (2014; 2017) approach, will be set aside, since they entail Preservation and so fail to properly account even for **Flipping for Heads**. §3.1 discusses the approaches to belief revision outlined in (Lin and Kelly 2012) and (Goodman and Salow 2018). Surprisingly, these approaches predict the counterexample to Preservation from §2.1, but fail to predict the counterexample to Anticipation in §2.2. In §3.2, I discuss more Goodman and Salow's more recent theory of belief revision, outlined in (Goodman and Salow 2021) and (Goodman and Salow ms). While their more recent approach can predict *other* counterexamples to Anticipation, it struggles to offer a predict the counterexample from **Flipping for Both**. We therefore need an alternative theory, outlined in §4.

3.1 Simple Normality Theories

I'll focus on Goodman and Salow's (2018) approach, which behaves relevantly similarly to Lin and Kelly's (2012).[17] Here's the rough idea. Worlds can be compared by how normal they are.[18] One is justified in ruling out a world either when it is incompatible with one's

[17]See footnote 20 for more details on how Lin and Kelly's approach works.

[18]I follow (Goodman and Salow 2023) in using the term "normality", but more commonly these orderings are thought of in terms of "plausibility", as with (Lin 2019). In either case, these terms should be seen as

evidence, or when it is sufficiently less normal than some other world compatible with one's evidence. We then represent one's strongest justified belief with a "belief set": the set of worlds one is not justified in ruling out.

Suppose my evidence fails to eliminate just three worlds: w_1, in which it is rainy, w_2, in which is it overcast but not rainy, and w_3 in which it is sunny. However, while w_1 and w_2 are not sufficiently less normal than any other world compatible with your evidence, w_3 is sufficiently less normal than w_1. Goodman and Salow's approach says that I am justified in ruling out w_3, and so my belief set consists of just w_1 and w_2. In other words, while I should be unsure as to whether it will rain, I am justified in believing it won't be sunny.

Importantly, Goodman and Salow allow for there to be differences in normality *even among* worlds in the belief set: w_1 may be more normal than w_2 even though one is not justified in ruling out either. This allows for Preservation to fail, for if the evidence received is compatible with some worlds in the belief set, but *in*compatible with the *most normal* worlds in that set, worlds that were previously ruled out on pain of being sufficiently less normal than those most normal worlds are reintroduced. For instance, if in our toy example w_2 is not sufficiently more normal than w_3, then learning that I am not in w_1 will mean that w_3 is now included in my belief set. This means, contra Preservation, that I'll be forced to give up my belief that it isn't sunny, even though I learnt something that I thought was possible (that it isn't raining).

Let's apply these ideas to **Flipping for Heads**. First, we define our set of worlds W, individuated with respect to which sequence of heads and tails is produced in them. Letting t^k denote the world in which the coin produces an initial sequence of k tails, W consists of worlds of the form t^k, where $0 \leqslant k \leqslant 1,000$. We then impose *two* normality orderings onto W. The first, denoted $\succeq$, tells us whether one world is *at least as normal as* another. We'll say that $\mathrm{t}^k \succeq \mathrm{t}^j$ iff $k \leqslant j$; that is, one world is at least as normal as another just in case the sequence of tails produced in the former is no longer than the sequence of tails produced in the latter. The second ordering, $\ggcurly$, tells us whether one world is *sufficiently* more normal than another. To keep things simple, let's make the unrealistic assumption that $\mathrm{t}^k \ggcurly \mathrm{t}^j$ just in case $j - k > 3$. This will result in your justified beliefs in **Flipping for Heads** being far stronger than we supposed in §2, but all that matters for the following argument is how your justified beliefs are structured, which is not effected by this definition for $\ggcurly$. Finally, where one's evidence is represented by a subset of W, e, we identify one's strongest justified belief with their "belief set" B_e: the set of worlds, consistent with one's evidence, that are not sufficiently less normal than any other world consistent with one's evidence. Formally:

---

technical notions for which intuitive considerations regarding what is more or less "normal" or "plausible" than what can be overridden by more pressing considerations concerning what one is justified in believing. Cf. Lewis's (1973) use of the term "similar".

Normal Belief
$B_e = \{w \in E : \forall v \in e, \neg(v \gg w)\}$.

We can represent the corresponding model diagrammatically. An arrow points from $\mathfrak{t}^k$ to $\mathfrak{t}^j$ iff $\mathfrak{t}^k \geq \mathfrak{t}^j$, with arrows that can be inferred from reflexivity and transitivity omitted. Worlds are placed inside the dotted box iff they belong to your belief set. This model predicts that, before seeing any flips, your strongest justified belief is that the coin will land on tails no more than three times:

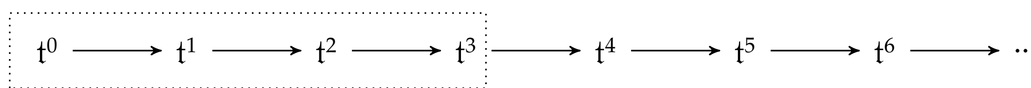

Figure 1: **Flipping for Heads**, before any flips.

To see how this model predicts the Preservation failure argued for in §2.1, note that on learning that the coin has landed tails at least 3 times, you gain evidence inconsistent with worlds $\mathfrak{t}^0$. Although this evidence is consistent with what you initially believed, your beliefs must now be revised since there are no longer worlds consistent with your evidence that are sufficiently more normal than $\mathfrak{t}^4$:

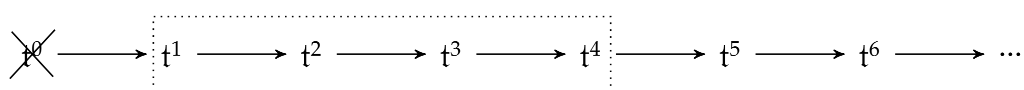

Figure 2: **Flipping for Heads**, on learning that the first flip lands tails.

Indeed—on the face of it, at least—this approach seems to vindicate No Gambler's Fallacy. So long as the coin has not yet landed on heads, you'll be justified in believing the same thing about how many *more* tails the coin may produce; in our simplified case, you'll always believe at most 3 *more* tails will occur. So far, so good! Given my arguments in §2.2, one would now expect Anticipation failures. Surprisingly, this does *not* happen.

To model **Flipping for Both**, we'll need to broaden W to include, instead of $\mathfrak{t}^0$, worlds of the form $\mathfrak{h}^k$, for $1 \leqslant k \leqslant 1,000$, which represent worlds in which the produced sequence consists of k heads in a row, followed by a tails (or 1,000 heads in a row). Keeping the normality ordering analogous to those used to model **Flipping for Heads**, for $x, y \in \{\mathfrak{h}, \mathfrak{t}\}$ we'll say that $x^k \geq y^j$ iff $k \leqslant j$, and that $x^k \gg y^j$ iff $j - k > 3$.

This model predicts that, before seeing the first flip, you are justified in believing that the coin will not land the same way more than 4 times in a row:[19]

[19]This is a surprising prediction: although we have kept the orderings defined analogously to our model for **Flipping for Heads**, in **Flipping for Heads** it was predicted that one can rule out the possibility of 4 tails followed by a heads, whereas now with respect to **Flipping for Both** it is predicted that one must leave that

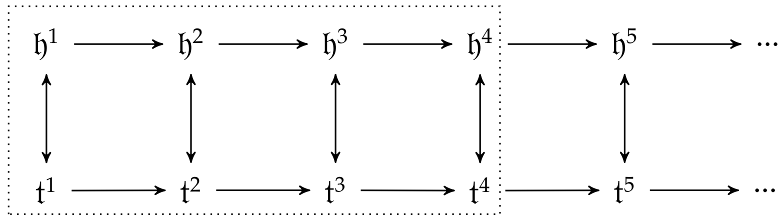

Figure 3: **Flipping for Both**, before any flips.

Matters get interesting on considering what you should believe upon learning that the first flip has landed tails (symmetric considerations apply to the case in which you learn that the first flip has landed heads). This case should be no different to **Flipping for Heads**: you should revise your strongest belief about how many tails there will be in a row. Yet this is not predicted on the current model. Even though your evidence has eliminated worlds in which the opening sequence consists of heads, your evidence has *not* eliminated $\mathfrak{t}^1$, one of the most normal worlds, meaning new worlds cannot be introduced into the belief set, and so no revision occurs:

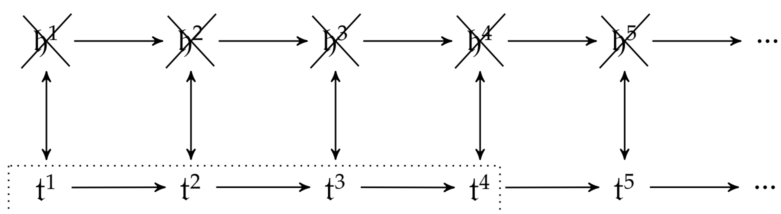

Figure 4: **Flipping for Both**, after the first flip lands tails.

The model predicts that, initially, you believe there will not be more than four tails in a row, yet upon learning that the first flip has landed tails, you *continue* to believe that there will not be four tails in a row. No counterexample to Anticipation is predicted.

This is a bad result. With respect to **Flipping for Heads**, the model makes the attractive prediction that upon seeing the first flip land tails, you revise your strongest belief about how many tails in a row there will be. Indeed, to do otherwise, as I argued in §2.1, would violate the plausible No Gambler's Fallacy principle. But this approach fails to extend this desirable feature to **Flipping for Both**. Here, the model predicts the *counterexamples* to No Gambler's Fallacy: on seeing the first flip land tails, you should now expect the first heads to occur sooner. So, insofar as we find this approach attractive because it seemed to track No Gambler's Fallacy, we now see that their approach fails to do so once applied to **Flipping for Both**. The predictions on this matter are therefore objectionably asymmetric:

possibility open. While my preferred approach in §4 will avoid this awkward consequence, I will not stake much on this objection. After all, we are not understanding these orderings of normality pre-theoretically, so the above awkwardness may just mean that, when modelling **Flipping for Both**, we need to redefine the orderings.

in **Flipping for Heads**, the gambler's-fallacy-like inference is forbidden, yet in **Flipping for Both** it is mandated.

For these reasons, Goodman and Salow's earlier (2018) approach, to which the approach of Lin and Kelly (2012) (2021) behaves relevant similarly,[20] should be rejected. The correct theory of belief revision should fully respect NO GAMBLER'S FALLACY, and should accordingly deny ANTICIPATION.

3.2 Goodman and Salow's Probabilistic, Question-Sensitive Theory

It's worth identifying the more general reason why the above approach cannot predict counterexamples to ANTICIPATION even though it can predict counterexamples to PRESERVATION. The reason is that it implicitly assumes the following weakening of PRESERVATION:

> **NORMAL PRESERVATION**
> If one is justified in believing p, and e is consistent with the most normal worlds consistent with one's evidence, then one would still justified in believing p were they to learn e as total information.

w is among the "most normal" world consistent with one's evidence e just in case, for all $v \in w$, $w \succeq v$. NORMAL PRESERVATION, when combined with NORMAL BELIEF, entails ANTICIPATION. For suppose that one is justified in believing p, but is not upon learning e. NORMAL PRESERVATION tells us that e is therefore incompatible with all of the most normal worlds. It follows that $\neg e$ *cannot also* be incompatible with all of the most normal worlds, and so that one would still be justified in believing p were one to learn $\neg e$, validating ANTICIPATION.

In order to model ANTICIPATION failures, we need an approach that does not implicitly assume NORMAL PRESERVATION. One way to do this is to adopt an approach on which the normality ordering of worlds *changes* on receiving new evidence. If we allow for this, then new worlds may get introduced into the belief set even if the evidence did not rule out all of the most normal worlds, so long as the new ordering imposed by the learnt evidence promotes a previously abnormal world to now be among the most normal.

This is the approach my theory in §4 will take. However, Goodman and Salow (2021) (ms) have also recently outlined a sophisticated probabilistic theory of normality in which normality orderings are dependent on evidence. And, indeed, this allows their approach

[20] In Goodman and Salow's terminology, Lin and Kelly propose that for two distinct worlds w and v, $w \gg v$ iff $\frac{P(w)}{P(v)} > t$ for specific probability function P and threshold $t \geqslant 1$. Assuming that P respects the initial objective chances in **Flipping for Heads** and **Flipping for Both**, that $t = 15$, and defining one's strongest belief in the same way as in the main text, it's easy to verify that we get the exact same results as in figures 1-4.

to predict other counterexamples to ANTICIPATION. However, they still face significant trouble when it comes to the counterexample contained in **Flipping for Both**.

Goodman and Salow's (2021) (ms) more recent approach introduces two new components: a probability function Pr defined over W, and a partition of W, or "salient question", Q. Introducing a probability function to their framework allows Goodman and Salow to propose the following evidence-dependent definitions for our two orderings of normality, $\succeq$ and $\ggcurly$. Where $Pr(p \mid e) = \frac{Pr(p\&e)}{Pr(e)}$ for $P(e) > 0$, and w_1, w_2 and v are elements of W:

PROBABILISTIC NORMALITY

- $w_1 \succeq_e w_2$ iff $Pr(\{w_1\} \mid e) > Pr(\{w_2\} \mid e)$; and
- $w_1 \ggcurly_e w_2$ iff $Pr(\{v : \neg(w_2 \succeq_e v)\} \mid \{v : w_1 \succeq_e v\}) \geqslant t$.

Here's the intuitive idea. For an agent with evidence e, whether world w_1 is more or less normal than w_2 (i.e. the relation $\succeq$) depends just on the probability of that world conditional on e. Worlds that are more probable are more normal. Whether w_1 is *sufficiently* more normal than w_2 (i.e. the relation $\ggcurly$) depends on the probability that the agent's actual situation is more normal than w_2, conditional on the agent's actual situation being at least as weird as w_1. If conditional on things being at least as weird as w_1, it is still sufficiently likely (i.e. at least t likely) that things are more normal than w_2, w_1 is sufficiently more normal than w_2.

Although we now have a notion of normality that is evidence-dependent, these new definitions *alone* do not allow us to predict failures of ANTICIPATION. For instance, it's easy to verify that, setting $t = \frac{15}{16}$, and assuming Pr aligns with the objective chances, the current approach will make the same predictions as in figures 3 and 4 regarding **Flipping for Both**. For example, before the first flip in **Flipping for Both**, the probability that things are more normal than t^5 conditional on things be as weird as t^1 is $\frac{15}{16}$, so $t^1 \ggcurly t^5$; yet, on learning that the first flip has landed on tails, this probability is *still* $\frac{15}{16}$, meaning will still have $t^1 \ggcurly t^5$. Indeed, Goodman and Salow prove that the above definitions validate ANTICIPATION in general (Goodman and Salow ms, p. 48). Failures of Anticiaption only arise once they introduce questions, which we'll now introduce.

Goodman and Salow maintain that whether the agent being modelled is justified in believing some proposition depends on what question is relevant, where a question Q is modelled as a partition of W.[21] Letting $[w]_Q$ designate the member of Q to which w belongs, they redefine the orderings as follows.

[21]See (Holguín forthcoming), (Blumberg and Lederman 2020), (Yalcin 2018) and (Leitgeb 2017) for other question-sensitive approaches to belief. I speak loosely here to remain neutral between semantic versions of question-sensitivity (e.g. (Goodman and Salow 2021) and (Holguín forthcoming)) and subject-sensitive versions (e.g. (Leitgeb 2017)).

Question-Sensitive Normality

- $w_1 \succeq_{e,Q} w_2$ iff $Pr([w_1]_Q \mid e) > Pr([w_2]_Q \mid e)$; and
- $w_1 \ggcurly_{e,Q} w_2$ iff $Pr(\{v : \neg(w_2 \succeq_{e,Q} v)\} \mid \{v : w_1 \succeq_{e,Q} v\}) \geqslant t$.

The intuitive idea being that w_1's normality is determined not by it's own probability, but by the probability that it is a member of the true answer to the salient question.

We now also need to re-define Normal Belief so that it takes into account the now posited question-sensitivity of belief, as follows:

Question-Sensitive Normal Belief
$B_{e,Q} = \{w \in e : \forall v \in e, \neg(v \ggcurly_{e,Q} w)\}$.

In words: an agent with evidence e must leave open a world w, with respect to Q, iff no other worlds consistent with their evidence are more normal than w. Note that I will sometimes drop the 'Q' parameter when the relevant Q is clear from context.

This more sophisticated framework raises three issues that need discussing. The first two—discussed in §3.2.1 and §3.2.2—concern the fact this framework allows Goodman and Salow to predict some counterexamples to Anticipation. I'll demonstrate that, this being so, they still do not offer us the resources to model Anticipation failure in **Flipping for Both**. The third is that introducing question-sensitivity potentially offers Goodman and Salow the resources to *explain away*, rather than predict, my counterexample to Anticipation. I argue against this in §3.2.3.

3.2.1 Cross-cutting Questions

The first kind of Anticipation failure occurs when the relevant agent receives evidence that cross-cuts the salient quesion. Goodman and Salow (ms, p. 28) consider:

> **Celebrity Hike.** 101 celebrities go on a hike in Runyon Canyon. A paparazzo shadowing them notices a hiking pole on the trail. On inspection, he notices some fingerprints. He knows that Michael Jackson and Beyoncé were on the hike, and that Michael always hikes wearing one glove on his right hand and Beyoncé always hikes wearing one glove on her left hand. After inspecting the pole further, the paparazzo discovers whether the fingerprints were made by a left or right hand.

Goodman and Salow consider the following model. Let W consist of 200 worlds, one for each hand that might have made the fingerprints on the pole: one for Michael's left hand, one for Beyoncé's right hand, and one for either hand of the remaining 98 celebrities. Let Pr be a uniform probability distribution over W, and Q the question *who dropped the pole*. Setting $t = 0.99$, it follows that, before looking at the fingerprints, the paparazzo is justified in believing p: *someone other than Michael or Beyoncé dropped the pole*. To see this, note first that the answers *Michael dropped the pole* and *Beyoncé dropped the pole* are half as likely as any other answer, and so count as less normal by Question-sensitive Normality. A world in which some other celebrity dropped the pole is therefore sufficiently more normal than a world in which (say) Michael dropped the pole: the probability that things are more normal than Michael having dropped the pole, conditional on things being at least as weird as someone other than Michael or Beyoncé having dropped the pole, is 0.99. So the paparazzo is initially justified in believing p. However, were the paparazzo to learn that the fingerprints on the pole were made by a left hand, or were he to learn the fingerprints were made by a right hand, the paparazzo must then give up his belief in p, as all the remaining answers—which will include just one of *Michael dropped the pole* or *Beyoncé dropped the pole*—will be equally likely. Hence Anticipation fails.

Still, we do not have the resources to provide a plausible model of **Flipping for Both**. This counterexample depends on having evidence that cross-cuts the answers to the salient question, but this does not happen in **Flipping for Both**—the relevant anti-Anticipation intuitions are solicited just by considering the question *What sequence of heads/tails will the coin produce*. Second, it's difficult to model Anticipation failure in **Flipping for Both** even if we model the case as one in which your evidence cross-cuts the relevant question. As outlined in the following footnote, it's possible to generate models that look *close* to the Anticipation failure argued for in §2.2, but they are not exactly the same, and otherwise require the use of highly gerrymandered partitions of W as the relevant question.[22]

[22]For instance, where W is as described in §2.2 for **Flipping for Both**, let:

$$Q = \{For\ 1 \leqslant n \leqslant 4 : \{\mathfrak{h}^n, \mathfrak{t}^n\}\} \cup \{For\ n \geqslant 5,\ x \in \{\mathfrak{h}, \mathfrak{t}\} : \{x^n, x^{n+1}, ..., x^{1{,}000}\}\}$$

That is, the possible answers are: *1 heads/tails in a row,..., 4 heads/tails in a row, More than 4 heads in a row*; *More than 4 tails in a row*. Setting $t = \frac{15}{16}$, Question-Sensitive Normality entails that one believes, before the first flip, that there will not be more than 4 heads/tails in a row, yet this belief must be given up regardless of whether one learns that the first flip landed heads or that it landed tails. (Crunching the numbers: before the first flip, the probability that things are more normal than $\mathfrak{h}^5$ conditional on things being as weird as $\mathfrak{h}^1$ is equal to $\frac{15}{16}$, so $\mathfrak{h}^1 \ggcurly \mathfrak{h}^5$. Similarly, we also get that $\mathfrak{t}^1 \ggcurly \mathfrak{t}^5$. If the first flip lands heads, $[\mathfrak{h}^4]_Q$ comes to have the same probability as $[\mathfrak{h}^5]_Q$, meaning $\mathfrak{h}^4$ becomes equally as normal as $\mathfrak{h}^5$, and the probability that things are more normal than $\mathfrak{h}^5$ conditional on things being as weird as $\mathfrak{h}^1$ is equal to $\frac{14}{16}$, so $\neg(\mathfrak{h}^1 \ggcurly \mathfrak{h}^5)$. If the first flip lands tails, we similarly also get that $\neg(\mathfrak{t}^1 \ggcurly \mathfrak{t}^5)$.)

Two problems. First, since any world in which there is streak of heads (tails) greater than 4 is placed into the same member of the relevant question, it predicts that on learning that the first flip landed heads (tails), one has the extreme reaction of now having thinking *any* number of heads (tails) is possible. Second, Q is highly gerrymandered; it is not the salient question when soliciting the anti-Anticipation intuitions concerning **Flipping for Both**.

3.2.2 De Se Questions

Goodman and Salow also describe a further kind of counterexample to Anticipation (2021, Appendix C) (ms, p. 36), for example:

> **Flipping for All Heads.** A coin flipper will simultaneously flip 100 fair coins until they all simultaneously land heads. Then he will flip no more.

Plausibly, one in justified in believing both that it will take at least a few simultaneous flips before they land on all heads, but also that the coins will all land on heads simultaneously *eventually*. The strongest proposition one is justified in believing is therefore that the number of trials required before the coin flipping ends falls in some interval $[n, m]$, with $n > 1$.

If that's right, Anticipation should fail. Consider the proposition p: *there will be between 2 and m simultaneous flips*. One is initially justified in believing this as it is entailed by one's strongest justified belief. However, were one to learn that some coins landed tails on the first simultaneous flip, then, abiding by an appropriate analogue of No Gambler's Fallacy, one should give up believing p as one should now take it to be possible that there will be $m + 1$ trials. At the same time, were one to learn that every coin landed heads on the first simultaneous flip, one should now think that the process took exactly 1 trial, therefore also giving up one's belief that p. Anticipation fails.

Modelling this case is tricky, and requires Goodman and Salow to introduce a further tool: *de se* questions. Skipping on the formal details—see (Goodman and Salow 2021, Appendix C) and (Goodman and Salow ms, §10)—here's the basic idea. Goodman and Salow model this kind of Anticipation failure as occuring when the salient question is *When will the coins all land heads together*, with the three answers being *extremely soon*; *an extremely long time from now*; or *in between*. Unlike with regular questions, which worlds belong to which members of this de se partition changes depending on the moment of time the agent is in. For instance, although the world in which the process takes $1,000,000$ attempts may initially count as *an extremely long time from now*, after one has observed $999,999$ attempts all fail to produce 100 simultaneous heads, the world in which is takes $1,000,000$ attempts now counts as *extremely soon*. Selecting an appropriate threshold, one will continue to believe the answer *in between* as their strongest belief for so long as the coin-flipping continues. However, Anticipation can fail since the propositional content of *in between* changes over time: at first *in between* may have propositional content $[n, m]$, but after the first flip has failed to land all heads, it will have content $[n + 1, m + 1]$.

I find this counterexample to Anticipation compelling and the framework of de se questions fruitful. However, it still do not offer us the resources to model Anticipation failure in **Flipping for Both**. There is no natural *de se* question for which Anticipation failure in **Flipping for Both** is predicted. For example, a natural de se question to consider here is: *How long a streak of heads/tails, will the coin produce from now?*. Answers to this question, such as *one more*, will have differing propositional content as the coin flipping in **Flipping for Both** unfolds: before any flip it will have content $\{\mathfrak{h}^1, \mathfrak{t}^1\}$, but after, say, the first flip lands heads, it will have content $\{\mathfrak{h}^1\}$. Setting $t = \frac{15}{16}$, Question-sensitive Normality predicts that before the first flip one believes that that there will be no more than in total four heads in a row, but also that this belief is preserved after learning the first flip has landed heads, meaning Anticipation failure is not predicted.

3.2.3 Appeals to Question-Sensitivity

Let's take stock. Goodman and Salow outline a more sophisticated version of the normality framework that utilizes probability functions and questions. Although this more sophisticated framework can predict *other* counterexamples to Anticipation, it nevertheless fails to provide resources for plausibly modelling the Anticipation failure that occurs in **Flipping for Both**.

Goodman and Salow may hope to utilise their question-sensitive framework in a different way. Perhaps they can use question-sensitivity to *explain away*, rather than predict, my counterexample to Anticipation, in the same way early contextualists about knowledge used context-shifts to explain away apparent failures of single-premise closure (DeRose 1995) (Lewis 1996).

Before outlining the details, I'll start with an objection. Given that Goodman and Salow predict *other* counterexamples to Anticipation, trying to explain away the counterexample from **Flipping for Both** looks ad hoc. This is especially so given that their **Flipping for All Heads** counterexample exploits intuitions similar to those grounding my **Flipping for Both** counterexample and No Gambler's Fallacy principle. It is therefore preferable to have a theory that can predict all of these alleged counterexamples to Anticipation, like the one I present in §4.

Nevertheless, here's how an appeal to question-sensitivity might help. We can distinguish between three questions:

Q_{Tails}: *How many consecutive tails will there be at the beginning of the sequence?*

Q_{Heads}: *How many consecutive heads will there be at the beginning of the sequence?*

Q_{Same}: *How long will the opening consecutive sequence, either of heads or tails, be?*

Diagrammatically, worlds in the same box belong to the same member of the salient question, and are (trivially) all equally as normal. When salient question is Q_{Tails}, following Goodman and Salow's above definitions for $\geq$ we get the following ordering:

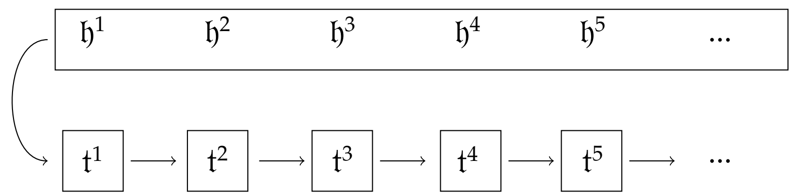

Figure 5: Q_{Tails}: How many tails will there be at the beginning of the sequence?

The diagram for Q_{Heads} is symmetric:

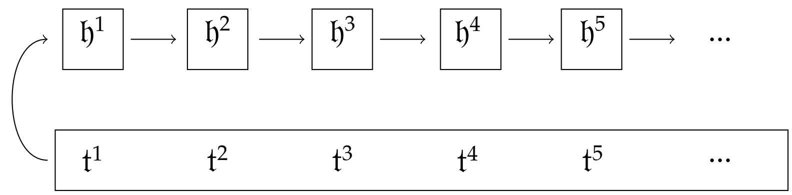

Figure 6: Q_{Heads}: How many heads will there be at the beginning of the sequence?

And for for Q_{Same} we get:

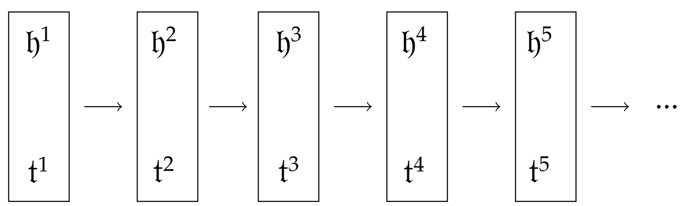

Figure 7: Q_{Same}: How long will the initial repeating sequence, either of heads or tails, be?

Setting $t = \frac{15}{16}$ for our definition of $\gg$, your strongest justified belief will then either be that there will be no more than 3 tails/heads in a row (in Q_{Tails}/Q_{Heads}), or that there no more than 4 of the same in a row (in Q_{Same}).

These predictions allow us to offer a debunking explanation of my argument against Anticipation in §2.2. Crucially, my argument relied on inferring (v) from (ii) and (iv) (see p. 11, with "3" subbed in for "20"):

(ii) You are justified in believing that the coin will not land tails 3 times in a row.

(iv) You are justified in believing that the coin will not land heads 3 times in a row.

(v) You are justified in believing that the coin will not land the same way 3 times in a row.

While this looks like an innocent application of our assumption that beliefs are deductively closed, there is no fact no *single* question where (v) can be legitimately derived. (ii) is true with respect to Q_{Tails}, but (iv) is not: with respect to Q_{Tails}, you leave any number of heads open. Likewise, (iv) is true with respect to Q_{Heads} but (ii) is not. As for Q_{Same}, *neither* (ii) nor (iv) hold: you can neither rule 4 tails nor 4 heads out. My argument therefore appears insensitive to subtle changes of the relevant question, illicitly deriving (v) from premises that are only true with respect to different questions.

I take this to be the most promising reply Goodman and Salow can make to my arguments in §3.1. Nevertheless, even ignoring the extent (mentioned above) to which it feels ad hoc, it is unsatisfactory. The problem is that the above three questions are not the only ones that might be relevant. We might instead be asking the question Q_{Seq}: *What sequence will the coin produce*, which is naturally interpreted as a trivial partition on W, consisting of a singleton-cell for each member of W. Yet, with respect to Q_{Seq}, Goodman and Salow's more sophisticated approach does not make importantly different predictions to those theories considered in §3.1. It therefore faces the same objection: we get counterexamples to No Gambler's Fallacy. In particular, setting the threshold t to $\frac{15}{16}$, the probability that things are more normal than t^5 conditional on things be as weird as t^1 is $\frac{15}{16}$, so $t^1 \ggg t^5$; yet, on learning that the first flip has landed on tails, this probability is *still* $\frac{15}{16}$, meaning will still have $t^1 \ggg t^5$. In words: while one is initially justified in thinking there will be no more than 4 *more* tails, one fails to be justified in believing this upon seeing the first flip land tails. To the extent that it is these kinds of worries that are motivating our initial departure from Preservation and Anticipation, Goodman and Salow's approach should still strike us as unsatisfactory, even given their more flexible question-sensitive framework.

4 A New Approach: Objective Normality

We need an alternative approach; one that properly vindicates the counterexample to Anticipation contained in **Flipping for Both**. Outlining one is the goal of this section. I'll first introduce a simple model that can make the right predictions for **Flipping for Both** in §4.1. Not only does this model generate the correct predictions, it is also naturally motivated by the popular idea that the function of belief is to simplify reasoning by ruling out sufficiently unlikely possibilities (see e.g. (Harsanyi 1985), (Lance 1995), (Lin 2013) and (Ross and Schroeder 2014)). However, this simple model faces two problems it must answer before it can be considered a genuine contender for a theory of belief revision. §4.2 develops a modified theory in light of those objections.

4.1 A Simple Model

To develop a model that can predict Anticipation fails in **Flipping for Both**, it's instructive to isolate the feature of Goodman and Salow's approach that prevents them from doing this. The key issue is that, on their approach, whether a world is normal enough that it cannot be ruled out depends on how its probability *compares* to the probability of other worlds. It's worth looking at their proposed model of **Flipping for Both** again. Initially, setting $t = \frac{15}{16}$, t^5 can be ruled out on the following grounds: conditional on things being at least as weird as t^1, it's sufficiently likely that things are more normal than t^5 (a probability of $\frac{15}{16}$). Now, on learning that the first flip lands tails, the probability of t^5 *increases* from $\frac{1}{32}$ to $\frac{1}{16}$; however, the probabilities of t^1-t^4 increase at the same rate. This means that it's *still* $\frac{15}{16}$ likely that things are more normal than t^5 conditional on things being at least as weird as t^1, and hence t^5 is still sufficiently abnormal that it can be ruled out, despite its increase in probability.

We may therefore make better predicitons if we reject the implicit idea in the above model that normality is a *relational* matter; that is, the idea that how normal a world is depends on how its probability compares to the probability of other worlds. Instead, we can judge the normality of a world by looking at how its probability compares to some independent, fixed value. If, for instance, we say that a world is sufficiently normal whenever it is at least $\frac{1}{16}$ likely, we should predict that t^5 goes from being abnormal before the first coin is flipped, with probability $\frac{1}{32}$, to sufficiently normal after the first flip lands tails, with probability $\frac{1}{16}$. More generally, let N_e be the set of "normal" worlds given our agents evidence e, and Pr the relevant probability function. We'll say that a world compatible with one's evidence counts as normal just in case it's probability reaches a certain threshold t: $N_e = \{w \in e \mid Pr(\{w\} \mid e) \geqslant t\}$. The idea is to then identify one's strongest justified belief, B_e, with N_e. This idea fits nicely with the idea that the role of belief is to allow agents to ignore possibilities that are sufficiently likely: the worlds inside the belief set are just those that too likely to be ignored in this way.

Setting $t = \frac{1}{16}$, we can construct an attractive model of **Flipping for Both**. It's illuminating to again depict our model diagrammatically. The conventions are the same as in the diagrams above, except that we no longer have a defined relation $\geq$ to depict, and next to each world is the probability assigned to it by Pr conditional on the relevant evidence e. We get the following result before the coin is flipped:

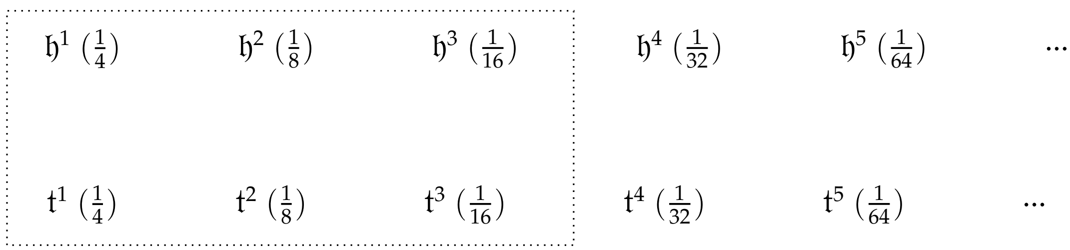

Figure 8: Alternative **Flipping for Both** model, before any flips.

That is, you can justifiably believe that there will be no more than 3 tails, no more than 3 heads, and no more than three of the same. Upon learning that the first flip lands on tails, we have:

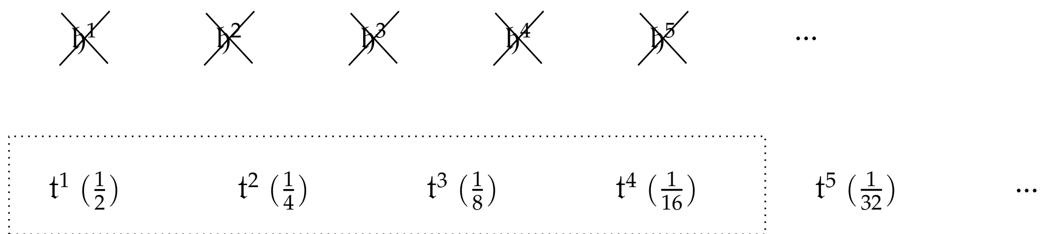

Figure 9: Alternative **Flipping for Both** model, after first flip.

That is, upon seeing the first flip land heads, you must now give up your belief that there will be no more than three heads, and thus also your belief that there will be no more than three of the same. A symmetric result holds in the case in which the first flip lands on heads, giving us the Anticipation failure argued for in §2.2: initially you believe that there will be no more than three of the same, and this belief will be revised both on learning that the first flip lands on tails and on learning that the first flip lands on heads. In turn, unlike with Goodman and Salow's approach, we predict no counterexamples to No Gambler's Fallacy.[23]

4.2 Objective Normality

Despite these attractive results, the simple model faces two issues that prevent it from forming the basis of a plausible theory of belief revision. For the first issue, consider:

> **Dime or Nickel.** I am about to perform the coin-flipping procedure described in **Flipping for Both**. However, I am unsure whether to use a dime or a nickel

[23]It's also easy to check that this approach can predict the other counterexamples to Anticipation outlines in §3.2.1 and §3.2.2, so long as, for the latter, we follow Goodman and Salow and introduce *de se* questions.

for my coin. I decide to roll a 6-sided die to decide: I'll use a dime if it lands even and a nickel otherwise.

Plausibly, the matter of whether I use a dime or a nickel is irrelevant to how long a streak of tails I might get. Both are fair coins. So there should be no difference in your beliefs in **Flipping for Both** and in **Dime or Nickel**. The simple model will not predict this if, for **Dime and Nickel**, we have to distinguish between worlds in which the same sequence is produced but by a different coin. For example, there will now be two worlds in which the coin lands tails 4 times, and each will have probability $\frac{1}{32}$. Assuming the threshold remains at $\frac{1}{16}$, this will mean you are permitted to have stronger beliefs in **Dime or Nickel** than you are in **Flipping for Both**: in the former, but not the latter, you can rule out the coin landing tails 4 times in a row.

The most promising response to problems of this kind is to follow Goodman and Salow—as well as Leitgeb (2017), who faces a similar problem (Staffel 2016, pp. 1731-2)—by endorsing a question-sensitive account of justified belief. The idea will be that, in **Dime or Nickel**, so long as the relevant question is, say, *What sequence will be produced*, then worlds containing the same sequence, even if produced by a different coin, will belong to the same answer to the relevant question. So, at least with respect to the relevant question, we'll get the prediction that your beliefs in **Dime or Nickel** should be the same as in **Flipping for Both**. Formally, we redefine our set of normal worlds, as follows, so that it is sensitive to the relevant question: $N_{e,Q} = \{w \in e \mid Pr([w]_Q \mid e) \geqslant t\}$. In other words, the normal worlds for an agent with evidence with e, realtive to Q, are those that belong to the sufficiently likely answers to Q.

The second issue is that, so far, nothing stops agents having justified beliefs in highly unlikely propositions. If the threshold t is $\frac{1}{3}$, and the relevant question Q consists of four answers, a, b, c and d, such that $Pr(a \mid e) = \frac{1}{3}$, yet $Pr(b \mid e) = Pr(c \mid e) = Pr(d \mid e) = \frac{2}{9}$. Here, $N_{e,Q} = a$ and so the relevant agent is justified in believing a proposition they take to only be $\frac{1}{3}$ likely. A particularly acute version of this problem occurs if *none* of the answers to Q have probability of at least t. In that scenario, $N_{e,Q}$ is the empty set, and so it is predicted that our agent is justified in believing the contradiction.[24]

The most obvious fix to this problem is to introduce a necessary threshold that any proposition must meet in order to be believed. However, we cannot simply say that one ought to have strongest belief (with respect to Q) equal to $N_{E,Q}$ only if $Pr(N_{E,Q}) \geqslant T$ for some $0 \leqslant T \leqslant 1$, as this fails to specify a strongest justified belief for cases in which $P(N_{E,Q}) < T$. That is, adding a threshold for when one can believe one is in a normal

[24]While theories of weak belief, such as those in (Dorst and Mandelkern forthcoming) and (Holguín forthcoming), will be happy with justified beliefs in unlikely propositions, they will not be happy with justified beliefs in contradictory propositions.

world ($N_{E,Q}$) gives no guidance regarding what to believe when it is not sufficiently likely one is in a normal world.

A natural idea here—one that follows the spirit of (Goodman and Salow 2023), but not the letter—is that when one is not in a position to justifiably believe that one is in the normal worlds, one should retreat, yet still take oneself to be in a world that is at least *pretty* normal world. If it's still sufficiently unlikely that one is among worlds that are at least pretty normal, one should retreat again and take oneself to be in a world that is at least *somewhat* normal, so on and so forth.

Formally, we introduce multiple thresholds t_i. t_1 will constitute the standard for being among the most normal worlds; t_2 the standard for being among at least the second most normal worlds, and so on. We can then define different ranks of normality, $N^i_{e,Q}$, as follows: $N^i_{e,Q} = \{w \in e \mid Pr([w]_Q \mid e) \geqslant t_i\}$. In words, the i^{th} rank of normality, with respect to Q and for an agent with evidence e, consists of all the worlds that belong to an answer to Q that is at least t_i likely conditional on the agents evidence.

We arrive at the following theory:

Objective Normal Belief

Let $t_1, t_2 ...$ and T be elements of the unit interval $[0, 1]$ such that $t_1 > t_2 > ...$ Where $N^i_{e,Q} = \{w \in e \mid Pr([w]_Q \mid e) \geqslant t_i\}$, the strongest justified belief, relative to Q, of an agent with evidence e, is:

$$B_{e,Q} = \begin{cases} N^1_{e,Q}, \text{ if } P(N^1_{e,Q} \mid e) \geqslant T \\ ... \\ N^i_{e,Q}, \text{ if } Pr(N^i_{e,Q} \mid e) \geqslant T \text{ and } Pr(N^{i-1}_{e,Q} \mid e) < T \\ ... \end{cases}$$

That is, one's strongest justified belief (with respect to e and Q) is that one is in a world among the first rank of normality that is sufficiently likely. Restating in procedural terms: if it's likely enough you are in one one of the most normal worlds, believe this is so. If not, check whether it is likely enough that you are among either the most or second most normal worlds: if it is, believe you are among those worlds; if not check whether it is likely enough that you are among either the most, second most, or third most normal worlds. So on and so forth.

Setting T high enough, Objective Normal Belief avoids the threshold problem: you'll never justifiably believe any proposition that is less than T likely. Yet it still preserves the same predictions regarding the coin flipping cases that made the simple model in §4.1 attractive. If we set T to $\frac{1}{2}$ and t_1 to $\frac{1}{16}$, then with respect to the question *What sequence* *will be produced*, we will get the same results as depicted in Figures 8 and 9 above, thereby having Anticipation failures, and thereby satisfying No Gambler's Fallacy.

Objective Normal Belief thus offers a natural picture of belief that makes the best predictions yet in these coin-flipping cases. It has various other advantages, too. For one, Stalnaker (1994) outlines a famous counterexample to Preservation—concerning beliefs about the nationality of three composers—which turn out also to constitute counterexamples to Anticipation.[25] Like with **Flipping for Both**, Goodman and Salow's sophisticated theory cannot predict this counterexample to Anticipation. Objective Normal Belief can, as I argue in (Pearson ms).[26] For another, Objective Normal Belief can be naturally translated into a normal-theoretic account of knowledge—such as those in (Stalnaker 2006) and (Goodman and Salow 2023)—on which one knows p iff p is true in all worlds at least as normal as the actual world. For we can say that, if the highest rank of normality one's actual world is in is N_e^i, then N_e^i is the strongest proposition one knows.[27]

The trade-off is that now our theory of belief revision is extremely weak. For the most part, it is as weak as the theory offered in (Goodman and Salow ms); it predicts similar counterexamples not only to weak principles of belief revision like Anticipation, but also to weak principles concerning how new beliefs can be formed.[28] One difference, however, concerns:

Cautious Monotinicty
If one is justified in believing p and in believing e, one is still justified in believing p on learning e.

Goodman and Salow's approach predicts counterexamples to Cautious Monotinicity,

[25]I take this observation from (Goodman and Salow ms), who credit the observation to Ginger Schultheis via personal communication.

[26]A brief explanation of Stalnaker's example and the proposed model. Our agent believes that Satie is French and that Verdi and Bizet are compatriots, but is 50/50 as to whether Verdi and Bizet are both Italian or both French. Our agent then learns that all three composers are compatriots. The judgement that causes trouble for various theories of belief revision is the judgement that our agent should now be ambivalent as to the three composers are all French or all Italian. As (Schultheis 2018) demonstrates, part of the problem is that, given natural modelling assumptions, learning that all three composers are compatriots does not reduce the probability that Satie is French, making it hard to see why this belief ought to be defeated. Objective Normal Belief *is* able to predict defeat in this case since, although the probability in the proposition Satie is French has not decreased, the probability in the *possibility* that all three composers are Italian may nevertheless increase to the extent that it can no longer be ignored.

[27]Alternatively, we can allow for KK-failure by saying that, for some margin-of-error m, when the highest rank of normality one's actual world is in is N_e^i, the strongest proposition one knows is N_e^{i+m}.

[28]For instance, we can generate similar counterexamples to the principle sometimes called "Frontloading" or "Inductive Anti-Dogmatism": that one can justifiably believe p on learning e only if one had a prior justified belief that $e \supset p$. See (Goodman and Salow 2023, §8) and the references therein for detailed discussion.

but only once we introduced *de se* questions into the framework (see §3.2.2)[29] OBJECTIVE NORMAL BELIEF predicts counterexamples to CAUTIOUS MONOTINICITY even without allowing for *de se* questions. Consider:

> **Premier League.** It's almost the end of the English Premier League season, and only three teams remain that have a chance of winning it. Manchester City are 60% likely to win, whereas Arsenal and Liverpool are each only 20% likely to win. Initially, you believe City will win. However, you then learn that Arsenal have lost a crucial game; only City and Liverpool remain in contention.

Here's an (overly) simple model of **Premier League** using OBJECTIVE NORMAL BELIEF. Let $Q = \{City\ win;\ Arsenal\ win;\ Liverpool\ win\}$. Let Pr be such that $Pr(City\ win) = \frac{6}{10}$, $Pr(Arsenal\ win) = \frac{2}{10}$ and $Pr(Liverpool\ win) = \frac{2}{10}$, and let $t_1 = \frac{1}{4}$ and $T = \frac{1}{2}$. Then initially, you are justified in believing that City will win. However, on learning that Arsenal will lose—call this proposition e—we have $Pr(Liverpool\ win \mid e) = \frac{1}{4}$, meaning you can no longer rule out Liverpool winning.

Despite the weak theory of belief revision that it ensues, I hope to have demonstrated that OBJECTIVE NORMAL BELIEF otherwise constitutes a plausible theory of belief, possessing distinctive advantages over its rivals.

5 Conclusion: Life Without Anticipation

I have argued that ANTICIPATION fails in cases like **Flipping for Both**, and that while the theories of belief revision offered so far fail to get this result, my OBJECTIVE NORMAL BELIEF constitutes a respectable theory that does. Still, one may remain uncomfortable with these results. Recall the considerable intuitive appeal of ANTICIPATION argued for in §1. If ANTICIPATION is false, I argued, bizarre and infelicitous assertions are licensed, and one can be rational in avoiding free evidence. Do we have to bite the bullet here, or can we avoid these awkward consequences?

My arguments concerning infelicitous assertions relied on contentious principle connecting belief revision and beliefs in conditionals. Triviality results, from e.g. (Lewis 1976) and (Gärdenfors 1986), give us reason to doubt there is any such neat connection here. But thinking there is *some* connection between belief revision and conditionals is irresistible. So, those who hope to ameliorate this awkward consequence of ANTICIPATION failures

[29]They discuss this principle under the same '' (Goodman and Salow ms). The counterexample arises from the **Flipping for All Heads Case**, as described in §3.2.2: one believes that it will take more than one simultaneous flip for all the coins to land heads, yet on learning this, there is some n for which one drops the belief that it will take less than n simultaneous flips for all the coins to land heads.

by denying there is a straightforward connection between belief revision and beliefs in conditionals will need to tell us what the not-so-straightforward connection is.

My arguments concerning rational evidence avoidance depended on justified beliefs playing a substantive role is rational decision making; in particular, as having their content usable as premises in practical reasoning. Perhaps justified beliefs play no such important role, meaning we can avoid this awkward consequence of Anticipation failure. However, so long as justified beliefs play *some* important role in rational decision making—one that cannot be reduced to the role played by rational credences—those who hope to pursue this strategy need to tell us what this role is. Alternatively, perhaps all that matters in rational action *is* one's rational credences (Jeffrey 1970). I'm inclined to resist this conclusion, but I cannot deny that my arguments forge a new path to reaching it.

References

Alchourrón, Carlos E., Peter Gärdenfors, and David Makinson (1985). "On the Logic of Theory Change: Partial Meet Contraction and Revision Functions". In: *Journal of Symbolic Logic* 50.2, pp. 510–530. doi: 10.2307/2274239.

Blumberg, Kyle and Harvey Lederman (2020). "Revisionist Reporting". In: *Philosophical Studies* 178.3, pp. 755–783. doi: 10.1007/s11098-020-01457-4.

Buchak, Lara (2010). "Instrumental Rationality, Epistemic Rationality, and Evidence-Gathering". In: *Philosophical Perspectives* 24.1, pp. 85–120. doi: 10.1111/j.1520-8583.2010.00186.x.

Carter, Sam and John Hawthorne (forthcoming). "Dogmatism & Inquiry". In: *Mind*.

Comesaña, Juan (2020). *Being Rational and Being Right*. Oxford, United Kingdom: Oxford University Press.

DeRose, Keith (1995). "Solving the Skeptical Problem". In: *Philosophical Review* 104.1, pp. 1–52. doi: 10.2307/2186011.

Dorr, Cian, Jeremy Goodman, and John Hawthorne (2014). "Knowing Against the Odds". In: *Philosophical Studies* 170.2, pp. 277–287. doi: 10.1007/s11098-013-0212-9.

Dorst, Kevin and Matthew Mandelkern (forthcoming). "Good Guesses". In: *Philosophy and Phenomenological Research*. doi: 10.1111/phpr.12831.

Fantl, Jeremy and Matthew McGrath (2009). *Knowledge in an Uncertain World*. Ed. by Matthew McGrath. New York: Oxford University Press.

Fine, Kit (1975). "Vagueness, Truth and Logic". In: *Synthese* 30.3-4, pp. 265–300. doi: 10.1007/bf00485047.

Freund, Micheal and Daniel Lehmann (1996). "On Negation Rationality". In: *Journal of Logic and Computation* 6.2, pp. 263–269. doi: 10.1093/logcom/6.2.263.

Gärdenfors, Peter (1986). "Belief Revisions and the Ramsey Test for Conditionals". In: *Philosophical Review* 95.1, pp. 81–93. doi: 10.2307/2185133.

Goodman, Jeremy and Bernhard Salow (2018). "Taking a Chance on KK". In: *Philosophical Studies* 175.1, pp. 183–196. doi: 10.1007/s11098-017-0861-1.

— (June 2021). "Knowledge from Probability". In: *Electronic Proceedings in Theoretical Computer Science* 335, pp. 171–186. issn: 2075-2180. doi: 10.4204/eptcs.335.15. url: http://dx.doi.org/10.4204/EPTCS.335.15.

— (2023). "Epistemology Normalized". In: *Philosophical Review*.

— (ms). "Belief Revision Normalized".

Harsanyi, John C. (1985). "Acceptance of Empirical Statements: A Bayesian Theory Without Cognitive Utilities". In: *Theory and Decision* 18.1, pp. 1–30. doi: 10.1007/bf00134075.

Hawthorne, John, Daniel Rothschild, and Levi Spectre (2016). "Belief is Weak". In: *Philosophical Studies* 173.5, pp. 1393–1404. doi: 10.1007/s11098-015-0553-7.

Hawthorne, John and Jason Stanley (2008). "Knowledge and Action". In: *Journal of Philosophy* 105.10, pp. 571–590.

Holguín, Ben (forthcoming). "Thinking, Guessing, and Believing". In: *Philosophers' Imprint*, pp. 1–34.

Huber, Franz (2013). "Belief Revision I: The Agm Theory". In: *Philosophy Compass* 8.7, pp. 604–612. doi: 10.1111/phc3.12048.

Jeffrey, Richard (1970). "Dracula Meets Wolfman: Acceptance Vs. Partial Belief". In: *Induction, Acceptance, and Rational Belief*. Ed. by Marshall Swain. D. Reidel, pp. 157–185.

Kelly, Kevin and Hanti Lin (2021). "Beliefs, Probabilities, and Their Coherent Correspondence". In: *Lotteries, Knowledge, and Rational Belief: Essays on the Lottery Paradox*. Ed. by IgorEditor Douven. Cambridge University Press, pp. 185–222. doi: 10.1017/9781108379755.011.

Kraus, Sarit, Daniel Lehmann, and Menachem Magidor (1990). "Nonmonotonic Reasoning, Preferential Models and Cumulative Logics". In: *Artificial Intelligence* 44.1-2, pp. 167–207. doi: 10.1016/0004-3702(90)90101-5.

Kripke, Saul A. (2011). "On Two Paradoxes of Knowledge". In: *Philosophical Troubles. Collected Papers Vol I*. Ed. by Saul A. Kripke. Oxford University Press.

Lackey, Jennifer (2008). *Learning From Words: Testimony as a Source of Knowledge*. Oxford: Oxford University Press.

Lance, Mark Norris (1995). "Subjective Probability and Acceptance". In: *Philosophical Studies* 77.1, pp. 147–179. doi: 10.1007/bf00996316.

Leitgeb, Hannes (2014). "The Stability Theory of Belief". In: *Philosophical Review* 123.2, pp. 131–171. doi: 10.1215/00318108-2400575.

Leitgeb, Hannes (2017). *The Stability of Belief: How Rational Belief Coheres with Probability*. Oxford University Press.

Lewis, David (1973). *Counterfactuals*. Blackwell.

— (1976). "Probabilities of Conditionals and Conditional Probabilities". In: *Philosophical Review* 85.3, pp. 297–315. DOI: 10.2307/2184279.

— (1980). "A Subjectivist's Guide to Objective Chance". In: *Studies in Inductive Logic and Probability, Volume II*. Ed. by Richard C. Jeffrey. Berkeley: University of California Press, pp. 263–293.

— (1996). "Elusive knowledge". In: *Australasian journal of Philosophy* 74.4, pp. 549–567.

Lin, Hanti (2013). "Foundations of Everyday Practical Reasoning". In: *Journal of Philosophical Logic* 42.6, pp. 831–862. DOI: 10.1007/s10992-013-9296-0.

— (2019). "Belief Revision Theory". In: *The Open Handbook of Formal Epistemology*. Ed. by Richard Pettigrew and Jonathan Weisberg. PhilPapers Foundation, pp. 349–396.

Lin, Hanti and Kevin Kelly (2012). "Propositional Reasoning That Tracks Probabilistic Reasoning". In: *Journal of Philosophical Logic* 41.6, pp. 957–981. DOI: 10.1007/s10992-012-9237-3.

Lindström, Sten (1996). "The Ramsey Test and the Indexicality of Conditionals: A Proposed Resolution of Gärdenfors' Paradox". In: *Logic, Action and Information*. Ed. by André Fuhrmann and Hans Rott. de Gruyter.

Mandelkern, Matthew and Justin Khoo (2019). "Against Preservation". In: *Analysis* 79.3, pp. 424–436. DOI: 10.1093/analys/any051.

Pearson, J.E. (ms). "Composing Composers".

Ramsey, Frank (1926). "Truth and Probability". In: *Philosophy of Probability: Contemporary Readings*. Ed. by Antony Eagle. Routledge, pp. 52–94.

Ross, Jacob and Mark Schroeder (2014). "Belief, Credence, and Pragmatic Encroachment". In: *Philosophy and Phenomenological Research* 88.2, pp. 259–288. DOI: 10.1111/j.1933-1592.2011.00552.x.

Salow, Bernhard and Arif Ahmed (2019). "Don't Look Now". In: *British Journal for the Philosophy of Science* 70.2, pp. 327–350. DOI: 10.1093/bjps/axx047.

Schultheis, Ginger (2018). "Belief and Evidence". PhD thesis. MIT.

Silva, Paul and Luis R. G. Oliveira (forthcoming). "Propositional Justification and Doxastic Justification". In: *Routledge Handbook of the Philosophy Evidence*. Ed. by Maria Lasonen-Aarnio and Clayton M. Littlejohn. Routledge.

Staffel, Julia (2016). "Beliefs, Buses and Lotteries: Why Rational Belief Can't Be Stably High Credence". In: *Philosophical Studies* 173.7, pp. 1721–1734. DOI: 10.1007/s11098-015-0574-2.

Stalnaker, Robert (1994). “What is a nonmonotonic consequence relation?” In: *Fundamenta Informaticae* 21, pp. 7–22.

— (2006). “On Logics of Knowledge and Belief”. In: *Philosophical Studies* 128.1, pp. 169–199. DOI: `10.1007/s11098-005-4062-y`.

— (2019). *Knowledge and Conditionals: Essays on the Structure of Inquiry*. Oxford University Press.

Williamson, Timothy (2000). *Knowledge and its Limits*. Oxford University Press.

Yalcin, Seth (2018). “Belief as Question-Sensitive”. In: *Philosophy and Phenomenological Research* 97.1, pp. 23–47. DOI: `10.1111/phpr.12330`.