Bayesians Commit the Gambler's Fallacy

Kevin Dorst
Massachusetts Institute of Technology

January 2024
11,294 words

Abstract

The *gambler's fallacy* is the tendency to expect random processes to switch more often than they actually do—for example, to think that after a string of tails, a heads is more likely. It's often taken to be evidence for irrationality. It isn't. Rather, it's to be expected from a group of Bayesians who begin with causal uncertainty, and then observe unbiased data from an (in fact) statistically independent process. Although they converge toward the truth, they do so in an *asymmetric* way—ruling out "streaky" hypotheses more quickly than "switchy" ones. As a result, the majority (and the average) exhibit the gambler's fallacy. If they have limited memory, this tendency persists even with arbitrarily-large amounts of data. Indeed, such Bayesians exhibit a variety of the empirical trends found in studies of the gambler's fallacy: they expect switches after short streaks but continuations after long ones; these nonlinear expectations vary with their familiarity with the causal system; their predictions depend on the sequence they've just seen; they produce sequences that are too switchy; and they exhibit greater rates of gambler's reasoning when forming binary predictions than when forming probability estimates. In short: what's been thought to be evidence for irrationality may instead be rational responses to limited data and memory.

Keywords: Gambler's fallacy, hot-hands fallacy, Bayesian updating, resource-rational analysis, memory limits, hierarchical Bayesian models

Mathematica Notebook: https://github.com/kevindorst/gamblers_fallacy_notebook

1 Introduction

Baylee is bored. The fluorescent lights hum. The spreadsheets blur. She needs air.

As she steps outside, she sees the Prius nestled happily in the front spot. Three days in a row now. The Jeep will probably get it tomorrow—she'll place her bet when she gets back.

Baylee's not alone. Each afternoon in Bay City, her fellow office-mates check who won the morning's battle for the front spot. It's always either a Prius or a Jeep, and each manages to get it about half the time. But they all wonder about the fluctuations: if the Prius won the last few days, what are the chances it'll win tomorrow? They place their bets with Bob.

In fact, the outcomes are statistically independent: each day, the Prius and the Jeep each have a 50% chance to get the spot, regardless of how the previous days have gone. Yet most of our workers seem to think otherwise: when the Prius is on a winning streak, they start to think that the Jeep is "due". (Vice versa when the Jeep is on a streak.) Some disagree: they think that if the Prius is winning, it's more likely to *continue* its "hot streak". But they're in the minority. Indeed, when Bob (the bookie) *averages* their opinions, he finds that as the Prius wins a few in a row, the workers increasingly think that the Jeep is due.

At this point, you might hazard a couple hypotheses.

First: *Baylee and company are very, very bored.* About that, you are right.

Second: *Baylee and company are bad at statistical reasoning.* Since the outcomes are statistically independent, rationally responding to the evidence would've taught them as much. Thus they must instead be using flawed heuristics to form their opinions. The majority believe in the "law of small numbers": they think that since the Prius wins 50% of the time in the *long* run, there must be a mechanism that balances the outcomes out (Tversky and Kahneman 1971). As a result, they commit the **gambler's fallacy:** the tendency to think that streaks of (in fact independent) outcomes are likely to switch (Kahneman and Tversky 1972). Meanwhile, the dissident minority commit the **hot-hands fallacy:** the tendency to think that streaks of (in fact independent) outcomes are likely to continue (Gilovich et al. 1985; Tversky and Gilovich 1989). Yet more examples of human irrationality, you think.

About that, you are wrong. Baylee and company are rational Bayesians. As I'll show, when either data or memory are limited, Bayesians who start out with causal uncertainty about an (in fact independent) process and then learn from unbiased evidence will typically commit either the gambler's or hot-hands fallacies. The majority—and, more surprisingly, the *average*—will commit the gambler's fallacy. Why? Although they'll get evidence that the process is neither "switchy" nor "streaky", they'll get *more* evidence against the latter—and this asymmetric convergence toward the truth will lead to the gambler's fallacy.

More is true. Baylee's predicament is just like ours. We're constantly (and reasonably) unsure whether everyday stochastic processes are independent: weather (rainy or sunny), recipes (tasty or bland), submissions (accepted or rejected), experiments (success or failure), moods (good or bad), conversations (productive or difficult), lectures (engaging or dull)—and even *coins* (heads or tails—see below).

Faced with such processes, people robustly exhibit a variety of trends:

1) **Nonlinear Expectations:** As a streak grows in length, people at first increasingly expect *switches*, but then (once it's long enough) increasingly expect *continuations*.[1]
2) **Experience-Dependence:** These nonlinear expectations take different shapes for different causal systems—people begin expecting continuations earlier for unfamiliar ones (e.g. stock markets) than for familiar ones (e.g. coins). Those with more experience come closer to treating outcomes independently.[2]
3) **Sequence-Dependence:** People's estimates depend on the type of sequence they've just seen—if it switched a lot, they expect a switch; if it switched rarely, they expect a continuation (Bloomfield and Hales 2002; Bao et al. 2023).

[1] Derks 1963; Altmann and Burns 2005; Asparouhova et al. 2009; Suetens et al. 2016; Gershman 2021; Rao and Hastie 2023.

[2] Burns and Corpus 2004; Asparouhova et al. 2009; Chen et al. 2016; Rao and Hastie 2023.

4) **Too Switchy:** When asked to produce random sequences, people produce sequences that switch more than 50% of the time; the switch-rate increases with streak-length.[3]
5) **Format-Dependence:** People exhibit greater rates of gambler's reasoning when forming *binary predictions* than when forming *probability estimates* (Rao and Hastie 2023).

That's how real people reason about random processes.

It's also how *Bayesians* reason about random processes. I'll show that all of these trends are to be expected from Bayesians who start out with causal uncertainty and have either limited data or limited memory. Thus I'll defend the following:

> **Causal-Uncertainty Hypothesis:** The gambler's and hot-hands fallacies are due to causal uncertainty combined with rational responses to limited data and memory.

Two clarifications. First: to avoid begging the question I'll avoid the term "fallacy"—instead refering to "predicting switches" as "gambler's *reasoning*" and to "predicting continuations" as "hot-hands *reasoning*".

Second, and relatedly: it's often suggested that the term "gambler's fallacy" applies only if the person *knows* the outcomes are independent and yet still expects streaks to switch. But that doesn't seem right. To know that the outcomes are independent *just is* to know that a switch is equally likely after a streak—so anyone who exhibits gambler's reasoning does *not* know as much, or at least doesn't activate that knowledge when forming their expectations. The empirical findings are simply trends like (1)–(5) above; if we restrict the term "gambler's fallacy" to cases of known independence, we have little evidence that those trends are instances of it (cf. Kelly 2004). Thus it's better to use the term neutrally, as I do, to capture the empirical evidence that needs to be explained.

That said, I think this suggestion is right about something: often people who exhibit gambler's reasoning are *in a position to know* that the outcomes are independent. First, they've often seen lots of data—if they had a perfect memory, they'd figure it out. This is why my results about limited-memory matter. Second, you might think people should be able to *reason* their way to knowledge of statistical independence. Coins have no memory, after all. And yet people exhibit gambler's reasoning for coins, slot machines, and roulette wheels. Surely *this* isn't rational. Right?

Not obviously. Slot machines *can* change their payout-probabilities (Richtel 2006). Roulette wheels *have* exhibited subtle biases (Fletcher 2018). And coins *do* have a memory—of a sort. When tossed and caught in the hand, coins exhibit *dynamical bias*: due to the fact that they precess (rotate like a pizza) as well as flip, they are on average 51%-likely to land the way that was initially face-up (Diaconis et al. 2007; Bartoš et al. 2023). Thus the probability of what happens next really *does* depend on what happened last (along with your flipping procedure: do you turn it over, or not, when you go to flip it again?). Given surprising facts like this—combined with the observation that most stochastic processes encountered in everyday life are *not* independent (Pinker 1997)—very few people should be *certain*, from the armchair, that a given random process is independent. The best way to figure it out is to learn from experience. That's where my explanation kicks in.

[3] Bar-Hillel and Wagenaar 1991; Rapoport and Budescu 1992, 1997; Budescu and Rapoport 1994; Chen et al. 2016.

Suppose I'm right. Why does it matter?

First, there is an ongoing debate about whether to understand the mind in *rational* or *mechanistic* terms: using strategies that make sense for solving the problems we face (Marr 1982; Anderson 1990), or instead using kludge-y heuristics that did well enough in our evolutionary past (Tversky and Kahneman 1974). The mechanistic approach was dominant in the late 20th century, spurred the development of behavioral economics, and seeped into the wider culture—often via claims that people are systematically and predictably irrational.[4] However, the rational approach has gained momentum in recent decades: resource-constrained optimality has been the dominant approach in explaining low-level cognitive processes like vision (Geisler 2011), motor control (Todorov 2004), and memory (Anderson 1990); and it's recently had successes in explaining high-level cognition as well.[5]. This trend in cognitive science has been matched with a trend in philosophy to criticize the normative presuppositions of many apparent empirical biases.[6]

Despite this momentum, there's a fly in the ointment: if people really do use sophisticated Bayesian reasoning to perform everyday feats of human cognition, why do they seem so *bad* at probabilistic reasoning in some of the easiest cases? The gambler's fallacy is one of the most pressing.[7] Although Bayesian explanations exist, as we'll see they all assume some sort of hard-wired error—for example, mistakenly (and dogmatically) ruling out the possibility that the outcomes are statistically independent—and then go on to explain how rational principles might explain people's behavior *given* that error. The initial error goes unexplained, lending support to the mechanistic, "kludge-y" approach to the mind. Thus in tackling the gambler's fallacy, I'll critique the mechanistic approach at its strongest: if even here a resource-rational explanation can better explain the empirical data, that should significantly shift our opinions in the overall plausibility of the rational approach.

To clear the ground, it's worth highlighting the reasons to doubt the standard mechanistic account. This theory says that people are unable to do proper statistical reasoning, mistakenly believing in the "law of small numbers": that short sequences will be representative of long ones (Tversky and Kahneman 1971). This theory arguably hasn't stood the test of time. (i) It struggles to explain detailed empirical predictions like (1)–(5) above. (ii) It leaves it puzzling why gambler's reasoning persists in high-stakes settings (Chen et al. 2016; Jiao 2017) and isn't weeded out by (evolutionary or market-based) selection processes. (iii) It fits poorly with the evidence that many cognitive processes exhibit striking degrees

[4] Kahneman and Tversky 1972; Tversky and Kahneman 1974; Kahneman et al. 1982; Kahneman and Tversky 1996; Fine 2005; Ariely 2008; Kahneman 2011; Thaler 2015; Lewis 2016; Mandelbaum 2018; Williams 2021.

[5] Gopnik 1996, 2020; Chater and Oaksford 1999; Tenenbaum and Griffiths 2006; Griffiths et al. 2008; Tenenbaum et al. 2011; Bonawitz et al. 2014b,a; Vul et al. 2014; Ruggeri and Lombrozo 2015; Icard 2016, 2017, 2023; Bhui and Gershman 2018; Bhui et al. 2021; Gershman 2021; Hartley 2022

[6] Cohen 1981; Kelly 2004, 2008; Crupi et al. 2008; Koralus and Mascarenhas 2013; Nebel 2015; Hahn and Harris 2014; O'Connor and Weatherall 2018, 2019; Weatherall and O'Connor 2020; Hedden 2019; Doody 2020; Fisher 2020; Fisher and Mandel 2021; Karlan 2021; Kinney and Bright 2021; Henderson and Gebharter 2021; Thorstad 2021; Dorst and Mandelkern 2022; Dorst 2023b

[7] Others include hindsight bias (Fischhoff and Beyth 1975—but see Hedden 2019), the conjunction fallacy (Tversky and Kahneman 1983—but see Dorst and Mandelkern 2022), apparent overconfidence (Lichtenstein et al. 1982—but see Moore and Healy 2008; Dorst 2023a), and the base-rate fallacy (Kahneman and Tversky 1973—but see Schwarz et al. 1991; Koehler 1996; Lejarraga and Hertwig 2021).

of statistical sophistication (see footnote 5). And (iv) it makes it mysterious why gambler's reasoning is *positively* correlated with cognitive sophistication: infants do hot-hands reasoning, while gambler's reasoning shows up in elementary schools and peaks among college students (Estes 1962; Craig and Myers 1963; Bogartz 1965; Derks and Paclisanu 1967; Chiesi and Primi 2009); meanwhile, reaction times are quicker for hot-hands than for gambler's reasoning, which becomes *less* common under time pressure or cognitive load (Diener and Thompson 1985; Militana et al. 2010; Tyszka et al. 2017; Braga et al. 2018). Upshot: there's both empirical and theoretical motivation for a rational explanation.

I'll offer one. In doing so, I'll illustrate a contrast between two very different types of Bayesian models of cognition. *Bounded-optimal models* impose architectural constraints on agents—such as limitations of memory (Icard 2021) or representational precision (Bhui and Gershman 2018)—and then derive predictions from the best that *can* be done (Thorstad 2023). Meanwhile, *error-optimal models* assume that the agent has made some sort of *mistake*—for example, ruling out the possibility that the outcomes are independent (Barberis et al. 1998; Rabin 2002; Rabin and Vayanos 2010) or ignoring *how often* a given subsequence has occurred (Hahn and Warren 2009)—and then derive predictions from what a Bayesian agent would do *given* that mistake. Although formally similar and often lumped together, these approaches are very different. The latter can be unifying—showing that looks like many errors are reducible to a single one—especially if there's an independent explanation of why the error would persist. Nevertheless, error-optimal models are not *rationalizing* models: they explain behavior without *justifying* it. Thus they lack many of the virtues of bounded-optimal models: they don't explain why the mistake persists; they (arguably) don't give a computational-level explanation for why the mind works that way (Marr 1982); and—from a philosophical point of view—they don't reconcile the manifest and scientific images (Sellars 1956). In particular, they are reflectively unstable: they don't offer an explanation of what we're up to that we can coherently hold in mind, as we feel the pull of gambler's reasoning. Perhaps this is as it should be—that pull *might* just be a "mental tick" that can't be turned off, a persistent illusion of the mind. But—unlike Müller-Lyer lines, phobias, or earworms—gambler's reasoning doesn't *feel* like that. It'd be nice to have a theory that explained why.

2 The Model, Informally

I'll first explain how the model works informally; future sections will dive into the details.

Consider any process that can have one of two outcomes in each trial—call one a "hit" (success, heads, 1, Prius,...), and the other a "miss" (failure, tails, 0, Jeep,...). Let's suppose the long-run hit rate is known to be 50%. (§6.1 shows that the reasoning generalizes.)

Focus on Baylee. She knows that the parking outcomes are "random" in the sense that they are hard to predict, and that in the long run, the Prius wins 50% of the time. That's leaves open three classes of hypotheses about how the likelihoods shift.

First, parking could be **Steady:** each day there's a 50% chance the Prius wins the spot, independently of the previous days. The standard example of a Steady process is a coin.

Second, parking could be **Switchy:** after the Prius wins a few, the Jeep becomes more likely to win; and after the Jeep wins a few, the Prius becomes more likely to win. In other words: after a streak, a switch becomes "due"—the process tends to balance out. So long as these shifts in probability are symmetric around 50%, the long-run hit rate will be 50%. Switchy processes are those that involve some sort of depletion, fatigue, or strategic response. The parking battle would be Switchy if each driver gets frustrated when they lose, and so leaves for work a bit earlier the next day. The standard example of a Switchy process is drawing from a deck of cards without replacement: as you draw a series of red cards, a black card becomes increasingly likely; and vice versa. (This is a *hypergeometric* distribution.)

Third, the parking could be "streaky"—or, as I prefer to call it, **Sticky:** after the Prius wins a few, it becomes more likely to win again; after the Jeep wins a few, *it* becomes more likely to win again. In other words: after a streak, the the current winner becomes "hot". So long as these shifts in probability are symmetric around 50%, the long-run hit rate will still be 50%: long streaks of Priuses are likely, but no more likely than comparably-long streaks of Jeeps. Sticky processes are those that involve some sort of accumulation or "momentum" in which success breeds success ("rich get richer"), or performing well (poorly) puts you into (out of) "the groove". The parking battle would be Sticky if each driver gets disheartened when they lose, and so tries less hard the next day. The standard example of a Sticky process is basketball shots: when players make a shot, they become "hot", so more likely to make the next one.[8]

Switchy and Sticky processes vary along many dimensions: their long-run hit rates; how quickly, how far, and how exactly their probabilities vary around that hit rate; and so on. But generalizing over them, Switchy and Sticky processes are quite common (see §3). Thus when Baylee confronts a new causal system, she should usually be unsure whether it's Switchy, Steady, or Sticky—but pretty sure it's one of the three. Since each would predict the same *long*-run hit rates, such information doesn't tell her anything about which it is.

Now: two mathematical observations.

First observation: whether she should do gambler's or hot-hands reasoning depends on the exact *balance* of her uncertainty between our three hypotheses. For instance, suppose there's been a streak of Priuses. How confident should she be that the streak will continue? Her opinion should be a weighted average of what it would be if she knew it were Switchy, Steady, or Sticky—with weights determined by how confident she is in each (see §3.1). If she knew it were Switchy, she'd be less than 50%-confident the streak will continue. If she knew it were Sticky, she'd be (symmetrically) *more* than 50%. If she knew it were Steady, she'd be exactly 50%. Thus whenever she's more confident in Switchy than Sticky, her uncertainty will average out to being *less* than 50%-confident the streak will continue—she'll do gambler's reasoning. Meanwhile, whenever she's less confident in Switchy than in Sticky, she'll average out to being *more* than 50%-confident it'll continue—she'll do hot-hands reasoning. How far her opinions will deviate from 50% will depend on both (1) the balance of her opinions between Switchy and Sticky (i.e. $P(Swi|Swi\ or\ Sti)$), and (2) her degree of confidence in Steady. The more imbalanced, the further her opinions in a continuation will

[8] You might think this has been shown to be a myth (Gilovich et al. 1985)—but it turns out that the classic analysis failed to take into account a subtle selection effect; when you do, the original data *supports* the existence of a hot hand (Miller and Sanjurjo 2018b).

deviate from 50%; but this imbalance has less effect the more confident she is in Steady (i.e. in $\neg(Swi\ or\ Sti)$).

Upshot: whether and to what degree she should do gambler's or hot-hands reasoning depends on the precise balance of her uncertainty between Switchy, Steady, and Sticky—and, therefore, on the precise data she observes (or remembers) about the process. Since it'll be rare for that data to be *precisely* balanced in how much it support Switchy and Sticky, small degrees of gambler's and hot-hands reasoning will be common.

Yet I've claimed more—namely, that both the *majority* and the *average* will commit gambler's reasoning. Why? The situation looks symmetric. Surprisingly, it's not.

Second observation: data generated from a Steady process tends to provide more evidence against Sticky than against Switchy. More precisely: the likelihood distribution over data generated by Steady is *closer to that generated by Switchy* than to that generated by Sticky; thus it tends to rule out Sticky more quickly than Switchy.

Why? Imagine we're observing outcomes: Prius, Prius, Jeep, Prius, Jeep, Jeep, Jeep. Given the current streak, both Switchy and Sticky deviate substantially from Steady's 50% likelihood of another Jeep. If there *is* another Jeep, both will stay deviated; but if there's a Prius, both will return toward Steady's likelihood (50%) on the *following* outcome. And notice: Switchy makes a Prius more likely than Sticky. In other words: although both hypotheses deviate from Steady's likelihoods, Switchy *wants to return toward them*, while Sticky *wants to stay deviated*. That is: both Switchy and Sticky will increasingly deviate from Steady in what to expect next as the current streak grows in length. But it takes time for these deviations to build up, and Switchy tends to keep the streaks relatively short, while Sticky tends to generate long streaks. As a result, Sticky deviates more.

This difference is hard to visualize if we plot the likelihoods of every sequence, but easy to spot if we look at summary statistics. For example, with sequences of length 50 (and using the versions of Switchy and Sticky discussed in §3), Figure 1 plots the likelihoods of various total numbers switches (top left), average streak lengths (top right), and total numbers of hits (bottom left). In each cases there is substantially more overlap between Steady and Switchy than there is between Steady and Sticky.

Why does this matter? It implies that if Bayesians observe data from a Steady process, they'll (on average) *quickly* get evidence against Sticky, and *slowly* get evidence against Switchy. Given enough data, they'll become confident of Steady, but along the way, they'll (on average) be more confident of Switchy than Sticky—they will *asymmetrically* converge toward the truth. Figure 1 (bottom right) plots the mean posteriors in Switchy, Steady, and Sticky for Bayesians who start out uniformly unsure between the three and then each update on outcomes of sequences from a Steady distribution. Notice that the average posterior in Switchy is always higher than the average posterior in Sticky.

Now recall: whenever a Bayesian assigns more probability to Switchy than to Sticky, they'll do gambler's reasoning. Thus, on average, our Bayesians will do so. If they have perfect memories, these deviations from a 50%-probability of switching will (persist but) become quite small, for they will eventually become quite confident of Steady. But—as we'll see below (§4)—if their memories are limited, a substantial degree of gambler's reasoning can persistent given arbitrarily large amounts of data.

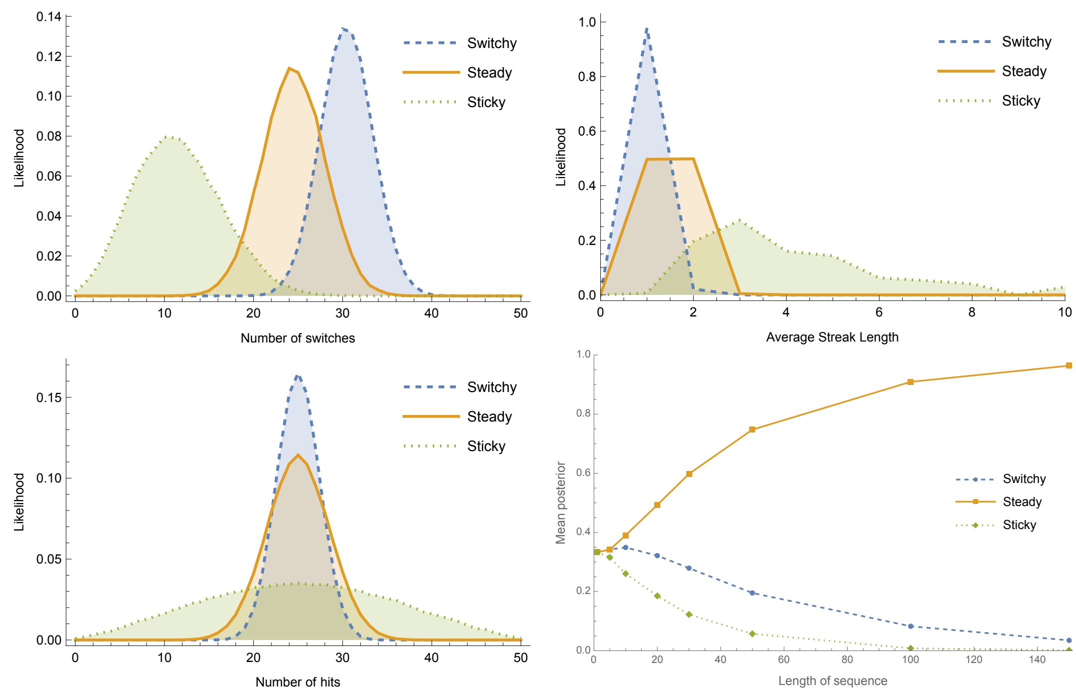

Figure 1: Among sequences of length 50, the likelihoods for total number of switches (top left), average streak length (top right), and total number of heads (bottom left). There is substantially more overlap between Steady and Switchy than between Steady and Sticky. As a result, when conditioning on the full sequence, the mean posteriors in Sticky fall faster than those in Switchy (bottom right).

More is true. Because of this asymmetric convergence, our Bayesians will exhibit our five empirical features of gambler's and hot-hands reasoning. Here's why.

Nonlinear Expectations: After they observe a couple Priuses in a row, they'll know that *if* the process is Switchy, it's very primed to switch—and so will at first increasingly expect a switch. But as the streak length grows, the streak itself provides more and more evidence for *Sticky*. Given a long-enough streak, this evidence will overwhelm their initial favoring of Switchy and lead them to increasingly perform hot-hands reasoning.

Experience-Dependence: For processes that they have little experience with, they'll only *slightly* favor Switchy over Sticky—thus even a short streak will provide enough evidence to lead to hot-hands reasoning. But for processes that they have a lot of experience with, they'll robustly favor Switchy over Sticky—thus it'll take a long streak before they do hot-hands reasoning.

Sequence-Dependence: Since sequences with lots of switches provide evidence for Switchy while those with few provide evidence for Sticky, they'll be more inclined to do gambler's reasoning after the former and hot-hands reasoning after the latter.

Too Switchy: They'll generate sequences with more than 50% switches because they're (on average) more confident of Switchy than Sticky, and so their probabilities of a switch will be higher than 50%. This deviation will get starker with longer streaks, because they know that *if* it's Switchy, longer streaks are more likely to switch.

Format-Dependence: The distribution of probability-estimates for a switch will be a fairly tight bell curve—meaning that even if the *average* probability of a switch is only slightly above 50%, far more than half of the subjects will have probabilities above 50%. Thus when making binary predictions the *proportion* of people doing gambler's reasoning will be well over 50%—inducing apparent format-dependence.

Upshot: given plausible assumptions, the causal-uncertainty hypothesis predicts a variety of subtle empirical trends found in studies of the gambler's fallacy.

2.1 Comparisons

How does the causal-uncertainty hypothesis compare to other theories? Its main advantages are that it can both (1) predict nonlinearities, and (2) offer a resource-rational explanation of *why* gambler's reasoning persists.

Its empirical predictions are closest to the models of Rabin (2002) and Rabin and Vayanos (2010), which assume that people have (dogmatic) beliefs both that the process is Switchy and the overall hit rate can vary over time. These are the only models I know of that also predict nonlinearities in expectations: a short streak makes a switch likely, but a long streak suggests the hit rate has changed. Predicting nonlinearity is often taken to be a significant mark in its favor (Asparouhova et al. 2009). Still, the model faces two main objections. First, it doesn't have a theory-internal way to predict experience- or format-dependence. Second, it both relies crucially on model-misspecification—people dogmatically hold onto an incorrect causal model of the process—and also offers no explanation for why people's (incorrect) causal model is a *depletion* (Switchy) one, rather than an accumulation (Sticky) one. In contrast, the causal-uncertainty hypothesis builds in no such asymmetries, nor does it assume dogmatic agents. Rather, it exploits mathematical asymmetries between Switchy and Sticky to explain why nonlinearities (etc.) would emerge.

Barberis et al. (1998) propose a "regime-shifting" model: people are (incorrectly) certain that the process is either Switchy or Sticky, but think it can shift back and forth. This model accurately predicts sequence-dependence and a shift from gambler's to hot-hands reasoning. But it can't predict *nonlinearity* in expectations (Asparouhova et al. 2009; Rao and Hastie 2023)—instead predicting that people's probability of a continuation will increase monotonically with streak-length. It also lacks a theory-internal way of predicting experience- and format-dependence. Of course, there's clearly something correct about this model: the parking battle might be Switchy at one time, and become Sticky later on. Thus a generalized version of the causal-uncertainty hypothesis should allow our Bayesians to be unsure whether and when the process *shifts* between Switchy, Steady, and Sticky. I avoid this complication to show that we don't need this complication to make our predictions.

Hahn and Warren 2009 give an explanation based on the *overlapping words paradox*: sequences that can overlap with themselves are less likely to occur in any given sequence, since they are more likely to occur multiple times if at all. For example, of the 16 different four-toss sequences, 'HHT' occurs in 4 of them (HHHT, HHTH, HHTT, THHT), while HHH occurs in only 3 (HHHT, THHH, HHHH) because it occurs twice in HHHH. Thus if people simply track *whether* (rather than how often) they've seen a given string in a certain interval of time, this will lead them to think that regular strings like HHH are less common than irregular ones like HHT, causing gambler's reasoning. Though fascinating, this proposal can't predict nonlinearity or format-dependence, and it doesn't explain *why* people would ignore how many times they've seen a string.

Miller and Sanjurjo 2018a exploit a subtle selection effect (Miller and Sanjurjo 2018b) to show that people who attend to the proportion of successes following a success—but partially neglect sample size (Griffin and Tversky 1992)—will do gambler's reasoning. Though interesting, this model can't explain nonlinearity or format-dependence, and doesn't explain why people would neglect sample size.

Pinker 1997, building off others (Gigerenzer and Hoffrage 1995; Cosmides and Tooby 1996), points out that many real-world "random" processes are Switchy, so that gambler's reasoning may be a reasonable overgeneralization. Though my model uses this idea, an important piece is missing: there are *also* many real-life processes that are *Sticky*. Absent a reason to think that Switchy mechanisms are more common, it's unclear why people's average- and modal responses would be *gambler's* rather than *hot-hands* reasoning. Using the above mathematical asymmetries, my theory explains why.

All of these theories have empirical support, and many of them likely play a role in explaining the full breadth of empirical findings. Nevertheless, I'll argue that the causal-uncertainty hypothesis stacks up well in its theoretical motivation, explanatory power, and empirical coverage. Don't dismiss it.

3 The causal hypotheses

Suppose—until §6.1—that Baylee knows that the long-run hit rate is 50%. This is consistent with Switchy, Steady, or Sticky.

Switchy and Sticky can vary in both their *extremity*—how far the probabilities ever deviate from 50%—and their *speed*—how quickly they build up to that probability. We can diagram them with *Markov chains*: a set of states along with transition-probabilities between states.[9] Letting the state encode the current streak, labeled arrows encode transition probabilities, and using a '1' for a Prius (hit) and a '0' for a Jeep (miss), the top row of Figure 2 gives examples of a 2-step-Switchy$_{70\%}$ (left) and a 2-step-Sticky$_{70\%}$ (right) chain, as well as the "2-step" Steady chain (whose probabilities don't shift). These diagrams are hard to read. Better are *(row-)stochastic matrices*, in which row i column j encodes the probability of transitioning from state i to state j. To do this, we need to order the states in a canonical way. Since in these Markov chains states track only the streak-length, let's put the longest miss-streak that it tracks on the left, the shortest streaks in the middle, and the longest hit-streak on the right; so 2-step chains are ordered $\langle 00, 0, 1, 11\rangle$. The bottom row of Figure 2 displays the same hypotheses in this notation.

Until §6, I'll focus on 5-step-90% versions of Switchy and Sticky, displayed in Figure 3. There are many variations of Switchy and Sticky that we could investigate. First, we've so far focused on *streak chains*: the probabilities are determined by the current streak. We could instead use *counting chains* that track the full (recent) sequence—for example, to precisely

[9]Thus I follow Oskarsson et al. 2009 in using Markov models to represent people's uncertainty about causal hypotheses. But instead of imputing a *single* model to them, I'll track their uncertainty over several.

2-step-Switchy$_{70\%}$: **2-step-Steady:** **2-step-Sticky$_{70\%}$:**

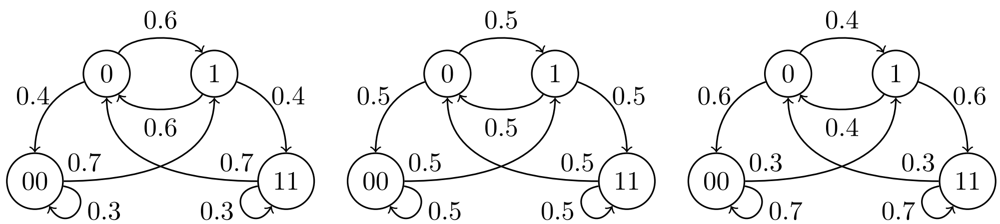

$$\begin{pmatrix} 0.3 & 0 & 0.7 & 0 \\ 0.4 & 0 & 0.6 & 0 \\ 0 & 0.6 & 0 & 0.4 \\ 0 & 0.7 & 0 & 0.3 \end{pmatrix} \quad \begin{pmatrix} 0.5 & 0 & 0.5 & 0 \\ 0.5 & 0 & 0.5 & 0 \\ 0 & 0.5 & 0 & 0.5 \\ 0 & 0.5 & 0 & 0.5 \end{pmatrix} \quad \begin{pmatrix} 0.7 & 0 & 0.3 & 0 \\ 0.6 & 0 & 0.4 & 0 \\ 0 & 0.4 & 0 & 0.6 \\ 0 & 0.3 & 0 & 0.7 \end{pmatrix}$$

Figure 2: *Top:* Markov diagrams for 2-step-Switchy$_{70\%}$, 2-step-Steady, and 2-step-Sticky$_{70\%}$ hypotheses. *Bottom:* stochastic-matrix representations of each hypothesis.

5-step-Switchy$_{90\%}$ **5-step-Sticky$_{90\%}$**

$$\begin{pmatrix} .1 & 0 & 0 & 0 & 0 & .9 & 0 & 0 & 0 & 0 \\ .18 & 0 & 0 & 0 & 0 & .82 & 0 & 0 & 0 & 0 \\ 0 & .26 & 0 & 0 & 0 & .74 & 0 & 0 & 0 & 0 \\ 0 & 0 & .34 & 0 & 0 & .66 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & .42 & 0 & .58 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & .58 & 0 & .42 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & .66 & 0 & 0 & .34 & 0 & 0 \\ 0 & 0 & 0 & 0 & .74 & 0 & 0 & 0 & .26 & 0 \\ 0 & 0 & 0 & 0 & .82 & 0 & 0 & 0 & 0 & .18 \\ 0 & 0 & 0 & 0 & .9 & 0 & 0 & 0 & 0 & .1 \end{pmatrix} \quad \begin{pmatrix} .9 & 0 & 0 & 0 & 0 & .1 & 0 & 0 & 0 & 0 \\ .82 & 0 & 0 & 0 & 0 & .18 & 0 & 0 & 0 & 0 \\ 0 & .74 & 0 & 0 & 0 & .26 & 0 & 0 & 0 & 0 \\ 0 & 0 & .66 & 0 & 0 & .34 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & .58 & 0 & .42 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & .42 & 0 & .58 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & .34 & 0 & 0 & .66 & 0 & 0 \\ 0 & 0 & 0 & 0 & .26 & 0 & 0 & 0 & .74 & 0 \\ 0 & 0 & 0 & 0 & .18 & 0 & 0 & 0 & 0 & .82 \\ 0 & 0 & 0 & 0 & .1 & 0 & 0 & 0 & 0 & .9 \end{pmatrix}$$

Figure 3: Transition matrices for the 5-step-90% versions of Switchy and Sticky used in main simulations.

model drawing (red/black) cards from a deck without replacement. Counting chains are natural and common, but they're harder to use because they're exponential in the length of sequence they track: for a 6-step counting chain, we'd need $2^6 = 64$ states; a 6-step streak chain needs only 12. (For a 52-card deck, we'd need 2^{52} states!) I'll show that our results generalize to counting chains in §6.4.

Second, we've focused on chains where the outcome-sequence fully determines the state, but in many systems it'll have a *stochastic* relationship to the state. The Prius-driver might *tend* to get up earlier if he misses the spot, but that tendency would be modulated by other factors. We could model this using *hidden* Markov models (Oskarsson et al. 2009). But we needn't. Instead, we can understand the probabilities in our chains as our agent's conditional *beliefs* about what'll happen next, conditional on the process being Switchy/Steady/Sticky—marginalizing over such more-complex, stochastic models.

Finally, we can imagine causal hypotheses that are neither Switchy, Steady, nor Sticky: perhaps winning *twice* in a row makes the Prius less likely to win, but winning *three* times in a row makes it *more* likely to win. Though possible, such processes are less natural and less common. It won't affect our results if our agents assign them some (smaller) amount of probability, so long as they are fairly confident the process is either Switchy, Steady, or Sticky. For simplicity, I'll consider agents who ignore these alternatives.

Switchy and Sticky processes are commonplace. Switchy processes are those that involve depletion, fatigue, strategy, or fairness: decks of cards, raffles, traffic lights, drawing lots, Q&As (who gets called on), electoral politics (voters get dissatisfied), races (the leader gets tired), board games (the leader gets teamed-up against), the stock market(?) (if big gains lead to corrections), and so on. Meanwhile, Sticky processes are those that involve accumulation, momentum, rich-get-richer ("Matthew") effects, or being in a "groove" vs. a "funk": reputation; trends; consumer loyalty; getting publications, jobs, or money; the stock market(?) (if big gains lead to a bullish market); skilled musical, athletic, or intellectual performances; mental or physical health; moods (Eldar et al. 2016); and even everyday life—if you stub your toe, spill your coffee, and get into an argument all before noon, the afternoon is likely to be rough. (As Deena Rose tells us: "Some days are diamonds, some days are stone.")

For modeling tractability, I'll focus on Bayesians who are certain that the process is either 5-step-Switchy$_{90\%}$, Steady, or 5-step-Sticky$_{90\%}$. §6 will show that the results are robust.

3.1 What a Bayesian expects

Baylee is unsure whether the parking war is 5-step-Switchy$_{90\%}$ (*Swi*), Steady (*Ste*), or 5-step-Sticky$_{90\%}$ (*Sti*). Thus if the Prius is on a streak, what she thinks will happen next depends on the *balance* of her uncertainty. By total probability and the Principal Principle (Lewis 1980), her probability for the Prius winning is:

$$\begin{aligned} P(Prius) &= P(Swi)\cdot P(Prius|Swi) &&+ P(Ste)\cdot P(Prius|Ste) &&+ P(Sti)\cdot P(Prius|Sti) \\ &= P(Swi)\cdot(0.5-c) &&+ P(Ste)\cdot(0.5) &&+ P(Sti)\cdot(0.5+c) \end{aligned}$$

The constant c is the same given Switchy and Sticky because they're symmetric in their deviations from 50% (or rather: because Baylee has no evidence for an asymmetry). Its value depends on how long the streak is—given our hypotheses, for a streak of 1, $c = 0.08$; for 2, $c = 0.16$, etc. Thus Baylee's probability for a continuation is a weighted average of $0.5-c$, 0.5 and $0.5+c$—with weights determined by her probabilities of Switchy, Steady, and Sticky. She'll be less than 50%-confident in a continuation—expecting a switch, and performing gambler's reasoning—iff she's more confident of Switchy than Sticky: $P(Swi) > P(Sti)$. And she'll be more than 50%-confident in a continuation—expecting a continuation, and performing hot-hands reasoning—iff she's less confident of Switchy than Sticky: $P(Swi) < P(Sti)$.[10]

Thus treating the process as *precisely* independent will be uncommon—it requires the rare event that Baylee's evidence (the sequence she's observed) should make her precisely equally confident in Switchy and Sticky. Most Bayesians will do *either* gambler's or hot-hands reasoning. How extreme that reasoning will be depends not only on the *balance* of their probabilities for Switchy and Sticky, but also on how confident they are of Steady: the more weight in the weighted-average that's put on the 50%-term, the less their probabilities can deviate from 50%. For example, if Baylee's 40-40-20 between Switchy-Steady-Sticky and she's seen 3 Priuses in a row, her probability of a continuation will be $0.4(0.26) + 0.4(0.5) +$

[10]This reasoning requires the c to be *precisely* equal on both sides. Often it won't be, for example when we move to hit-rates other than 50% (§6.1). But the same dynamics will apply: whenever she's *sufficiently* more confident in Switchy than Sticky, she'll do gambler's reasoning; and vice versa.

$0.2(0.74) \approx 0.45$. But if—holding $\frac{P(Switchy)}{P(Sticky)}$ fixed—she's 10-85-5 between Switchy-Steady-Sticky, her probability of a continuation will be roughly 0.49.

Upshot: whether and to what degree Bayesians commit gambler's or hot-hands reasoning depends on the balance and degree of their uncertainty between Switchy, Steady, and Sticky. Still, we have yet to find an *asymmetry*. Why would the majority—or the average—commit gambler's reasoning?

4 Asymmetric convergence and memory limits

If the process is Steady, the sequences of outcomes is determined by the Steady likelihoods: $P(\cdot|Ste)$. Suppose that our Bayesians begin uniform $({}^1\!/_3, {}^1\!/_3, {}^1\!/_3)$ over (Swi, Ste, Sti), and then each update by conditioning on a (different) sequence generated by Steady likelihoods. **Fact:** *on average*, they'll rule out Sticky faster than they will rule out Switchy.

Why? When Switchy deviates from 50%-likelihood of a Prius on the next day, it makes a switch more likely—which in turn would return the likelihoods back toward 50%. Meanwhile, when Sticky deviates from this 50%-likelihood, it makes a *continuation* more likely, which in turn would keep the likelihoods far from 50%. As a result, Switchy's distribution over sequences is more similar to Steady's than Sticky's is.

It's difficult to visualize this across all possible sequences (though see the Mathematica notebook); instead, we can look at measures of the distance between the likelihood distributions. The top of Figure 4 does this for Euclidean distance and KL divergence, as the sequence-length grows. As can be seen, Switchy is much closer to Steady than Sticky is.[11] The KL divergence from the true distribution P to another Q estimates how much more Bayesians' confidence will (on average) increase in P over Q when presented with data from P.[12] So as our Bayesians see more data, they're confidence in Steady will (on average) increase more over Sticky than over Switchy (Figure 4, bottom left)—and in fact the ratio of the average posteriors in Switchy to Sticky grows exponentially (Figure 4, bottom right). As we'll see, this asymmetry predicts our empirical trends.

However, recall that as Bayesians become more confident of Steady, the *degree* to which they'll exhibit gambler's or hot-hands reasoning diminishes. In the above plots, note that by the time they've seen 150 outcomes, the average posterior in Steady is already around 95%. Thus, for instance, if Baylee has seen 150 outcomes and then sees the Prius win three times in a row, her probability for a continuation will (on average) be around 49%—she'll do gambler's reasoning, but barely. So you might object: surely most people have seen more than 150 coin tosses! The causal-uncertainty hypothesis seems to make implausible predictions about the degree to which people will exhibit gambler's reasoning. Right?

[11] To calculate these, list all 2^n possible sequences of a given length n in some order, and then write the likelihood distribution as a vector of length 2^n where the ith entry is the probability of getting the ith sequence. The Euclidean distance between $\vec{x} = (x_1, ..., x_m)$ and $\vec{y} = (y_1, ..., y_m)$ is the distance of the line between the two vectors in Euclidean space: $d(\vec{x}, \vec{y}) = \sqrt{(x_1 - y_1)^2 + \cdots + (x_m - y_m)^2}$. The KL divergence from distribution $\vec{x}$ to distribution $\vec{y}$ is $\vec{y}$'s expectation of the logarithmic difference between $\vec{y}$ and $\vec{x}$: $D_{KL}(\vec{y}||\vec{x}) = \sum_i y_i \log(\frac{y_i}{x_i})$.

[12] See e.g. Cover and Thomas 2005, Ch. 11; many thanks to Yonathan Fiat for help here.

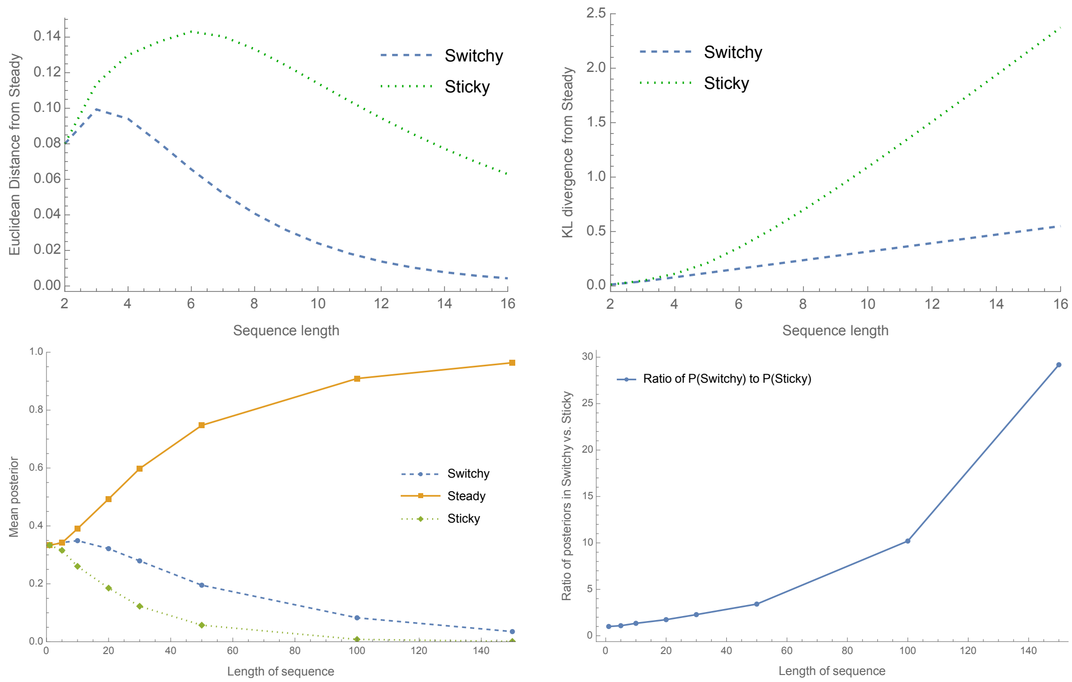

Figure 4: Limited data, asymmetric closeness and convergence. *Top:* Euclidean distances (left) and KL divergences (right) from Steady likelihoods to Switchy and Sticky likelihoods. *Bottom:* Since the Switchy distribution is closer to the Steady one, it on average gets ruled out slower—note that the dashed line is always higher than the dotted line (left), and in fact that the ratio grows exponentially (right).

Wrong. First, many processes for which people exhibit gambler's reasoning are ones where they haven't seen a large number of outcomes—horse races, board games, elections, urn draws, etc. Second, clearly those who *have* seen many outcomes (coin tosses or die rolls or slot machines) *don't remember the full sequence*. This is both obvious, and obviously not fallacious—nobody's claimed to discover the "fallacy of forgetting". Memory limits are clear cases of a cognitive-architectural resource constraint; the normative question is whether people form sensible beliefs *given* their limited memories (Aronowitz 2019; Thorstad 2023).

Our question: will rational processing given limited memory to lead to gambler's reasoning? From a high level, we can see that the answer will be *yes*, regardless of how we model limited memories. Memory is a way of *answering a (set of) question(s)* about what you've seen. A question is a partition of the possibilities (Hamblin 1976)—a way of dividing the possible sequences you might've observed into cells that each agree on the answer to the question. Equivalently, a question is a *summary statistic*: a function from data to a number that tells you some feature of that data—for example, X = *the number of tosses that landed heads*, or Y = *the length of the longest streak*. Thus having a limited memory about a sequence is going to function very much like observing limited data. (Indeed, observing n outcomes is equivalent to having a limited "memory" of the full sequence—namely, by having only the answer to the question "What were the first n outcomes?".) Since Steady's full distribution is closer to Switchy's than to Sticky's, it follows that limited memory will also lead to asymmetric convergence, whichever summary statistic (model of limited memory) we use.

For concreteness, I'll work with the following simple model—call it the **unpacking model** of memory. Baylee observes an arbitrarily long sequences S of outcomes—say 1000. The *1-unpacked* memory of S is the answer to the question, "How many hits were there in all 1000 outcomes?". The *2-unpacked* memory of S is the answer to the conjunctive question, "How many hits were there in the first half (500) of the outcomes, *and* how many hits were there in the second half?". The *3-unpacked* memory of S is the answer to the question, "How many hits were there in the first third (333) of the outcomes, *and* how many hits were there in the second third, *and* how many hits were there in the third third?" And so on.

Though simplistic, this model lets us smoothly vary how detailed Baylee's memory is. 1-unpacking gives a very coarse summary; 1000-unpacking recovers the entire sequence; and we can vary between these extremes. Moreover, this model seems to capture what people *do* tend to remember about "random" processes like coins—they can't remember what happened on any particular toss, but they can remember things like "it lands heads roughly half the times" and "in most of the sequences I've seen, the proportions of heads have been close to 50%". Indeed, there's empirical evidence that people are good at learning the proportions and means of a series of observations (Peterson and Beach 1967).[13]

Does the unpacking-model lead to asymmetric convergence? Yes. We can now *visualize* how Steady likelihoods are closer to the Switchy ones than to the Sticky ones—this was displayed in Figure 1 (bottom left) on page 8. Again, we can plot the Euclidean distances and KL divergences for the sub-parts of our unpacked memory (the numbers of hits in subsequences of varying lengths); the results are displayed in the top of Figure 5. Finally—and most importantly—we can simulate Bayesians who start out uniform over Switchy, Steady, and Sticky, update on various depths of memory-unpacking from (different) 1000-toss sequences, and then plot their mean posteriors. They are on average more confident of Switchy than Sticky (Figure 5, bottom left), and the degree to which they favor Switchy over Sticky grows as the memory is further unpacked (Figure 5, bottom right).

Importantly, even with deep unpacking we should still expect a substantial *degree* of gambler's-fallacy reasoning. For example, the mean 20-unpacked agent—who remembers the proportion of heads in 20 separate sets of 50 tosses—gives a probability of 47.2% to a streak of 2 continuing and 45.8% to a streak of 3 continuing.

That's asymmetric-convergence. Why does it predict our empirical findings?

5 Fallacious features

Recall our five empirical findings: (1) as a streak grows in length, people have *nonlinear expectations* for the probability of it continuing; (2) these nonlinear expectations are *experience-dependent*, exhibiting different shapes for different causal systems; (3) people's predictions are *sequence-dependent*, varying with the number of switches in the sequence

[13] Worry: perfectly remembering the hit-rate of a long sequence takes more memory than perfectly remembering the hit-rate of a short sequence, right? Right. But *approximately* remembering the hit-rate of a sequence turns out to be largely independent of sequence-length (see Icard 2023, §7; Alon et al. 1996). For simplicity I'll ignore this complication. It's an interesting and open question how different models of memory would affect our results (e.g. Icard 2023, Ch. 3).

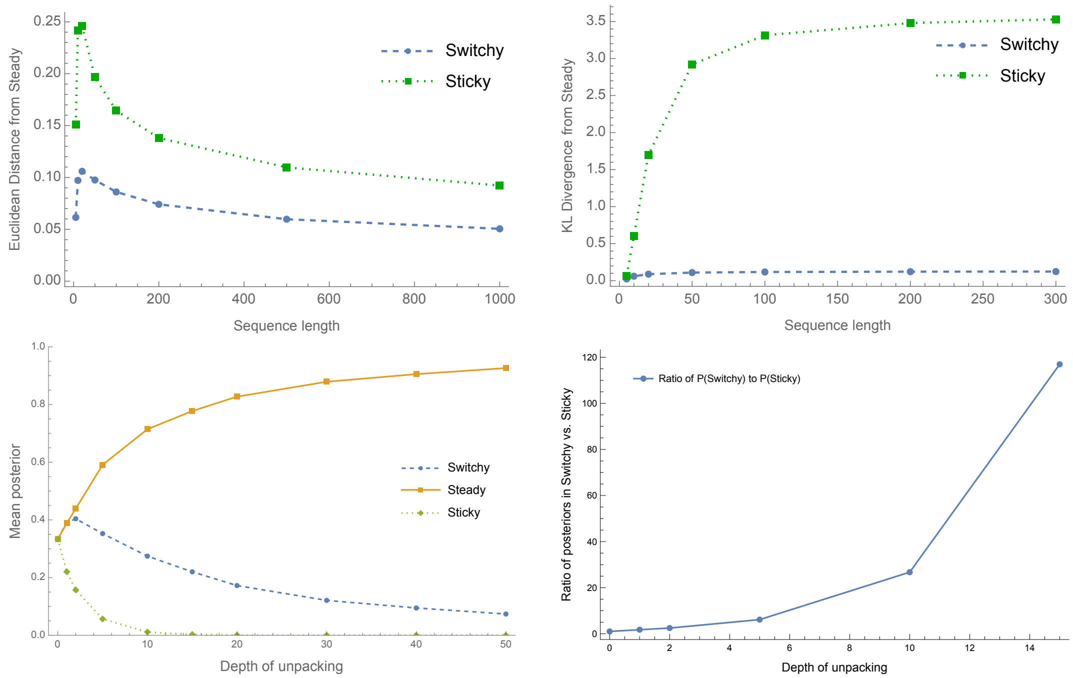

Figure 5: Limited memory, asymmetric closeness and convergence. *Top:* Euclidean distances (left) and KL divergences (right) from Steady likelihoods to Switchy and Sticky likelihoods. *Bottom:* Since the Switchy distribution is closer to the Steady one, it on average gets ruled out slower—note that the dashed line is always higher than the dotted line (left), and in fact that the ratio grows exponentially (right).

they just saw; (4) people generate "random" sequences that are *too switchy*; and (5) people's predictions appear *format-dependent*, showing starker rates of gambler's reasoning for binary prediction than for probability estimates.

I'll now show why the causal-uncertainty hypothesis predicts each of these patterns. For robustness, I'll show both simulations of limited data (conditioning on the full outcome of a short sequence) and limited memory (conditioning on a partially-unpacked long sequence). I won't attempt any quantitative model-fitting or -comparisons. Instead I'll focus on showing a robust qualitative match to empirical data, and that there are parameterizations which give reasonable quantitative predictions. I leave model-fitting for future work.

5.1 Nonlinear Expectations

The findings: One of the most intriguing, widely-confirmed, and difficult-to-model findings is that people display nonlinear dynamics in their expectations about whether a streak will continue. As streaks grow in length, they first increasingly exhibit gambler's reasoning (predicting a switch); but as the streak grows long enough, their gambler's-reasoning first tapers off and then is replaced by increasing rates of hot-hands reasoning (predicting a continuation). This has been found in many settings for many causal systems, ranging from coins to stock markets to lotteries (see footnote 1). It's difficult to predict, and is the main empirical argument in favor of Rabin 2002's model.

The most careful studies I know of are Asparouhova et al. 2009 and Rao and Hastie 2023, which both (in highly controlled settings) elicit probability estimates of a continuation given streaks of varying lengths. Their findings are reconstructed in the top left of Figure 6—Asparouhova et al.'s results reconstructed from their Table 6, and Rao and Hastie's from the raw data (https://osf.io/2d93t/).[14]

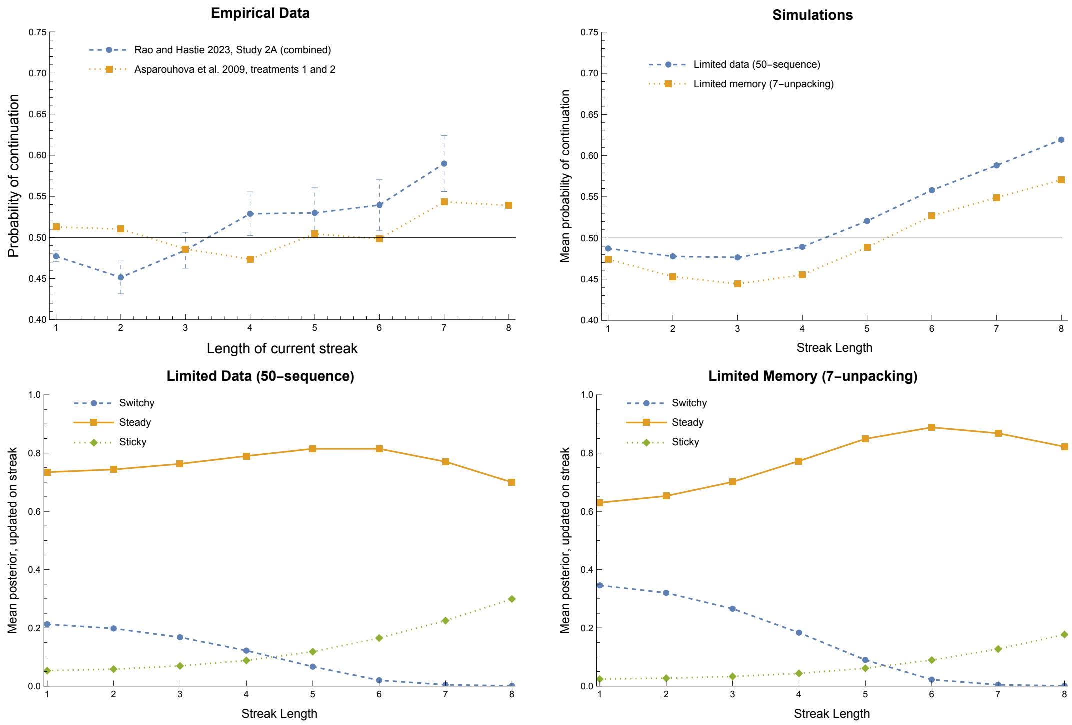

Figure 6: Nonlinear expectations, data and simulations. *Top left:* Combined empirical data from Rao and Hastie 2023, Study 2A (bars are 80%-confidence intervals) and from Asparouhova et al. 2009, treatments 1 and 2. *Top right:* Simulation results for mean probability of a continuation by streak length *Bottom:* Mean posteriors after updating on streaks, for limited data (left) and limited memory (right).

The simulations: To simulate this, first have our Bayesians update on limited data (or limited memory). Then present them with streaks of varying lengths, and see how likely they think the streaks are to continue. These streaks provides new evidence, since Sticky is more likely to generate streaks than Switchy or Steady. Thus when we update their beliefs on the streak, that shifts their posteriors over the causal hypotheses—which, in turn, shifts their probability estimates for what will happen next.

The results are displayed in the top right of Figure 6. With limited data (length-50 sequence), the mean probability for a continuation begins at 48.8%, falls slightly to 47.7% after a streak of 3, and then rises to 61.8% after a streak of 8. With limited memory (7-unpacking), the mean probability for a continuation begins at 47.5%, falls to 44.5% with a streak of 3, and then gradually rises to 57.1% after a streak of 8.

[14] Rao and Hastie 2023 find even clearer evidence for non-linearities when they elicit *binary predictions*, rather than probability estimates—see §5.5 below.

Why does this happen? Our agents start out (on average) favoring Switchy over Sticky, due to asymmetric convergence. At first, a growing streak means that the process is increasingly primed to *either* switch (if Switchy) or continue (if Sticky); since they're still more confident of Switchy, this first leads to *increasing* rates of gambler's reasoning. But as the streak grows, it eventually provides enough evidence for Sticky to swamp the initial asymmetry—they become more confident of Sticky than Switchy, and so begin to (increasingly) predict continuations. The effects of a streak on their beliefs about Switchy/Steady/Sticky are displayed on the bottom of Figure 6. As the streak grows, eventually they become more confident of Sticky than Switchy (the dotted line crosses the dashed line)—this is where they transition from gambler's to hot-hands reasoning.

Upshot: like real people, as our Bayesians see longer streaks they first increasingly do gambler's reasoning, and then increasingly do hot-hands reasoning.

5.2 Experience-Dependence

The findings: Related to nonlinearity, the degrees to which people exhibit gambler's and hot-hands reasoning—and how quickly they transition from the former to the latter as a streak grows—depends on their experience with the causal system. When people have more experience, they tend to deviate less from the true hit rate of 50% in their predictions (Chen et al. 2016), and they take longer to transition from gambler's to hot-hands reasoning (Burns and Corpus 2004; Asparouhova et al. 2009; Rao and Hastie 2023).

Once again, the most carefully-collected data is from Asparouhova et al. 2009 and Rao and Hastie 2023. They elicit people's probabilities of streaks of varying lengths continuing when faced with both (i) *familiar* random processes (a fair coin, and an urn drawn with replacement) and (ii) *unfamiliar* ones (a firm's earnings, an analyst's performance, and a company's stock). I've plotted Asparouhova et al. 2009's results in the top left of Figure 7 (again from their Table 6, this time separated by causal process), and Rao and Hastie 2023's results in the top right (again taken from their Study 2A raw data at https://osf.io/2d93t/, this time separated by causal process). Both find two trends: (1) with familiar processes (dotted), people deviate less from 50% in their probability estimates than they do with unfamiliar ones (dashed); and (2) familiar processes take much longer to shift from gambler's to hot-hands reasoning, if they do so at all.

The simulations: To model differing levels of familiarity (experience with) the causal system, we can give different amounts of information to different groups. With limited data: unfamiliar agents condition on a length-40 sequence, while familiar ones condition on a length-90 sequence. With limited memory: unfamiliar agents do 4-unpacking, while familiar agents do 20-unpacking. The results are displayed in the bottom of Figure 7; both replicate the qualitative trends that, facing familiar processes, agents (1) deviate less from 50%, and (2) take longer to transition from gambler's to hot-hands reasoning.

Why does this happen? Whether agents do gambler's or hot-hands reasoning depends on whether they're more confident of Switchy or Sticky, and the *degree* to which they do so is modulated by how confident they are of Steady. As we've seen above (Figures 4 and 5), as they're exposed to more data or unpack their memories further, they both (i) become more

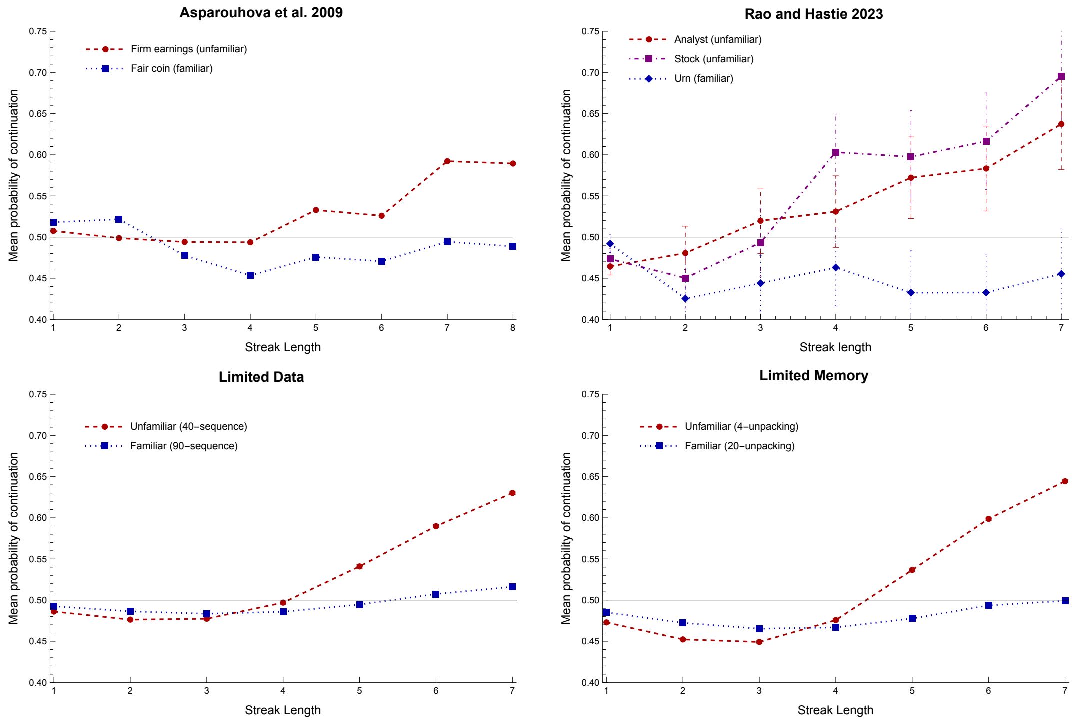

Figure 7: Experience-dependence, data and simulations. *Top:* Empirical data separated by causal process from Asparouhova et al. 2009 (left) and Rao and Hastie 2023 (right) Study 2A (bars are 80%-confidence intervals). *Bottom:* Simulation results for familiar vs. unfamiliar processes—the left shows limited data (given a length-40 vs. length-90 sequence), the right shows limited memory (given 4-unpacking vs. 20-unpacking).

confident in Steady, and (ii) increasingly favor Switchy over Sticky. The former explains trend (1): with familiar processes, they're more confident of Steady—preventing them from deviating as far from 50%. The latter explains trend (2): with familiar processes, they heavily favor Switchy over Sticky—meaning that it takes a longer streak before they start favoring Sticky and predicting continuations.

Upshot: like real people, our Bayesians' expectations take different shapes depending on how familiar they are with the causal system.

5.3 Sequence-dependence

The findings: An intuitive finding comes from Bloomfield and Hales 2002 and Bao et al. 2023, motivated by Barberis et al. 1998's model. Subjects are presented with sequences that vary in their number of switches, and asked for their probability that the current streak will continue. Their answers depend on how switchy the sequence is. Bloomfield and Hales 2002's results are plotted in Figure 8: for sequences with few switches, the mean continuation probability was 61.2%; for those with some, it was 50.9%; for those with many, it was 43.1%.

The simulations: To simulate this, first have our Bayesians update on limited data (or limited memory). Then present each agent with the sequences Bloomfield et al. used, have them update their opinions about the process in light of that sequence, and then elicit their probabilities for the sequence continuing. If the sequence has few switches, that provides evidence for Sticky; if the sequence has many switches, that provides evidence for Switchy. As a result, we expect that they'll be more likely to predict a continuation given the former and less likely given the latter.

They do. The results are plotted in Figure 8. The qualitative predictions are correct and robust, though of course the quantitative ones will depend on many details. Given limited data (length-30 sequence), the mean probability of continuing was 62.7% for sequences with few switches, 48.2% for those with a middling number, and 46.2% for those with many. Given limited memory (4-unpacking), the mean probability of continuing was 59.4% for sequences with few switches, 47.2% for those with a middling number, and 44.2% for those with many.

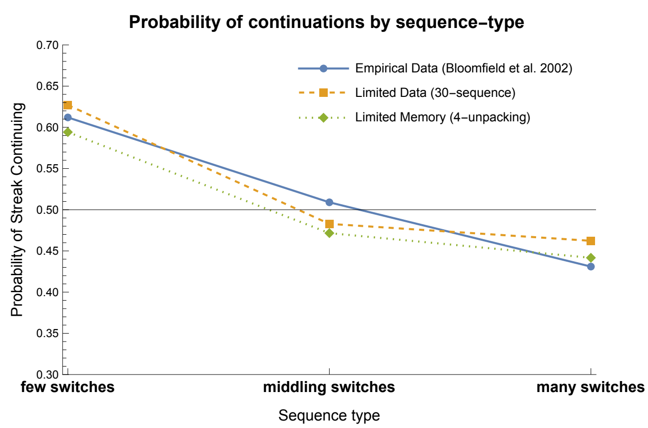

Figure 8: Sequence-Dependence, data and simulations. Empirical and simulated data of mean probability of sequence continuation, sequences grouped by how many switches they contained.

Upshot: like real people, our Bayesians' probability estimates about a streak continuing depend on the sequence they'd seen leading up to it.[15]

5.4 Too switchy

The findings: A classic finding is that when asked to produce sequences of random outcomes, people generate sequences that (i) switch too often—around 60% of the time, rather than 50%—and (ii) are increasingly likely to switch as streak-length grows. Since we've so far talked about the probability of a streak *continuing* (rather than switching), let's turn this around. Thus the findings are that people produce sequences that (i) *continue* less often

[15]One issue with this simple model is that once people see a switchy streak, that significantly changes their opinions about how switchy *all* future sequences will be. But in reality, people pay more attention to their local environment: if the parking battle has switched a lot for weeks, but then appears to shift to be acting Sticky, our Bayesians will update relatively quickly. As mentioned above, these dynamics could easily be modeled by incorporating Barberis et al. 1998's proposal that different instances of the process (or the same process across time) can vary in how Switchy or Sticky they are, using a hierarchical Bayesian model. Similar remarks apply to our other empirical predictions.

(around 40% of the time) than they should, and (ii) are *decreasingly* likely to continue as streak-length grows.

Though found by many authors in many settings (see fn. 3) the best quantitative data I know of is from Rapoport and Budescu 1997, via Rabin 2002. When asking subjects to produce random sequences from a fair coin, they find overall streak-continuation rates of 41.5%. After a streak of 1, the rate is 46%; after 2, it's 38%; and after 3, it's 29.8%. These rates are displayed in the top left of Figure 9.

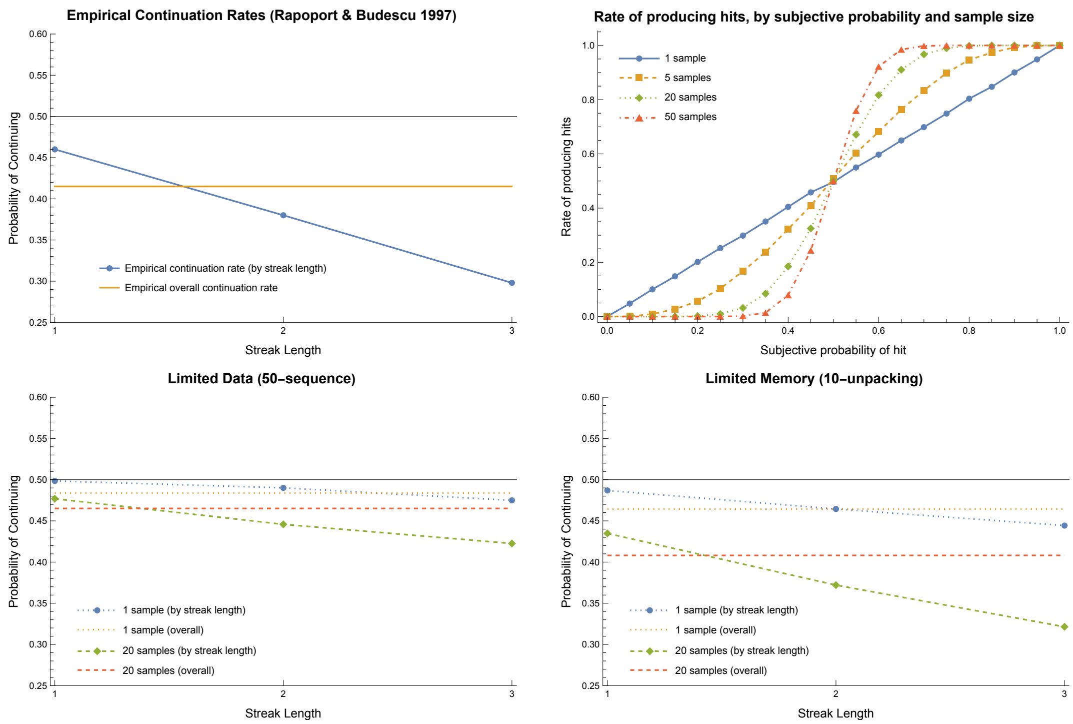

Figure 9: Too Switchy, data and simulations. *Top left:* empirical continuation rates in production tasks. *Top right:* rates of generating hits as a function of subjective probability when sampling from posteriors. *Bottom:* simulated continuation rates under limited data (left) and limited memory (right).

The simulations: To simulate this, first have our Bayesians update on limited data (or limited memory), and then use their posteriors to produce hypothetical sequences. But how to produce sequences? They should *not* always produce the outcome they think is most likely—for then every agent more confident of Switchy than Sticky would produce a perfectly-alternating sequence (0101010....), while every agent less confident of Switchy than Sticky would produce a perfectly-repeating sequence (1111111...). Instead, they should take *samples from their posteriors* to decide which sequence to generate. For example, suppose, given her uncertainty, Baylee is 60%-confident that the next outcome will be a hit. "Sampling from her posterior" means generating n outcomes each with an (independent) 60% chance to be a hit, and then writing down the most-common outcome from these samples. If she takes 1 sample, this is equivalent to *probability-matching:* writing down a hit with the probability (60%) that she assigns to their being one. If she takes many samples, she's increasingly likely to write down the outcome she thinks is most likely—see the top right of Figure 9. The idea that people use sampling to generate predictions is well-studied, theoretically motivated, and fits with a wide range of empirical data.[16]

The results for limited data (given a length-50 sequence) are summarized in the bottom left of Figure 9. 1-sampling (probability matching) generates an overall continuation rate of 48.4% (statistically-significantly below 50%). After streaks of 1, the rate is 49.9%; after 2, it's 48.8%; and after 3, it's 47.3%. 20-sampling generates an overall continuation rate of 46.5% (significantly below 50%). After streaks of 1, the rate is 47.2%; after 2, it's 43.9%, and after 3, it's 41.2%.

The results for limited memory (with 10-unpacking) are summarized in the bottom right of Figure 9. 1-sampling (probability matching) generates an overall continuation rate of 46.5% (significantly below 50%). After streaks of 1, the rate is 48.8%; after 2, it's 46.7%; after 3, it's 44.5%. 20-sampling generates an overall continuation rate of 40.6% (significantly below 50%). After streaks of 1, the rate is 43.7%; after 2, it's 37.4%; after 3, it's 32.5%. Note that these rates are similar to the empirical rates (compare bottom right to top left).

Limited data and limited memory both lead to the correct qualitative predictions; and given limited memory and many samples, the *quantitative* predictions are realistic. Notably, the data from Rao and Hastie 2023 provides empirical evidence that people are *not* simply probability-matching (drawing 1 sample), but instead tend to draw many samples. For they find (see §5.5) that when people's mean *probability estimates* for a continuation are around 45%, the *rate at which they predict* a continuation is around 20–30%—like the 20- and 50-sample lines (and unlike the 1- or 5-sample lines) in the top right of Figure 9.

Upshot: like real people, our limited Bayesians would produce "random" sequences that are too switchy, with the deviation from 50% increasing with streak-length.

5.5 Format-Dependence

The findings: Our final empirical finding is that when people are asked to *estimate a probability* of a continuation, their rates of gambler's reasoning look less stark than when asked to *make a binary prediction* (continue vs. switch). For processes that are known to have a long-run hit rate of 50%, the mean probability estimate is usually between 45–55%, while the rate of predicting a continuation can drop as low as 20%. The evidence for this comes most directly from Rao and Hastie 2023, though it also is supported by comparing the rates of gambler's reasoning in studies that use binary-prediction or production tasks (e.g. Bar-Hillel and Wagenaar 1991) to those that use probability-estimation tasks (e.g. Asparouhova et al. 2009; Dohmen et al. 2009; Miller and Sanjurjo 2018a).

Rao and Hastie 2023 find that across all trials the mean probability estimate for a continuation was just below 50%(at 49.2%; 95%-confidence interval = $[48.2, 50.0]$), but the rate of *predicting* a continuation was 40.0% (95%-confidence interval = $[38.7, 41.3]$). This difference can be seen if we break down the rates by streak-length, as done in the top left of Figure 10: the dashed line reproduces mean probability estimates (already shown in the top

[16]Gelman et al. 2013; Bonawitz et al. 2014b,a; Vul et al. 2014; Icard 2016; Sanborn and Chater 2016; Bhui and Gershman 2018; Zhu et al. 2020.

left of Figure 6), and the dotted line shows the rates of (binary) prediction. The deviations from 50% and the nonlinearities-in-streak-length are much more pronounced for the latter.

What explains this format-dependence? Rao and Hastie 2023 speculate that probability estimation may both trigger analytic thinking, and make salient to participants that they've been told that the hit rate is 50%. This likely plays a role, but there's another (complementary) explanation: our model predicts format-dependence *within* probability-estimate data. Precisely: if you take people's probability estimates and *dichotomize* them—count them as guessing "continuation" if their probability of a continuation is above 50% and "switch" if it's below 50% (dropping estimates of exactly 50%)—we expect their dichotomized prediction-rates to deviate farther from 50% than their mean probabilities. Why? The distribution of people's probabilities of a continuation will be a unimodal bell curve, so that small shifts in the mean probability above and below 50% can lead to dramatic swings in the *proportion* of people above or below 50%. This is illustrated in the top middle of Figure 10 using our limited memory (5-unpacked) agents. After a streak of 1, their mean probability is 47.4% but the proportion who predict "continuation" is 31.3%; after a streak of 2, their mean probability is 45.1% but the proportion who predict "continuation" is 25.1%; after a streak of 7, their mean probability is 60.5% but the proportion who predict "continuation" is 78.9%. In all cases, we predict that the deviation from 50% is larger when we look at the *rate of predictions* than when we look at the *mean probability*.

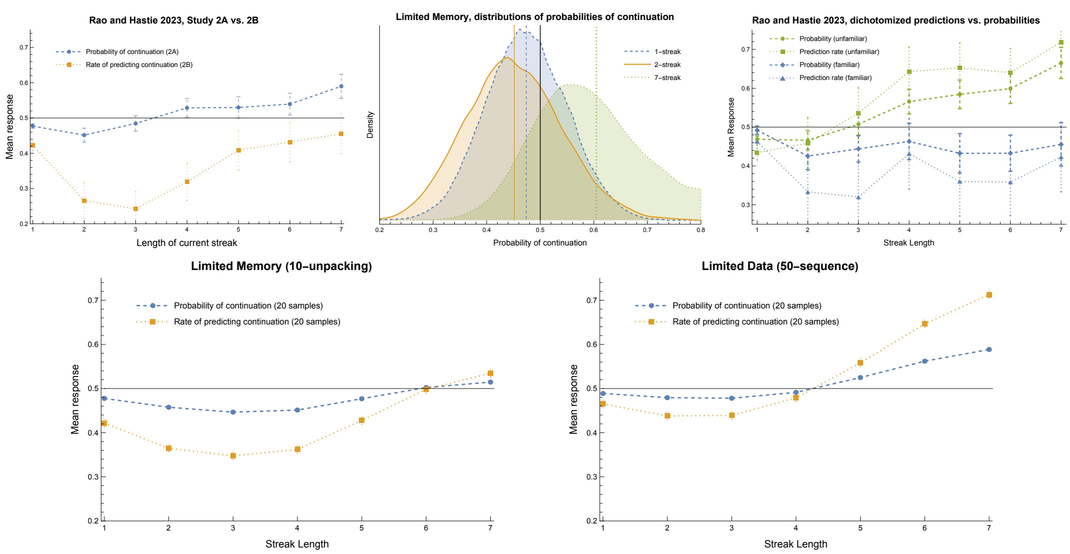

Figure 10: Format-dependence, data and simulations. *Top left:* Empirical contrast between probability-estimate and binary-prediction responses. *Top middle:* Histograms of our Bayesians' probabilities of a continuation facing streaks of various lengths. *Top right:* Comparison of mean probabilities and prediction rates in Rao and Hastie's study—note that the dotted lines (prediction rates) deviate further from 50% than the dashed lines (mean probabilities). *Bottom:* Comparison of probability estimates to binary-prediction rates amongst our Bayesians, given limited memory (left) and limited data (right).

This prediction is borne out in Rao and Hastie's data: even within their probability-estimate data (Study 2A), we find evidence for format-dependence. First, dichotomizing the data reveals that the rate at which people predict continuations deviates further from

50% than their mean probabilities. Overall, the mean probability estimate is slightly below 50% (49.2%), while the dichotomized rate of predicting continuations is lower (47.2%; $t(3726) = 1.82,\ p = 0.034$, one-sided). Though only a marginally-significant effect, this is likely due to the fact that the mean probability is close to 50%. This overall rate hides much heterogeneity: in some $\langle$process, streak-length$\rangle$-pairs, the mean probability was well below 50%; in others, it was well above 50%.

To probe this, we can separate the experimental conditions ("which causal process, and how long a streak?") into those with mean probability estimates *above* 50%, and those with means *below* 50%. Our prediction is that, in both cases, dichotomized rates of prediction will deviate further from 50% than mean probabilities. It does. In conditions with mean probabilities above 50%, the rate at which people predicted continuations (65.6%) was significantly above the mean probability of a continuation (59.4%; $t(708) = 2.36,\ p = 0.009$). And in conditions with mean probabilities below 50%, the rate at which people predicted a continuation (43.5%) was significantly below the mean probability of a continuation (47.2%; $t(3014) = 3.18,\ p < 0.001$, one-sided). This can be visualized, for instance, by dividing processes into *familiar* (drawing from an urn) and *unfamiliar* (analyst, stock price)—which, recall, have different rates of crossing the 50% line (§5.2)—and comparing the proportion of people predicting continuations to their mean probabilities at each $\langle$process, streak-length$\rangle$-pair. As the top right of Figure 10 shows, the proportions predicting continuations are uniformly further from 50% than the mean probabilities.

The simulations: The natural way to simulate binary predictions is to again use sampling (§5.4). Have our Bayesians sample outcomes from their posteriors, set their "realized" probability equal to the hit rate within their sample, and predict a continuation if the sampled rate of continuations is above 50% and a switch if it's below (randomizing if equals 50%). Since such samples generate an unbiased estimator of their true probabilities, it has no effect on their mean probability estimates. As above, if they only draw 1 sample then this is probability matching; the more samples they draw and the further their underlying probabilities are from 50%, the less likely they are to make a prediction opposing what their underlying probability says is most likely.

The results are plotted on the bottom of Figure 10. With limited data (50-sequence), overall mean probability estimate for continuations was 48.8%, while the mean rate of predicting continuations was 46.1% (95%-confidence interval = $[45.3, 47.1]$). How the two differ depends on how far the mean probability is from 50%; as shown in the bottom right plot of Figure 10, the prediction rate is uniformly farther from 50% than the mean probability. With limited memory (10-unpacking), the results are similar—though a bit better match to the empirical discrepancy between probability-estimates and binary predictions (bottom left). The overall mean probability estimate for continuations was 46.6%, while the overall rate of predicting continuations was 39.3%.

Nonlinearities emerge in prediction-rates because the closer people's probabilities are to 50%, the more likely (due to noisy sampling) they'll predict a result opposite to their probabilities. Thus the predicting-continuation rate first drops as the probability moves further below 50%, and then rises as that probability returns to and crosses the 50%-threshold.

Upshot: like real people, our Bayesians exhibit greater rates of gambler's reasoning when giving binary predictions than probability estimates.

5.6 New predictions

The causal-uncertainty hypothesis both explains the persistence gambler's and hot-hands reasoning, and also has empirical coverage that rivals the best extant models. I believe it is the only extant model that predicts all five of our qualitative trends: (1) people's expectations for how likely a streak is to continue are nonlinear in the streak length; (2) the *rate* and *degree* of this nonlinear change depends on how much experience they have with the process; (3) their predictions depend on how switchy the sequence they've just seen was; (4) they produce sequences that are too switchy; and (5) their responses exhibit greater deviations from 50% when asked for binary predictions than when asked for probability estimates.

Much remains to be done. First, quantitive model-fitting should be used to compare the causal-uncertainty hypothesis to its main rivals. Second, new predictions should formulated and tested. Let me highlight two predictions.

The most obvious prediction is that changing people's background beliefs about the causal mechanism should have direct effects on when and to what degree their predictions deviate from 50%. In particular, as people see more data, we should expect both (1) their probabilities will overall deviate less from 50%, and (2) they'll become slower to transition from gambler's to hot-hands reasoning. We could test this by introducing a new random process and exposing different subjects to different amounts of (Steady) data.

A second, more-subtle prediction is based on what drives nonlinear expectations. Our Bayesians are *learning* about the process by observing its outputs—long streaks provide evidence for Sticky. In contrast, those who are *producing* sequences (§5.4) are *not* learning. Because of this, we might initially expect that the causal-uncertainty hypothesis predicts that there'll be much lower continuation-rates for long streaks in production tasks than in prediction tasks—after all, in the former they've received no evidence from the streak, while in the latter they've gotten evidence for Sticky. This leads to the hypothesis that those *producing* (rather than predicting) a sequence won't exhibit a nonlinearity at all, and will instead be *monotonically* more likely to produce a switch as streak-length continues. Right?

Wrong. Surprisingly, when we simulate our Bayesians producing sequences, the average continuation-rate *does* exhibit nonlinearities. Why? Our Bayesians vary in how convinced of Switchy they are, and there's a *selection effect* when we consider longer streaks: agents who are convinced of Switchy will hardly ever generate long streaks, while those convinced of Sticky often will. So even though those who favor Switchy are a majority, they'll generate a *minority* of the long streaks. Since this selection effect grows with streak length, it'll look like a nonlinearity when averaging across streaks of a given length.

There are two ways to control for this selection effect. First, we could have *each* subject do both a production and a prediction task, and then compare the within-subject rates of continuations at various streak lengths. Prediction: at a given streak length, the mean person will be more likely to *predict* than to *produce* a continuation—for in the former case they've just gotten evidence for Sticky, while in the latter case they haven't. Second, we can dampen the evidential effects of seeing a long sequence. For example, we could tell subjects that we observed an extremely long sequences of outcomes and then selected a particular streak of length n; then we ask them for their probability that that sequence continued. This provides much less evidence for Sticky: in a long enough series of outcomes, Switchy, Steady, and Stick all agree that it's likely there'll be *some* long streaks—they just disagree about how common they'll be. Prediction: those presented with such a selected long streak will update less than those who see a long streak happen in real time, so the former will exhibit less nonlinearity and more gambler's reasoning for longer streaks.

As this illustrates, the fact that the causal-uncertainty hypothesis is a generative model means it's easy to investigate new predictions by simulation.[17] Thus it offers a potentially-fruitful way forward: if you think the causal-uncertainty hypothesis makes a (plausible or implausible) prediction, first confirm that it does by simulation, then test whether real people exhibit it. I'll be just as curious as you to hear the results.

6 Objections

I'll now address some objections by demonstrating the robustness of our predictions.

6.1 Unsure about the hit rate?

I've focused on the case where our Bayesians know that the long-run hit rate is 50%. But in realistic cases, they'll often know the hit rate is some other number, or be unsure what it is.

In general, the same dynamics will emerge—but we must be careful in how we construct the Markov chains so that their long-run (stationary) distributions have the desired hit rates. For example, if we want our hit rate to be 40%, then (1-step) Steady is of course the chain $\begin{pmatrix} 0.6 & 0.4 \\ 0.6 & 0.4 \end{pmatrix}$; but adding/subtracting the same number on both sides won't give us the correct Switchy/Sticky versions. For example, $\begin{pmatrix} 0.7 & 0.3 \\ 0.5 & 0.5 \end{pmatrix}$ has a long-run hit rate of 37.5%, not 40%. The "*Objection*: Non-50% long-run hit rates" subsection of the Mathematica notebook gives a function for generating the desired chains.

Two questions emerge. First, do we still observe asymmetric convergence of our Bayesians when they start out knowing that the hit rate is (say) 40%, rather than 50%? Yes—the "Asymmetric convergence for known non-50% stationary" section confirms as much.

Second, and more interestingly: when our agents start out *unsure* about the long-run hit rate, do we still observe asymmetric convergence? Yes—the "Asymmetric convergence when unsure of the long-run hit rate" subsection confirms as much. Suppose our Bayesians begin uniformly unsure whether the hit rate is 45%, 50%, or 55%. Conditional on each, they are $\frac{1}{3}$-confident in each of Switchy, Steady, and Sticky; thus they start out $\frac{1}{9}$-confident in each particular $\langle$hit rate, shiftiness$\rangle$ hypothesis, and $\frac{1}{3}$-confident in each overall type (Switchy, Steady, or Sticky). Data is generated from the 50%-Steady distribution. Figure 11 displays the results. The top row shows the results for limited data. Left is the mean posteriors in each hypothesis as a function of how long a sequence they observe. Solid lines are Steady

[17] Again, here's the Mathematica notebook: https://github.com/kevindorst/gamblers_fallacy_notebook.

hypotheses, dashed are Switchy, and dotted are Sticky; note that the dashed lines tend to be higher than the dotted ones. The overall results are easier to see if we marginalize over hit rates and look at the posteriors in Switchy/Steady/Sticky (top right). The bottom row shows the parallel results for limited memory; the main difference is that when agents track the overall proportion-true, they very quickly rule out alternative hit rates.

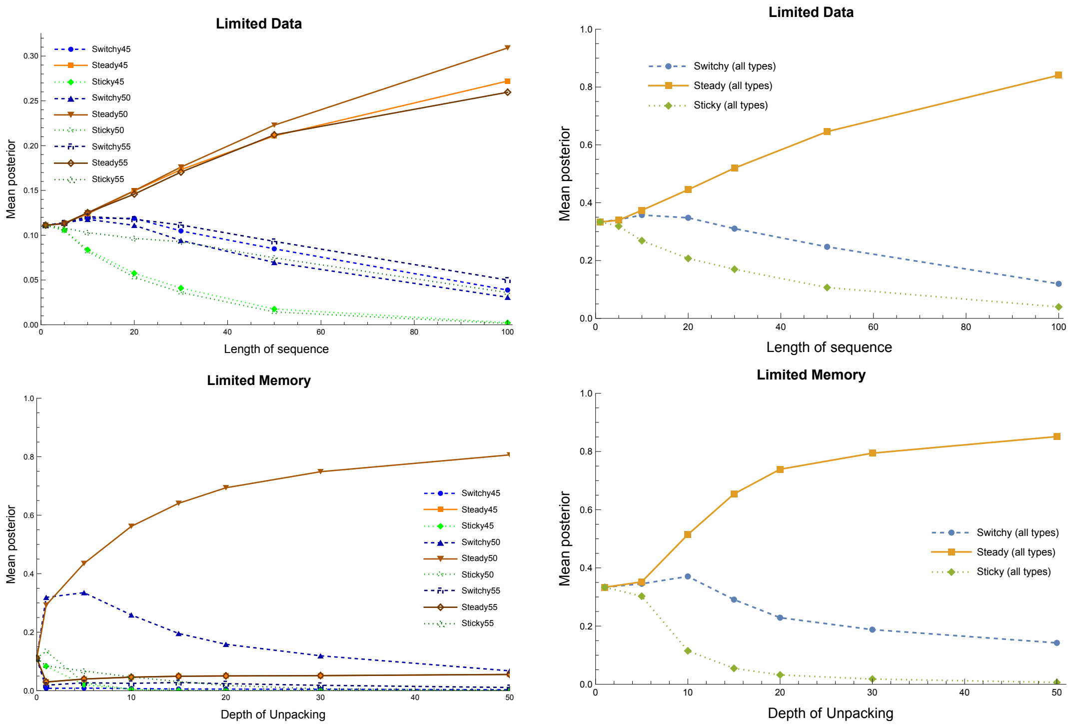

Figure 11: Asymmetric convergence with hit-rate uncertainty. *Top:* Mean posteriors in each hypothesis after seeing varying amounts of limited data (left); the trends are clearer if we marginalize over hit rates (right). *Bottom:* Mean posteriors in each hypothesis at various depths of memory-unpacking (left); the trends are clearer if we marginalize over hit rates (right).

This asymmetric convergence once again leads to the same empirical predictions—for instance, the Mathematica notebook finds similar degrees of nonlinear expectations.

6.2 Unsure about the shiftiness?

So far I've presupposed that the three live options were (1) *quite* Switchy (to degree 90%), (2) Steady, or (3) *quite* Sticky (to degree 90%). Might this have induced an unrealistic picture? Real people are are both unsure about whether the process is Switchy or Sticky, and unsure about *how* Switchy or Sticky it is. And if the sequence is in fact independent, then chains that are less shifty will be harder to rule out. So although we've established that Switchy-to-degree-90% is closer to Steady than Sticky-to-degree-90%, might it still be farther from Steady than Sticky-to-degree-*80%*? In other words: perhaps our agents will quickly learn that *if* it's Switchy or Sticky, it's only mildly so.

To test this, run simulations in which our agents know that the hit rate is 50%, but are initially (uniformly) unsure whether it's shiftiness (if Switchy or Sticky) builds to degree 90%, to 75%, or to 60%. The results are as expected: more-extreme rates get ruled out more quickly, but not *that* quickly. Even with 30-, 50- or 100-unpacking, the 75% and 90%-Switchy hypotheses have nontrivial probability. Indeed, as can be seen from the plot on the left of Figure 12, Switchy-to-degree-90% is closer to Steady than Sticky-to-degree-*75%*—note that the black dashed line is everywhere above the medium-green dotted line. As a result, the predictions are similar—for instance, we again see nonlinear expectations (right).

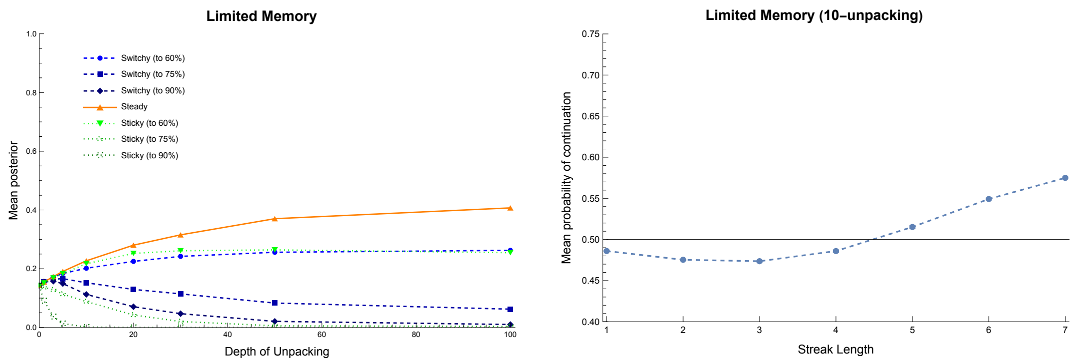

Figure 12: Asymmetric convergence with shiftiness uncertainty. Asymmetric convergence when agents are uncertain about how shifty the process is (left), leading to nonlinear expectations (right).

6.3 Unsure about the process's speed?

I've focused on processes that require a streak of 5 to build up their maximum shiftiness. But realistic processes will vary in how long of a streak is needed, so people will be uncertain about this as well. Perhaps *this* has induced an unrealistic picture?

It hasn't. If anything, the focus on 5-step chains seems to make the deviations from independence *less* extreme than they would otherwise be. To test this, simulate agents who start out uniformly unsure whether—supposing the process is Switchy or Sticky to 90%—it builds up over a streak of 3 (short), 5 (medium), or 7 (long). The convergence results for limited memory are displayed in the left of Figure 13—note that every version of Switchy is always more plausible than any version of Sticky. As a result, this leads to even more extreme rates of gambler's reasoning (Figure 13, right).

6.4 Different types of Switchy/Sticky chains?

I've focused on *streak chains*, wherein the probability of the next toss only depends on the length of the current streak—as soon as the streak is broken, the probabilities revert back to near 50%. As mentioned above (§3), this has been for simplicity. Many realistic processes' probabilities depend on the *full (recent) sequence* of outcomes. For example, when drawing cards from a deck, a series of reds will raise the probability of a black, but then a series of blacks will only *slowly* move the probabilities back towards 50%. This requires a *counting chain*: one that tracks the counts of each type of outcome in the last n outcomes.

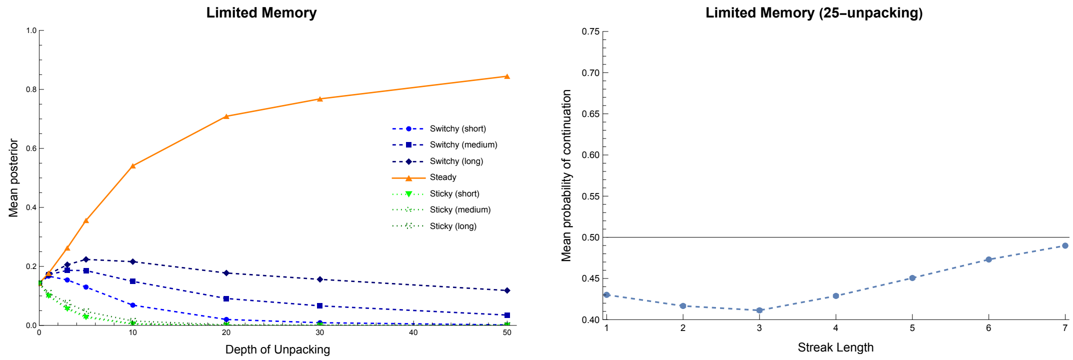

Figure 13: Asymmetric convergence with speed uncertainty. Asymmetric convergence when agents are uncertain about the speed with which the process's probabilities deviate from the long-run hit rate (left), leading to extreme nonlinear expectations even with large depths of memory-unpacking (right).

The modeling difficulty is that the number of states in such chains will exponentially explode in the length of recent sequence we track. This makes such chains tedious to code and difficult to simulate, but there's not much new conceptually that emerges. The "*Objection:* Counting chains" section of the Mathematica notebook shows how we can replicate our results with counting chains—the qualitative dynamics are similar to streak chains. In fact, counting chains may be able to help account for further empirical trends—such as the fact that people's probability estimates depend on more than just the current streak, but also on other recent outcomes (Rabin and Vayanos 2010, Table 1).

7 Conclusion

So Bayesians—if they are data- or memory-limited—commit the gambler's and hot-hands fallacies. Indeed, they do so in ways that are remarkably similar to the way that *we* do so. What to make of this?

First, and most obviously, we can strike the gambler's and hot-hands fallacies from the list of clear pieces of evidence for human irrationality. If anything—since we always knew that people were data- and memory-limited—the high-level empirical trends seem to be exactly what we should expect if people are being *rational*. This paper is an instance of a general trend: what looks like irrationality is often perfectly sensible under known resource limitations (Brav and Heaton 2002; Lieder and Griffiths 2019; Icard 2023). Data and memory are prime examples of such limitations (Wilson 2014; Dasgupta and Gershman 2021). I hope the models and methods I've used here will suggest similar approaches to other apparent irrationalities in human reasoning. Moreover, the fact that the causal-uncertainty hypothesis comes with generative models suggests a potentially-fruitful approach to generating and testing new predictions. There is much still to be done in terms of both predictions and quantitative model-fitting and comparisons.

Second, I have critiqued the "mechanistic" approach to the mind at one of its strongest points. Even when it comes to apparent basic violations of statistical reasoning, a rational account is well-suited to explain the data. Moreover, this is the first *bounded*-optimal model of the gambler's fallacy with such empirical coverage that I know of—extant empirically-adequate models are *error*-optimal, meaning that they rely on a hard-coded (and usually unexplained) error at their core. Since my model does not do so, it gives a vindication of gambler's reasoning: it is a model we can have in mind as we feel the pull of gambler's reasoning, without undermining that pull nor inducing incoherence.

Finally: as this discussion illustrates, questions about human rationality cross-cut traditional disciplinary boundaries and methods. As I hope to have shown, we can make progress on them by being ecumenical, pulling on work from philosophy and economics as well as psychology. As we try to sort out how the mind works, conceptual, mathematical, experimental, and computational approaches all have their place.[18]

References

Alon, Noga, Matias, Yossi, and Szegedy, Mario, 1996. 'The space complexity of approximating the frequency moments'. *Proceedings of the Annual ACM Symposium on Theory of Computing*, Part F129452:20–29.

Altmann, Erik M. and Burns, Bruce D., 2005. 'Streak biases in decision making: Data and a memory model'. *Cognitive Systems Research*, 6(1):5–16.

Anderson, John R, 1990. *The Adaptive Character of Thought.* Erlbaum Associates.

Ariely, Dan, 2008. *Predictably Irrational.* Harper Audio.

Aronowitz, Sara, 2019. 'Memory is a modeling system'. *Mind and Language*, 34(4):483–502.

Asparouhova, Elena, Hertzel, Michael, and Lemmon, Michael, 2009. 'Inference from Streaks in random outcomes: Experimental evidence on beliefs in regime shifting and the law of small numbers'. *Management Science*, 55(11):1766–1782.

Bao, Te, Corgnet, Brice, Hanaki, Nobuyuki, Riyanto, Yohanes E, and Zhu, Jiahua, 2023. 'Predicting the unpredictable: New experimental evidence on forecasting random walks'. *Journal of Economic Dynamics and Control*, 146:104571.

Bar-Hillel, Maya and Wagenaar, Willem A., 1991. 'The perception of randomness'. *Advances in Applied Mathematics*, 12(4):428–454.

Barberis, Nicholas, Shleifer, Andrei, and Vishny, Robert, 1998. 'A model of investor sentiment'. *Journal of Financial Economics*, 49(3):307–343.

Bartoš, František, Sarafoglou, Alexandra, Godmann, Henrik R, Sahrani, Amir, Leunk, David Klein, Gui, Pierre Y, Voss, David, Ullah, Kaleem, Zoubek, Malte J, Nippold, Franziska, Aust, Frederik, Vieira, Felipe F, Islam, Chris-gabriel, Zoubek, Anton J, Shabani, Sara, Petter, Jonas, Roos, Ingeborg B, Finnemann, Adam, Lob, Aaron B, Hoffstadt, Madlen F, Nak, Jason, Ron, Jill De, Derks, Koen, Huth, Karoline, Terpstra, Sjoerd, Bastelica, Thomas, Matetovici, Magda, Ott, Vincent L, Zetea, Andreea S, Karnbach, Katharina, Donzallaz, Michelle C, John, Arne, Moore, Roy M, Assion, Franziska, Bork, Riet Van, Leidinger, Theresa E, Zhao, Xiaochang, Motaghi, Adrian Karami, Pang, Ting, Armstrong, Hannah, Peng, Tianqi, Bialas, Mara, Pang, Joyce Y, Fu, Bohan, Yang, Shujun, Lin, Xiaoyi, Sleiffer, Dana, Bognar, Miklos, Aczel, Balazs, and Wagenmakers, Eric-Jan, 2023. 'Fair Coins Tend to Land on the Same Side They Started: Evidence from 350,757 Flips'. 1–12.

Bhui, Rahul and Gershman, Samuel J., 2018. 'Decision by sampling implements efficient coding of psychoeconomic functions'. *Psychological Review*, 125(6):985–1001.

Bhui, Rahul, Lai, Lucy, and Gershman, Samuel J., 2021. 'Resource-rational decision making'. *Current Opinion in Behavioral Sciences*, 41:15–21.

Bloomfield, Robert and Hales, Jeffrey, 2002. 'Predicting the next step of a random walk: Experimental evidence of regime-shifting beliefs'. *Journal of Financial Economics*, 65(3):397–414.

Bogartz, Richard S, 1965. 'Sequential dependencies in children's probability learning.' *Journal of Experimental Psychology*, 70(4):365.

Bonawitz, Elizabeth, Denison, Stephanie, Gopnik, Alison, and Griffiths, Thomas L., 2014a. 'Win-Stay, Lose-Sample: A simple sequential algorithm for approximating Bayesian inference'. *Cognitive Psychology*, 74:35–65.

Bonawitz, Elizabeth, Denison, Stephanie, Griffiths, Thomas L., and Gopnik, Alison, 2014b. 'Probabilistic models, learning algorithms, and response variability: Sampling in cognitive development'. *Trends in Cognitive Sciences*, 18(10):497–500.

Braga, João N, Ferreira, Mário B, Sherman, Steven J, Mata, André, Jacinto, Sofia, and Ferreira, Marina, 2018. 'What's next? Disentangling availability from representativeness using binary decision tasks'. *Journal of Experimental Social Psychology*, 76:307–319.

[18]Thanks to Dmitri Gallow, Peter Gerdes, Sam Gershman, Caspar Hare, Thomas Icard, Joshua Knobe, Matt Mandelkern, Jack Spencer, Elliot Thornley, Kevin Zollman, and especially Zach Barnett and Yonathan Fiat for much help figuring out what was (and wasn't) interesting about this idea. Thanks also to audiences at MIT, the Pittsburgh Formal Epistemology Workshop, the University of Birmingham, NUS, FAX3, the Gershman Lab, and many anonymous blog commenters.

Brav, Alon and Heaton, J. B., 2002. 'Competing Theories of Financial Anomalies'. *Review of Financial Studies*, 15(2 SPEC.):575–606.
Budescu, David V and Rapoport, Amnon, 1994. 'Subjective randomization in one-and two-person games'. *Journal of Behavioral Decision Making*, 7(4):261–278.
Burns, Bruce D and Corpus, Bryan, 2004. 'Randomness and Inudctions from Streaks: "Gambler's Fallacy" Versus "Hot Hand"'. *Psychological Bulletin & Review*, 11(1):179–184.
Chater, Nick and Oaksford, Mike, 1999. 'Ten years of the rational analysis of cognition'. *Trends in Cognitive Sciences*, 3(2):57–65.
Chen, Daniel L., Moskowitz, Tobias J., and Shue, Kelly, 2016. 'Decision making under the Gambler's fallacy: Evidence from asylum judges, loan officers, and baseball umpires'. *Quarterly Journal of Economics*, 131(3):1181–1242.
Chiesi, Francesca and Primi, Caterina, 2009. 'Recency effects in primary-age children and college students'. *International Electronic Journal of Mathematics Education*, 4(3):259–279.
Cohen, L. Jonathan, 1981. 'Can human irrationality be experimentally demonstrated?' *Behavioral and Brain Sciences*, 4(3):317–331.
Cosmides, Leda and Tooby, John, 1996. 'Are humans good intuitive statisticians after all? Rethinking some conclusions from the literature on judgment under uncertainty'. *cognition*, 58(1):1–73.
Cover, Thomas M. and Thomas, Joy A., 2005. *Elements of Information Theory*. John Wiley & Sons.
Craig, Grace J and Myers, Jerome L, 1963. 'A developmental study of sequential two-choice decision making'. *Child Development*, 483–493.
Crupi, Vincenzo, Fitelson, Branden, and Tentori, Katya, 2008. 'Probability, confirmation, and the conjunction fallacy'. *Thinking & Reasoning*, 14(2):182–199.
Dasgupta, Ishita and Gershman, Samuel J., 2021. 'Memory as a Computational Resource'. *Trends in Cognitive Sciences*, 25(3):240–251.
Derks, Peter L, 1963. 'Effect of run length on the "gambler's fallacy".'
Derks, Peter L and Paclisanu, Marianne I, 1967. 'Simple strategies in binary prediction by children and adults.' *Journal of Experimental Psychology*, 73(2):278.
Diaconis, Persi, Holmes, Susan, and Montgomery, Richard, 2007. 'Dynamical bias in the coin toss'. *SIAM Review*, 49(2):211–235.
Diener, Don and Thompson, W Burt, 1985. 'Recognizing randomness'. *The American journal of psychology*, 433–447.
Dohmen, Thomas, Falk, Armin, Huffman, David, Marklein, Felix, and Sunde, Uwe, 2009. 'Biased probability judgment: Evidence of incidence and relationship to economic outcomes from a representative sample'. *Journal of Economic Behavior and Organization*, 72(3):903–915.
Doody, Ryan, 2020. 'The Sunk Cost "Fallacy" Is Not a Fallacy'. *Ergo*, 6(40):1153–1190.
Dorst, Kevin, 2023a. 'Being Rational and Being Wrong'. *The Philosophers' Imprint*, To appear.
———, 2023b. 'Rational Polarization'. *The Philosophical Review*, 132(3):355–458.
Dorst, Kevin and Mandelkern, Matthew, 2022. 'Good Guesses'. *Philosophy and Phenomenological Research*, 105(3):581–618.
Eldar, Eran, Rutledge, Robb B., Dolan, Raymond J., and Niv, Yael, 2016. 'Mood as Representation of Momentum'. *Trends in Cognitive Sciences*, 20(1):15–24.
Estes, William Kaye, 1962. 'Learning theory'. *Annual review of psychology*, 13(1):107–144.
Fine, Cordelia, 2005. *A Mind of Its Own: How Your Brain Distorts and Deceives*. W. W. Norton & Company.
Fischhoff, Baruch and Beyth, Ruth, 1975. '"I knew it would happen": Remembered probabilities of once-future things.'
Fisher, Sarah A., 2020. 'Rationalising framing effects: At least one task for empirically informed philosophy'. *Critica-Revista Hispanoamericana de Filosofia*, 52(156):5–30.
Fisher, Sarah A. and Mandel, David R., 2021. 'Risky-choice framing and rational decision-making'. *Philosophy Compass*, 16(8):1–11.
Fletcher, Anne, 2018. *From the Mill to Monte Carlo: The Working-class Englishman who Beat the Monaco Casino and Changed Gambling Forever*. Amberley Publishing Limited.
Geisler, Wilson S, 2011. 'Contributions of ideal observer theory to vision research'. *Vision research*, 51(7):771–781.
Gelman, Andrew, Carlin, John B, Stern, Hal S, Dunson, David B, Vehtari, Aki, and Rubin, Donald B, 2013. *Bayesian data analysis*. CRC press.
Gershman, Samuel, 2021. *What Makes Us Smart: The Computational Logic of Human Cognition*. Princeton University Press.
Gigerenzer, Gerd and Hoffrage, Ulrich, 1995. 'How to improve Bayesian reasoning without instruction: Frequency formats.' *Psychological review*, 102(4):684.
Gilovich, Thomas, Vallone, Robert, and Tversky, Amos, 1985. 'The hot hand in basketball: On the misperception of random sequences'. *Cognitive psychology*, 17(3):295–314.
Gopnik, Alison, 1996. 'The Scientist as Child'. *Philosophy of Science*, 63(December):485–514.
———, 2020. 'Childhood as a solution to explore–exploit tensions'. *Philosophical Transactions of the Royal Society B: Biological Sciences*, 375(1803).
Griffin, Dale and Tversky, Amos, 1992. 'The Weighing of Evidence and the Determinants of Confidence'. *Cognitive Psychology*, 24:411–435.
Griffiths, Thomas L., Kemp, Charles, and Tenenbaum, Joshua B., 2008. 'Bayesian Models of Cognition'. In R. Sun, ed., *The Cambridge Handbook of Computational Psychology*, 59–100. Cambridge University Press.
Hahn, Ulrike and Harris, Adam J.L., 2014. 'What Does It Mean to be Biased.: Motivated Reasoning and Rationality.' In *Psychology of Learning and Motivation - Advances in Research and Theory*, volume 61, 41–102. Academic Press.
Hahn, Ulrike and Warren, Paul A., 2009. 'Perceptions of Randomness: Why Three Heads Are Better Than Four'.

Psychological Review, 116(2):454–461.

Hamblin, Charles L, 1976. ‘Questions in montague english’. In *Montague grammar*, 247–259. Elsevier.

Hartley, Catherine A., 2022. ‘How do natural environments shape adaptive cognition across the lifespan?’ *Trends in Cognitive Sciences*, 26(12):1029–1030.

Hedden, Brian, 2019. ‘Hindsight Bias is Not a Bias’. *Analysis*, 79(1):43–52.

Henderson, Leah and Gebharter, Alexander, 2021. ‘The role of source reliability in belief polarisation’. *Synthese*, 199(3-4):10253–10276.

Icard, Thomas, 2016. ‘Subjective Probability as Sampling Propensity’. *Review of Philosophy and Psychology*, 7(4):863–903.

———, 2017. ‘Bayes, Bounds, and Rational Analysis’. *Philosophy of Science*, 694837.

———, 2021. ‘Why Be Random?’ *Mind*, 130(517):111–139.

Icard, Thomas F, 2023. *Resource Rationality.*

Jiao, Peiran, 2017. ‘Belief in Mean Reversion and the Disposition Effect: An Experimental Test’. *Journal of Behavioral Finance*, 18(1):29–44.

Kahneman, Daniel, 2011. *Thinking Fast and Slow.* Farrar, Straus, and Giroux.

Kahneman, Daniel, Slovic, Paul, and Tversky, Amos, eds., 1982. *Judgment Under Uncertainty: Heuristics and Biases.* Cambridge University Press.

Kahneman, Daniel and Tversky, Amos, 1972. ‘Subjective probability: A judgment of representativeness’. *Cognitive psychology*, 3(3):430–454.

———, 1973. ‘On the psychology of prediction.’ *Psychological review*, 80(4):237.

———, 1996. ‘On the reality of cognitive illusions.’ *Psychological Review*, 103(3):582–591.

Karlan, Brett, 2021. ‘Reasoning with heuristics’. *Ratio*, 34(2):100–108.

Kelly, Thomas, 2004. ‘Sunk costs, rationality, and acting for the sake of the past’. *Nous*, 38(1):60–85.

———, 2008. ‘Disagreement, Dogmatism, and Belief Polarization’. *The Journal of Philosophy*, 105(10):611–633.

Kinney, David and Bright, Liam Kofi, 2021. ‘Risk aversion and elite-group ignorance’. *Philosophy and Phenomenological Research*, 1–23.

Koehler, Jonathan J., 1996. ‘The base rate fallacy reconsidered: Descriptive, normative, and methodological challenges’. *Behavioral and Brain Sciences*, 19(1):1–53.

Koralus, Philipp and Mascarenhas, Salvador, 2013. ‘The erotetic theory of reasoning: bridges between formal semantics and the psychology of deductive inference’. *Philosophical Perspectives*, 27:312–365.

Lejarraga, Tomás and Hertwig, Ralph, 2021. ‘How Experimental Methods Shaped Views on Human Competence and Rationality’. *Psychological Bulletin*, 147(6):535–564.

Lewis, David, 1980. ‘A subjectivist’s guide to objective chance’. In Richard C Jeffrey, ed., *Studies in Inductive Logic and Probability*, volume 2, 263–293. University of California Press.

Lewis, Michael, 2016. *The Undoing Project: A Friendship that Changed the World.* Penguin UK.

Lichtenstein, Sarah, Fischhoff, Baruch, and Phillips, Lawrence D., 1982. ‘Calibration of probabilities: The state of the art to 1980’. In Daniel Kahneman, Paul Slovic, and Amos Tversky, eds., *Judgment under Uncertainty*, 306–334. Cambridge University Press.

Lieder, Falk and Griffiths, Thomas L., 2019. ‘Resource-rational analysis: Understanding human cognition as the optimal use of limited computational resources’. *Behavioral and Brain Sciences.*

Mandelbaum, Eric, 2018. ‘Troubles with Bayesianism: An introduction to the psychological immune system’. *Mind & Language*, 1–17.

Marr, David, 1982. *Vision.* MIT Press.

Militana, Erika, Wolfson, Elissa, and Cleaveland, J Mark, 2010. ‘An effect of inter-trial duration on the gambler’s fallacy choice bias’. *Behavioural processes*, 84(1):455–459.

Miller, Joshua Benjamin and Sanjurjo, Adam, 2018a. ‘How Experience Confirms the Gambler’s Fallacy when Sample Size is Neglected’. *OSFPreprints*, 1–12.

Miller, Joushua B. and Sanjurjo, Adam, 2018b. ‘Surprised by the Hot Hands Fallacy? A Truth in the Law of Small Numbers’. *Econometrica*, 86(6):2019–2047.

Moore, Don A and Healy, Paul J, 2008. ‘The Trouble with Overconfidence.’ *Psychological Review*, 115(2):502.

Nebel, Jacob M., 2015. ‘Status quo bias, rationality, and conservatism about value’. *Ethics*, 125(2):449–476.

O’Connor, Cailin and Weatherall, James Owen, 2018. ‘Scientific Polarization’. *European Journal for Philosophy of Science*, 8(3):855–875.

———, 2019. *The Misinformation Age: How False Beliefs Spread Cailin O’Connor James Owen Weatherall.* Yale Press.

Oskarsson, An T., Van Boven, Leaf, McClelland, Gary H., and Hastie, Reid, 2009. ‘What’s Next? Judging Sequences of Binary Events’. *Psychological Bulletin*, 135(2):262–285.

Peterson, CAMERON R. and Beach, LEE R., 1967. ‘Man As an Intuitive Statistician’. *Psychological Bulletin*, 68(1):29–46.

Pinker, Steven, 1997. *How the Mind Works.* W. W. Norton & Company.

Rabin, Matthew, 2002. ‘Inference By Believers in the Law of Small Numbers’. *The Quarterly Journal of Economics*, 117(3):775–816.

Rabin, Matthew and Vayanos, Dimitri, 2010. ‘The gambler’s and hot-hand fallacies: Theory and applications’. *Review of Economic Studies*, 77(2):730–778.

Rao, Kariyushi and Hastie, Reid, 2023. ‘Predicting Outcomes in a Sequence of Binary Events: Belief Updating and Gambler’s Fallacy Reasoning’. *Cognitive Science*, 47(1).

Rapoport, Amnon and Budescu, David V, 1992. ‘Generation of random series in two-person strictly competitive games.’ *Journal of Experimental Psychology: General*, 121(3):352.

Rapoport, Amnon and Budescu, David V., 1997. ‘Randomization in individual choice behavior.’ *Psychological Review*, 104(3):603–617.

Richtel, Matt, 2006. ‘From the back office, a casino can change the slot machine in seconds’.

Ruggeri, Azzurra and Lombrozo, Tania, 2015. 'Children adapt their questions to achieve efficient search'. *Cognition*, 143:203–216.

Sanborn, Adam N. and Chater, Nick, 2016. 'Bayesian Brains without Probabilities'. *Trends in Cognitive Sciences*, 20(12):883–893.

Schwarz, Norbert, Strack, Fritz, Hilton, Denis, and Naderer, Gabi, 1991. 'Base rates, representativeness, and the logic of conversation: The contextual relevance of "irrelevant" information'. *Social Cognition*, 9(1):67–84.

Sellars, Wilfrid, 1956. 'Empiricism and the Philosophy of Mind'. *Minnesota studies in the philosophy of science*, 1(19):253–329.

Suetens, Sigrid, Galbo-Jørgensen, Claus B., and Tyran, Jean Robert, 2016. 'Predicting Lotto Numbers: a Natural Experiment on the Gambler'S Fallacy and the Hot-Hand Fallacy'. *Journal of the European Economic Association*, 14(3):584–607.

Tenenbaum, Joshua B and Griffiths, Thomas L, 2006. 'Optimal Predictions in Everyday Cognition'. *Psychological Science*, 17(9):767–773.

Tenenbaum, Joshua B, Kemp, Charles, Griffiths, Thomas L, and Goodman, Noah D, 2011. 'How to Grow a Mind: Statistics, structure, and Abstraction'. *Science*, 331(6022):1279–1285.

Thaler, Richard H., 2015. *Misbehaving: The Making of Behavioural Economics*. Penguin.

Thorstad, David, 2021. 'The accuracy-coherence tradeoff in cognition'. *British Journal for the Philosophy of Science*, To appear.

———, 2023. 'Why bounded rationality (in epistemology)?' *Philosophy and Phenomenological Research*.

Todorov, Emanuel, 2004. 'Optimality principles in sensorimotor control'. *Nature neuroscience*, 7(9):907–915.

Tversky, Amos and Gilovich, Thomas, 1989. 'The "Hot Hand": Statistical Reality or Cognitive Illusion?' *CHANCE*, 2(4):31–34.

Tversky, Amos and Kahneman, Daniel, 1971. 'Belief in the law of small numbers'. *Psychological Bulletin*, 76(2):105–110.

———, 1974. 'Judgment under uncertainty: Heuristics and biases'. *Science*, 185(4157):1124–1131.

———, 1983. 'Extensional Versus Intuitive Reasoning: The Conjunction Fallacy in Probability Judgment.' *Psychological Review*, 90(4):293.

Tyszka, Tadeusz, Markiewicz, Łukasz, Kubińska, Elżbieta, Gawryluk, Katarzyna, and Zielonka, Piotr, 2017. 'A belief in trend reversal requires access to cognitive resources'. *Journal of Cognitive Psychology*, 29(2):202–216.

Vul, Edward, Goodman, Noah, Griffiths, Thomas L., and Tenenbaum, Joshua B., 2014. 'One and done? Optimal decisions from very few samples'. *Cognitive Science*, 38(4):599–637.

Weatherall, James Owen and O'Connor, Cailin, 2020. 'Endogenous epistemic factionalization'. *Synthese*, 1–23.

Williams, Daniel, 2021. 'Socially adaptive belief'. *Mind and Language*, 36(3):333–354.

Wilson, Andrea, 2014. 'Bounded Memory and Biases in Information Processing'. *Econometrica*, 82(6):2257–2294.

Zhu, Jian-Qiao, Sanborn, Adam N, and Chater, Nick, 2020. 'The Bayesian sampler: Generic Bayesian inference causes incoherence in human probability judgments.' *Psychological Review*.