Tempering inference to the best explanation

Abstract

What does it mean to apply inference to the best explanation (IBE) in probabilistic inference? Incompatibilist accounts treat IBE as an alternative to Bayesian conditionalization. The leading incompatibilist inference rule is the Bump Rule, a rule initially developed by critics of IBE. I develop a series of challenges for the Bump Rule, then propose an alternative: tempered updating. I show how tempered updating avoids the challenges facing the Bump Rule while inheriting many of its desired behaviors. I conclude by reflecting on the prospects for incompatibilist approaches to IBE.

1 Introduction

What does it mean to apply inference to the best explanation (IBE) in probabilistic inference? Two types of answers have been proposed.

Compatibilists (Lange 2022; Lipton 2004; Okasha 2000; Poston 2014) think that IBE is compatible with Bayesian conditionalization. Bayesian updating is driven by likelihoods which capture the degree to which hypotheses explain observed data. Updating is also driven by prior probabilities of hypotheses, which can be influenced by explanatory virtues such as simplicity (Lombrozo 2007; Sober 2015) and unification (Kitcher 1981; Lange 2014). Explanatory considerations may further serve as heuristic guides to Bayesian inference (McGrew 2003; Dellsén 2018) and hypothesis selection (Psillos 2000).

Incompatibilists (Douven 1999, 2013; Douven and Wenmackers 2017; Douven 2020, 2022) think that IBE is a distinct inference rule from Bayesian conditionalization. IBE is often treated as a distinct inference rule governing coarse-grained attitudes such as belief and acceptance (Dellsén forthcoming; Harman 1986; Musgrave 1988). Incompatibilists seek to preserve this same distinctness in a probabilistic setting.

By contrast with the wide variety of coarse-grained versions of IBE on offer, the space of incompatibilist probabilistic updating rules is relatively under-explored. By far the most-discussed incompatibilist updating rule is the Bump Rule (Douven 2013, 2020, 2022; Trpin and Pellert 2019), proposed by Bas van Fraassen (1989) in the process of criticizing IBE. I characterize the Bump Rule in Sections 2-3.

It is often not a good strategy to allow oneself to be defined by one's opponents. In Section 4, I argue that the Bump Rule faces several challenges. To meet these challenges, I propose and characterize an alternative explanationist update rule: tempered updating (Section 5). I argue that tempered updating meets the challenges facing the Bump Rule (Section 6).

Much of the support for the Bump Rule comes from a series of probabilistic simulations in which the Bump Rule strikes an attractive balance between learning speed and accuracy.[1] I extend these simulations to show how tempered updating captures a similarly attractive balance between learning speed (Section 7) and accuracy (Section 8). As a result, many of those attracted to the Bump Rule based on its performance in probabilistic simulations may also be attracted to tempered updating (Section 9).

Section 10 concludes by asking how these results bear on the case for incompatibilist explanationism. On the one hand, they provide some support for incompatibilist explanationism by developing a novel incompatibilist rule which avoids some challenges facing the Bump Rule and shows potentially attractive performance in simulations. On the other hand, tempered updating does not address all traditional concerns for incompatibilist explanationism, and our discussion will reveal some under-appreciated concerns facing both the Bump Rule and tempered updating. In this way, while tempered updating may improve the prospects for incompatibilist explanationism, it does not fully settle the debate between compatibilists and incompatibilists.

[1] See especially (Douven 2013; Douven and Wenmackers 2017; Douven 2020, 2022). Following Pettigrew (2021), I do not discuss social learning simulations (Douven and Wenmackers 2017) in this paper as Bayesians may not endorse the social learning rules used in them.

2 Bayes and the bump: The basic idea

Suppose you are concerned with a range $\mathcal{H} = \{H_i\}$ of hypotheses. You have prior credences $p(H_i)$ in the hypotheses. Then you encounter some evidence E. How should your opinions about the hypotheses update? Let us consider the matter semi-formally before looking at the details in Section 3.

Bayesians say that your opinions should update to

$$p_E(H_i) = p(H_i|E) =_{df} \frac{p(H_iE)}{P(E)}. \tag{1}$$

Equivalently, we can view Bayesian conditionalization as a matter of updating prior credences $p(H_i)$ in proportion to the likelihoods $P(E|H_i)$ that hypotheses assign to observed evidence. That is,

$$p_E(H_i) = \frac{p(H_i)p(E|H_i)}{\sum\limits_{H' \in \mathcal{H}} p(H')p(E|H')}. \tag{2}$$

The Bump Rule adjusts this form of likelihood-based updating by giving an additive bump to the hypotheses that best explain the evidence.

For some fixed constant c, let $\textsc{Bump}(H_i) = c$ if H_i best explains the agent's total evidence, c/n if H_i is one of n explanations tied for best, and 0 otherwise. Upon observing E, the Bump Rule instructs agents to update to:

$$p_E(H_i) = \frac{p(H_i)p(E|H_i) + \textsc{Bump}(H_i)}{\sum\limits_{H' \in \mathcal{H}} p(H')p(E|H') + \textsc{Bump}(H')}. \tag{3}$$

To illustrate, consider an example from Douven (2013).

Suppose that an agent observes flips of a fair coin. She wants to decide between the hypotheses $H_0, \ldots, H_{10}$ where H_i holds that the coin has bias $0.1i$ towards heads. She currently assigns equal probability to each H_i. Say that the best explanation is the hypothesis or hypotheses whose proposed bias is closest to the current frequency of observed heads. Set the explanationist bump c equal to 0.1. Let T_k be the proposition that

Hypothesis	Bayes T_1	Bayes $T_1, \neg T_2$	Bump T_1	Bump $T_1, \neg T_2$
$\mathbf{H_0}$	0.182	0.000	0.318	0.000
$\mathbf{H_1}$	0.164	0.055	0.136	0.039
$\mathbf{H_2}$	0.146	0.097	0.121	0.069
$\mathbf{H_3}$	0.128	0.127	0.106	0.091
$\mathbf{H_4}$	0.110	0.146	0.091	0.104
$\mathbf{H_5}$	0.091	0.152	0.076	0.394
$\mathbf{H_6}$	0.073	0.146	0.061	0.104
$\mathbf{H_7}$	0.055	0.127	0.045	0.091
$\mathbf{H_8}$	0.037	0.097	0.030	0.069
$\mathbf{H_9}$	0.019	0.055	0.015	0.039
$\mathbf{H_{10}}$	0.000	0.000	0.000	0.000

Table 1: Posterior credences in hypotheses H_i after observations T_1 and $T_1, \neg T_2$ on Bayesian conditionalization and Bump Rule

the k-th flip is observed to be tails.

Suppose the agent observes T_1, then $\neg T_2$. A Bayesian agent updates to push moderate probability mass towards the center (Table 1, Figure 1). For example, $p(H_5)$ updates to 0.152 and $p(H_4)$ updates to 0.146. By contrast, the Bump Rule updates much more aggressively towards the central hypothesis H_5, based on the fact that H_5 best explains the observed flips. For example, $p(H_5)$ updates to 0.394, whereas $p(H_4)$ updates only to 0.104.

That is the basic idea. But the details are a bit more complicated. Some of these complications will matter a great deal.

3 Bayes and the bump: Details

3.1 Bayesian updating

Let Ω be a finite set of states and $\mathcal{A}$ be an algebra over Ω. Let $\Delta(\mathcal{A})$ be the set of probability functions on $\mathcal{A}$.

Classically, an update rule takes two inputs: a prior probability function $p \in \Delta(\mathcal{A})$, and

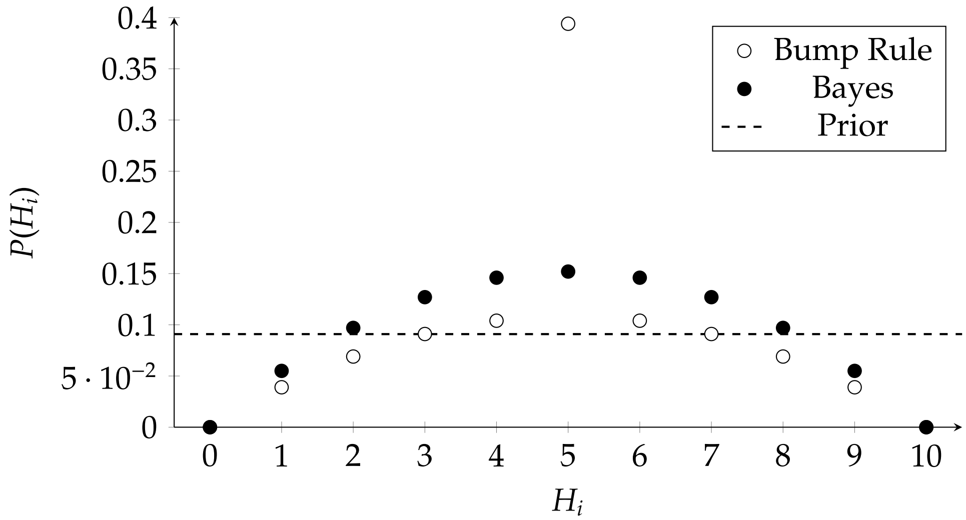

Figure 1: Posterior probabilities after observing $T_1, \neg T_2$.

an item of new evidence $E \in \mathcal{A}$. The rule returns a posterior probability function. That is, an *update rule* is a function $\mathcal{U} : \Delta(\mathcal{A}) \times \mathcal{A} \longrightarrow \Delta(\mathcal{A})$.

Bayesian conditionalization takes as input a prior probability $p \in \Delta(\mathcal{A})$ and an item of evidence $E \in \mathcal{A}$, returning the posterior probability function $p(*|E)$:

$$\mathcal{U}_B(p, E)(H) = p(H|E) =_{df} \frac{p(HE)}{p(E)}. \tag{4}$$

Compatibilists (Lange 2022; Lipton 2004; Okasha 2000) take Bayesian updating to be fully compatible with IBE. But incompatibilists (Douven 1999, 2013; Douven and Wenmackers 2017; Douven 2020, 2022) want to go further.

The Bump Rule is the leading incompatibilist proposal. It turns out that the Bump Rule takes us further afield from Bayesian updating than it might appear.

3.2 History- and hypothesis-dependent updating

To state the Bump Rule, we need to extend the classical framework in two ways. First, on the Bump Rule, explanation is a relation between a hypothesis and an agent's total evidence. As a result, the Bump Rule is history-dependent. While only the most recent item of evidence is used to assess likelihoods, all evidence received throughout the agent's history is used to assess explanatory fit.

Second and less obviously, we need to restrict the domain of propositions being updated. Bayesian conditionalization tells agents how to update credences over any algebra $\mathcal{A}$ of propositions - a set closed under complements, unions and intersections. The Bump Rule tells us how to update credences over a *partition* $\mathcal{H}$ of hypotheses - a set of pairwise incompatible hypotheses that jointly exhaust the state space Ω. We will see in Section 3.4 that the Bump Rule can be characterized as an update rule over a restricted algebra generated by the hypotheses $\mathcal{H}$. But the Bump Rule will not be definable over arbitrarily rich algebras of propositions. The relevant algebra is always tied to a hypothesis partition.

This form of hypothesis-dependence is an under-appreciated difference between the Bump Rule and Bayesian conditionalization. Hypothesis-dependent updating rules are common in statistical applications, where a hypothesis partition is rich enough to capture many modelers' interests. However, many philosophical Bayesians have the *ambitious Bayesian* project of characterizing the static and dynamic behavior of a probability function over all propositions of interest to an agent. Hypothesis-dependent update rules may not be suitable for this ambitious Bayesian project, since there is no agreed-upon way to embed them within the ambitious Bayesian project. I return to this point in Section 10.3.

3.3 Formalizing history- and hypothesis-dependence

Let a *history* in $\mathcal{A}$ be a finite sequence $\mathbf{E} = (E_1, \ldots, E_n)$ of elements of $\mathcal{A}$. The intended interpretation is that an agent with history $\mathbf{E}$ has received total evidence $\{E_1, \ldots, E_n\}$ in order $E_1, E_2, \ldots, E_n$. In our example, $\mathbf{E}$ records the observed coinflips. Let $\mathcal{E}(\mathcal{A})$ be the set of histories on $\mathcal{A}$.

If $\mathbf{E}$ is a history, and E_{n+1} is an element of $\mathcal{A}$, let $\mathbf{E} \oplus E_{n+1} = (E_1, \ldots, E_n, E_{n+1})$ be the history formed by appending E_{n+1} to $\mathbf{E}$. The intended interpretation is that $\oplus$ represents the effect of learning through appending new items of evidence to the agent's history. In our example, $\oplus$ appends the most recently observed coinflip to the evidence history.

A set $\mathcal{H}$ of hypotheses is a finite partition $\{H_1, \ldots, H_n\}$ of Ω. Closing $\mathcal{H}$ under unions, intersections and complements generates a sub-algebra $\mathcal{A}(\mathcal{H})$ of $\mathcal{A}$. In our example, $\mathcal{A}(\mathcal{H})$ consists of all possible claims about the truth and falsity of one or more hypotheses $H_0, \ldots, H_{10}$ about the coin's bias.

Because the agent's prior $\hat{p}$ on $\mathcal{A}(\mathcal{H})$ has restricted domain, we cannot define likelihoods directly in terms of the prior. For example, we cannot let the likelihood of E on H be defined as $\hat{p}(EH)/\hat{p}(H)$, because $\hat{p}(EH)$ is often undefined. Nevertheless, we can often define likelihoods theoretically based on our understanding of what hypotheses claim. Let a *likelihood function* be a map $\pi : \mathcal{A} \times \mathcal{A} \longrightarrow \mathbb{R}$, with the intended interpretation that $\pi(E, H)$ is the likelihood of E given H. Denote $\pi(E, H)$ as $\pi(E|H)$ and let $\Pi(\mathcal{A})$ be the set of likelihood functions over $\mathcal{A}$.

Once a hypothesis partition $\mathcal{H}$ is fixed, a **history-dependent $\mathcal{H}$-relative update rule** is a function $\hat{U} : \Delta(\mathcal{A}(\mathcal{H})) \times \Pi(\mathcal{A}) \times \mathcal{A} \times \mathcal{E}(\mathcal{A}) \longrightarrow \Delta(\mathcal{A}(\mathcal{H}))$. $\hat{U}$ takes four inputs: a prior $\hat{p}$ over the sub-algebra $\mathcal{A}(\mathcal{H})$, a likelihood function π over the full algebra, an item of evidence $E \in \mathcal{A}$, and an evidence history $\mathbf{E} \in \mathcal{E}(\mathcal{A})$. $\hat{U}$ returns a posterior $\hat{q}$ over $\mathcal{A}(\mathcal{H})$. In our example, history-dependent updating takes a prior $\hat{p}$ over claims about bias hypotheses H_i, likelihoods $\pi(T_k|H_i) = 1 - 0.1i$, an observed coinflip E and a history $\mathbf{E}$ of previously observed coinflips, then returns a posterior $\hat{q}$ over claims about bias hypotheses.

3.4 The Bump Rule

Relative to a given hypothesis partition $\mathcal{H}$, the Bump Rule partially characterizes a history-dependent $\mathcal{H}$-relative update rule by stating posterior probabilities over the hypothesis partition $\mathcal{H}$:

$$\hat{U}(\hat{p}, \pi, E, \mathbf{E})(H) = \frac{\hat{p}(H)\pi(E|H) + \textsc{Bump}(E, \mathbf{E}, H)}{\sum_{H' \in \mathcal{H}} \hat{p}(H')\pi(E|H') + \textsc{Bump}(E, \mathbf{E}, H')}. \tag{5}$$

Here we extend the previous definition of the explanatory Bump to make explicit its dependence on both the most recent evidence and the previous evidence history, as:

$$\textsc{Bump}(E, \mathbf{E}, H) = \begin{cases} c/n & H \text{ is one of the n best explanations of } \mathbf{E} \oplus E \\ \\ 0 & \text{else} \end{cases} . \tag{6}$$

Note that (5) only characterizes $\hat{U}$ over the hypothesis partition $\mathcal{H}$. We can extend this characterization to $\mathcal{A}(\mathcal{H})$ using disjoint additivity.[2]

This concludes our characterization of Bayesian updating and the Bump Rule. How should we evaluate the Bump Rule?

4 Challenges for the Bump Rule

The Bump Rule is a normalized linear combination of two updating rules. The first is Bayesian conditionalization. The second is empirical risk minimization, the rule which puts full probability mass on the current best-fitting hypothesis or hypotheses.

Taking linear combinations of updating rules is a risky business. Combinations often inherit the pathologies of both rules and the coherence properties of neither rule.

To illustrate these risks, let us consider three challenges for the Bump Rule. The first involves incoherence: the Bump Rule makes updating order-dependent (Section 4.1). The next two challenges involve pathologies inherited from empirical risk minimization. The Bump Rule assigns nonzero probability to hypotheses which give zero likelihood to the evidence (Section 4.2), and shifts probability mass against the direction of likelihoods (Section 4.3). These challenges are not meant to be exhaustive, but rather to illustrate the kinds of challenges facing the Bump Rule.

4.1 Order-dependence

Suppose you tell me that a coin with unknown bias has been flipped twice. One flip landed heads, and one landed tails. You ask for my views about the bias of the coin. I

[2] Set $\hat{q}(\emptyset) = 0$ and note that remaining elements of $\mathcal{A}(\mathcal{H})$ have the form $\bigcup_{i \in \mathcal{I}} H_i$ for some non-empty index set $\mathcal{I} \subseteq \{0, 1, \ldots, 10\}$. Disjoint additivity forces $\hat{q}(\bigcup_{i \in \mathcal{I}} H_i) = \sum_{i \in \mathcal{I}} \hat{q}(H_i)$.

respond that I cannot tell you until I know the order of the coinflips. You will probably think something is amiss with me. Views about the bias of a coin have nothing to do with the order of coinflips.

More generally, many theorists would like an agent's final credence to be independent of the order in which the evidence is received. We can express this requirement by considering finite streams $\mathbf{E}_{\text{new}}$ of incoming evidence, regarded as evidence histories.

Standard update rules $\mathcal{U} : \Delta(\mathcal{A}) \times \mathcal{A} \longrightarrow \Delta(\mathcal{A})$ on single items of evidence can be extended to update rules $\mathcal{U} : \Delta(\mathcal{A}) \times \mathcal{E}(\mathcal{A}) \longrightarrow \Delta(\mathcal{A})$ on evidence streams by iteratively applying $\mathcal{U}$ to items of evidence in the order that they are presented. For example $\mathcal{U}(p, (T_1, \neg T_2))$ gives the probability function that results from updating p on a tails flip followed by a heads flip, with both updates conducted according to $\mathcal{U}$.

One way for the order of evidence to matter is through failures of exchangeability, in which previous flips affect the hypothesis-relative likelihoods of subsequent flips.[3] For example, if $p(T_2|H_i \wedge T_1)$ is not the same as $p(T_2|H_i \wedge \neg T_1)$, then it matters a great deal whether the first flip lands heads or tails. But that is not what happened in our example of coinflips. Conditional on any bias hypothesis, the likelihood of subsequent flips does not depend on past flips.

More generally, say that:

> Evidence stream $\mathbf{E}$ is **$\mathcal{H}$-conditionally-independent** if for any $H_i \in \mathcal{H}$ and any $E_j \in \mathbf{E}$ we have $\pi(E_j|H_i \wedge \bigcap_{k<j} E_k) = \pi(E_j|H_i)$.

Here π denotes the likelihood function, which exhibits its standard behavior in the Bayesian setting. $\mathcal{H}$-conditional independence is a general form of the requirement that likelihoods do not depend on past observations.

Many theorists would like updating over $\mathcal{H}$-conditionally-independent evidence streams to be order independent.

[3]This paper works with a likelihood-based consequence of standard exchangeability requirements because likelihoods are sometimes defined separately from priors.

(Standard Order Independence) For all priors p on $\mathcal{A}$, all $\mathcal{H}$-conditionally-independent evidence streams $\mathbf{E}_{\text{new}}$, and all permutations σ, $U(p, \sigma(\mathbf{E}_{\text{new}})) = U(p, \mathbf{E}_{\text{new}})$.[4]

For example, Standard Order Independence forces $\mathcal{U}(p, (T_1, \neg T_2)) = \mathcal{U}(p, (\neg T_1, T_2))$. It is a familiar result that Bayesian conditionalization satisfies Standard Order Independence.[5]

In exactly the same way, history-dependent $\mathcal{H}$-relative update rules $\hat{U} : \Delta(\mathcal{A}(\mathcal{H})) \times \Pi(\mathcal{A}) \times \mathcal{A} \times \mathcal{E}(\mathcal{A}) \longrightarrow \Delta(\mathcal{A}(\mathcal{H}))$ on single items of evidence can be extended to update rules on evidence streams $\hat{U} : \Delta(\mathcal{A}(\mathcal{H})) \times \Pi(\mathcal{A}) \times \mathcal{E}(\mathcal{A}) \times \mathcal{E}(\mathcal{A}) \longrightarrow \Delta(\mathcal{A}(\mathcal{H}))$. For example, $\hat{U}(\hat{p}, \pi, \mathbf{E}_{new}, \mathbf{E}_{old})$ gives the results of updating $\hat{p}$ with evidence history $\mathbf{E}_{old}$ on the new stream of evidence $\mathbf{E}_{new}$ one item at a time, according to $\hat{U}$.[6] In this vocabulary, order-independence requires:

(Extended Order Independence) For all priors $\hat{p}$ on $\mathcal{A}(\mathcal{H})$, likelihoods π, **$\mathcal{H}$-conditionally-independent** evidence streams $\mathbf{E}_{\text{old}} \oplus \mathbf{E}_{\text{new}}$, and permutations σ, $\hat{U}(\hat{p}, \pi, \sigma(\mathbf{E}_{\text{new}}), \mathbf{E}_{\text{old}}) = \hat{U}(\hat{p}, \pi, \mathbf{E}_{\text{new}}, \mathbf{E}_{\text{old}})$.

For example, Extended Order Independence forces the agent to update in the same way upon observing a tails and then a heads as she would upon observing a heads and then a tails, no matter her priors or previous evidence.

The Bump Rule does not satisfy Extended Order Independence. The easiest way to see this is to consider the sequence $H^{19}T$ of 19 heads followed by a single tails. Updating on the first nineteen flips drives the probability of H_{10} very high. However, the last observation provides an explanatory bump to H_9, which now explains observed frequencies just as well as H_{10} does. The resulting probability of H_{10} is 0.495, with the bulk of the remaining probability going to H_9 at a final probability of 0.503.

[4]To rigorously define permutation of evidence streams, evidence streams $\mathbf{E} = (E_1, \ldots, E_n)$ should be unpacked into indicator random variables for the relevant events as $\mathbf{E} = (1_{E_1} = x_1, \ldots, 1_{E_n} = x_n)$ so that $\sigma(\mathbf{E}) = (1_{E_1} = x_{\sigma(1)}, \ldots, 1_{E_n} = x_{\sigma(n)})$.

[5]If a proof is desired, take $\lambda = 1$ in Equation (16).

[6]It is important to update the likelihood function, so that updating on E_j uses likelihoods $\pi(E_j|H_i \wedge \bigcap_{k<j} E_j)$ rather than $\pi(E_j|H_i)$. While this distinction will not matter for $\mathcal{H}$-conditionally independent evidence streams, it will be important in later sections.

By contrast, consider the sequence TH^{19} of a single tails followed by nineteen heads. Updating on the first flip drives the probability of H_1 to zero, where it remains until the very last flip. Here, H_{10} again shares an explanatory bump with H_9, but because of its poor start, H_{10} receives a final probability of 0.050, with H_9 now receiving probability 0.903. This is a sizable order effect.

4.2 The zero likelihood principle

What is striking about the previous example is not only that the agent assigns different credence to H_{10} after observing $H^{19}T$ rather than TH^{19}. Another striking feature is that she assigns any nonzero credence at all to H_{10} after either sequence of observations. After all, H_{10} says that the coin has bias 1 towards heads. This is standardly taken to imply that any finite sequence involving a tails flip has zero probability. And that, in turn, is standardly taken to imply that an agent's credence in H_{10} should fall to zero after observing either $H^{19}T$ or TH^{19}.

A natural way to express this requirement is to say that if E has zero likelihood on some hypothesis H, then updating on E should drive the probability of H to zero. Expressed as a requirement on standard update rules, this means:

> **(Zero-Likelihood Principle)** For all priors $p \in \Delta(\mathcal{A})$ and propositions $E, H \in \mathcal{A}$, if $p(E|H) = 0$ then $\mathcal{U}(p, E)(H) = 0$.

Bayesians respect the Zero-Likelihood Principle, because $p(H|E) = p(E|H)\frac{p(H)}{p(E)}$ is forced to zero when $p(E|H) = 0$.[7]

Expressed as a requirement on history-dependent, hypothesis-relative update rules, the Zero-Likelihood Principle becomes:

[7] A separate argument is needed for the case that $p(E) = 0$. While the details of that argument differ among leading approaches to updating on probability-zero events (Hájek 2003; Jeffrey 1965; Popper 1959), to my knowledge all leading approaches respect the Zero-Likelihood Principle. But perhaps the easiest bugfix would be to restrict the Zero-Likelihood Principle to the case $p(E) > 0$ and note that all of our examples satisfy this condition.

(Extended Zero-Likelihood Principle) For all priors $\hat{p} \in \Delta(\mathcal{A}(\mathcal{H}))$, likelihoods $\pi \in \Pi(\mathcal{A})$, hypotheses $H \in \mathcal{H}$, evidence $E \in \mathcal{A}$ and histories $\mathbf{E} \in \mathcal{E}(\mathcal{A})$, if $\pi(E|H) = 0$ then $\hat{\mathcal{U}}(\hat{p}, \pi, E, \mathbf{E})(H) = 0$.

As we saw above, the Bump Rule violates the Extended Zero-Likelihood Principle because an explanatory bump can be provided to hypotheses which assign zero likelihood to the evidence. For example, our agent assigns nonzero probability to H_{10} after observing nineteen heads flips followed by a tails flip. However, that tails flip has likelihood zero on H_{10}.

4.3 Minimal likelihoodism

Suppose I tell you that I am five times more confident in the hypothesis H_5 that a coin is fair than in the hypothesis H_9 that it has bias 0.9 towards heads. Then I see another coinflip. The coin lands heads. Aha! Now, I explain, I am ten times more confident in H_5 than in H_9.

This is puzzling behavior. My latest observation was much more enthusiastically predicted by H_9 than it was by H_5. It therefore seems strange for me to become more confident in H_5, relative to H_9, after observing an outcome much more enthusiastically predicted by H_9. But the Bump Rule allows just this sort of behavior.

Given hypotheses H, H' and evidence E, let the likelihood ratio be defined as:

$$\mathcal{L}(H, H', E) = \frac{\pi(E|H)}{\pi(E|H')} \tag{7}$$

For generality, we use the notation $\pi(E|H)$ for likelihoods, with the understanding that $\pi(E|H) =_{df} p(E|H)$ in standard Bayesian models.

Bayesians think that probability ratios should update proportionally to likelihood ratios.[8] That is:[9]

[8]Note that Bayesians may reject some stronger versions of the Likelihood Principle (Fitelson 2007; Steel 2007).

[9]This excludes the case $\pi(E|H') = 0$, where $\mathcal{U}_B(p, E)(H') = 0$ and both sides of (8) are undefined.

$$\textbf{(Likelihoodism)}\ \frac{\mathcal{U}_B(p,E)(H)}{\mathcal{U}_B(p,E)(H')} \propto \mathcal{L}(H,H',E). \tag{8}$$

While many are attracted to Likelihoodism, Likelihoodism is not a principle that incompatibilist explanationists will uncontroversially accept. After all, the entire point of incompatibilist explanationism is to go beyond likelihood ratios in updating.

However, a more minimal likelihoodist requirement may be costlier to deny. Suppose that E has higher likelihood on some hypothesis H than on a competing hypothesis H', so that $\mathcal{L}(H,H',E) > 1$. Then it is natural to assume that H should remain at least as plausible, relative to H', as it was before. That is, working now in terms of history- and hypothesis-dependent update rules:

> **(Minimal Likelihoodism)** For all hypotheses $H, H' \in \mathcal{H}$ and evidence $E \in \mathcal{A}$, if $\mathcal{L}(H,H',E) > 1$, then for all histories $\mathbf{E}$, $\frac{\hat{U}(\hat{p},\pi,E,\mathbf{E})(H)}{\hat{U}(\hat{p},\pi,E,\mathbf{E})(H')} \geq \frac{p(H)}{p(H')}$.

After all, if the only thing we learn is something that is more favorable to H than to H', then it seems we should not become less favorably disposed towards H relative to H'.

Table 1 shows that the Bump Rule violates Minimal Likelihoodism. After observing a tails flip, the ratio $p(H_9)/p(H_5)$ of the probabilities assigned to biases 0.9 and 0.5 is 0.197. Then the agent observes a heads flip, which has substantially higher likelihood on H_9 than it does on H_5. Nonetheless, the ratio $p(H_9)/p(H_5)$ updates downwards to 0.099, so that H_9 has only about half of its previous relative probability against H_5, despite H_9 having made the second flip more likely than H_5 did.

It is certainly open to explanationists to reject this form of Minimal Likelihoodism. Perhaps some will insist that violations of Minimal Likelihoodism are not pathologies when they reflect reassessments of what best explains the agent's total evidence. Minimal Likelihoodists will think that they look more like double-counting (Sober 2008; van Fraassen 1980, 1989), since prior features of the agent's evidence history are re-used to update against the direction of the most recent evidence. I leave it to readers to decide what conclusion to draw.

4.4 Taking stock

We began this section with the idea that linear combinations of two learning rules often inherit the coherence of neither and the pathologies of both. In partial vindication of this view, we saw that the Bump Rule is order-dependent (Section 4.1), a form of incoherence shared by neither component rule. We also saw that the Bump Rule violates the Zero-Likelihood Principle (Section 4.2) and Minimal Likelihoodism (Section 4.3), both pathologies associated with empirical risk minimization. These examples are intended to illustrate, rather than exhaust, the common wisdom, and to point the way towards where further challenges might be found.

I hope that this discussion goes some way towards reinforcing the common wisdom against summing learning rules of different types. Is there a different rule which gives explanationists much of what they wanted without resorting to this summative strategy?

5 Tempered updating

Bayesians often argue that there is no need to modify Bayes' rule to incorporate IBE, since explanatory considerations are already built into likelihoods, and perhaps also to priors (Bird 2017; Huemer 2009; Weisberg 2009). The likelihood $p(E|H)$ expresses the degree to which H would explain E. This explanatory power is rewarded in updating. In fact, when we express Bayesian conditionalization as:

$$p_E(H_i) = \frac{p(H_i)p(E|H_i)}{\sum_{H' \in \mathcal{H}} p(H')p(E|H')}. \tag{9}$$

we see that the *only* force driving Bayesian updating away from the prior is the comparative power of hypotheses to explain the evidence, as expressed in the likelihoods $p(E|H)$.

However, even if we grant the Bayesian contention that likelihoods encode explanatory strength, we do not have to grant that Bayes' rule gives correct weight to explanatory considerations. Suppose we take seriously the explanationist contention that explanatory considerations should be given additional weight during updating. One natural way to capture this contention is through tempered updating (Grünwald and van Ommen 2017; Holmes and Walker 2017).[10]

Expressed as a hypothesis-dependent, but history-independent update rule, standard Bayesian updating requires:

$$\hat{\mathcal{U}}_B(\hat{p}, \pi, E)(H) = \frac{\hat{p}(H)\pi(E|H)}{\sum_{H' \in \mathcal{H}} \hat{p}(H')\pi(E|H')}. \tag{10}$$

Tempered updating modifies (10) by weighting posteriors by a constant *temperature parameter* λ:

$$\hat{\mathcal{U}}_B(\hat{p}, \pi, E)(H) = \frac{\hat{p}(H)\pi(E|H)^{\lambda}}{\sum_{H' \in \mathcal{H}} \hat{p}(H')\pi(E|H')^{\lambda}}. \tag{11}$$

The temperature parameter λ controls the aggressiveness of updating towards the hypotheses which best explain the data. Taking $\lambda = 1$ recovers Bayesian updating. Higher temperatures $\lambda > 1$ provide an explanationist boost, updating more aggressively in the direction of explanatory considerations embedded in likelihoods. Lower temperatures $\lambda < 1$ provide a conservative cooling, pulling updates closer towards the prior.

Tempered updating is a well-studied inference rule (Grünwald and van Ommen 2017; Holmes and Walker 2017; Wu and Martin 2023). While no alternative to conditionalization will have every property that Bayesians might desire, we will see that tempered updating avoids the challenges facing the Bump Rule (Section 6) while retaining many of its attractive features (Sections 7-9).

But first, let us introduce two additional characterizations of tempered updating. These characterizations will help us to see what tempered updating amounts to and why tempered updating may be attractive to explanationists.

[10] Tempered updating is naturally understood in the context of generalized Bayesian updating (Bissiri et al. 2016). Tempered updating has been studied in other applications such as marginal likelihood approximation (Friel and Pettit 2008) and discounting historical data (Ibrahim and Chen 2000), though with distinct interpretations.

5.1 Learning rate interpretation

The first approach is to think of λ as a *learning rate* (Wu and Martin 2023) in the sense familiar from approaches such as reinforcement learning (Sutton and Barto 2018) and Carnap's (1950) λ-continuum. Raising the learning rate λ leads agents to generalize faster based on the explanatory considerations encoded in their current evidence. This often increases the chance of fast convergence, but also often raises the risk of spurious convergence. Lowering the learning rate λ leads agents to generalize more slowly, often reducing convergence speeds but decreasing the risk of spurious convergence.

At the slowest possible learning rate, $\lambda = 0$ and the agent refuses to learn at all. Here we recover:

(Dogmatism) $Pr_E(H) = Pr(H)$.

At the other extreme, as $\lambda \longrightarrow \infty$ the agent tends towards complete concentration of probability mass on the current best-fit.

(Empirical Risk Minimization) $Pr_E(H) = \begin{cases} 1 & \text{H is the best fit to E} \\ 0 & \text{else} \end{cases}$

As before, Empirical Risk Minimization can be modified to split probability mass among hypotheses that are tied for best.

On the learning rate interpretation, both Bayesians and explanationists pick learning rates that steer comfortably between the extremes of Dogmatism and Empirical Risk Minimization. Explanationists differ from Bayesians in that they pick somewhat higher learning rates λ, bringing them closer to Empirical Risk Minimization and further away from Dogmatism.

5.2 Loss minimization interpretation

A second way to think about λ is as a compromise between two aims (Bissiri et al. 2016). On the one hand, we want to *minimize empirical loss*. That is, we want a rule that fits observed data as closely as possible. On the other hand, we want to *regularize*. That is, we want to rein in the ability to overfit accidental features of observed data so that our model will perform well on unobserved data.

More concretely, suppose you have priors $\hat{p}$ over $\mathcal{A}(\mathcal{H})$ and assign likelihoods π. You receive some evidence E. You need to choose a posterior $\hat{q}$.

If you only had the first goal of minimizing empirical loss, you would pick $\hat{q}$ to make the data as likely as possible. A popular way of operationalizing this requirement is to say that the expected log-likelihood of the data should be maximized, or equivalently that the negative expected log-likelihood of the data should be minimized. That is, you would pick $\hat{q}$ to minimize:

$$-\sum_{H \in \mathcal{H}} \hat{q}(H) log(\pi(E|H)). \tag{12}$$

(12) is minimized by Empirical Risk Minimization, which assigns probability one to the current best-fitting hypothesis.

If you did this, you would fit the existing data very well. But you might be in for a nasty surprise as new data handed you a large loss. To avoid overfitting in this way, it is standard to choose posteriors which trade off empirical loss minimization against a regularization penalty chosen to prevent overfitting.

Bayesians regularize by pulling posteriors back towards the prior. That is, let $KL(\hat{q}||\hat{p}) = \sum_H \hat{q}(H) log(\hat{q}(H)/\hat{p}(H))$ be the Kullback–Leibler divergence between the posterior $\hat{q}$ and the prior $\hat{p}$, a generalization of the notion of distance between posterior and prior (Kullback and Leibler 1951; Staffel 2020). Bayesians choose posteriors to minimize:

$$-[\sum_{H \in \mathcal{H}} \hat{q}(H) log(\pi(E|H))] + KL(q||p). \tag{13}$$

The Bayesian posterior minimizes this expression (Bissiri et al. 2016).

In this sense, we might think of Bayesian updating as an equally-weighted compromise between the goals of fitting data and regularizing to avoid overfitting. In this perspective, we might ask what happens if we make a different compromise between these goals. It turns out (Grünwald and van Ommen 2017) that tempered updating minimizes:

$$-\lambda[\sum_{H\in\mathcal{H}} \hat{q}(H)log(\pi(E|H))] + KL(q||p). \tag{14}$$

Here the temperature parameter λ controls the relative importance of fitting data versus regularizing towards the prior. In this sense, we might see tempered updating as placing relatively more weight on the goal of data fitting as opposed to regularization.

This characterization dovetails nicely with recent suggestions that IBE might be seen as a more risk-seeking alternative to standard Bayesianism (Pettigrew 2021). Tempered updating accepts a greater risk of overfitting by extracting spurious trends from data, in exchange for a greater ability to extract genuine trends from data. As we will see in Section 8, this increases the risk of large unanticipated losses, but often leads to accuracy improvements. Just as Lara Buchak (2013) suggests that different risk attitudes may sometimes be permissible in nonmental action, so too we might see the defender of tempered updating as suggesting that different risk attitudes may sometimes be appropriate in the mental action of updating credences.

6 Avoiding challenges

So far, we have seen what tempered updating requires and how the temperature parameter can be understood as a learning rate (Section 5.1) or as an emphasis on fitting data over regularizing (Section 5.2). The next order of business is to see what can be said in favor of tempered updating. Let us lead with the best news: tempered updating avoids the challenges raised in Section 4.

6.1 Order dependence

The first challenge was avoiding order-dependent updating. Tempered updating is order-independent in the natural way.

As before, for any history $\mathbf{E} = (E_1, \ldots, E_n)$, let $\hat{U}(\hat{p}, \pi, \mathbf{E})$ be the result of applying $\hat{U}$ iteratively from E_1 to E_n. Then for any permutation σ, tempered updating gives:

$$\hat{U}(\hat{p}, \pi, \sigma(\mathbf{E}))(H) = \frac{\hat{p}(H) \prod_{i=1}^{n} \pi(E_{\sigma(i)}|H \wedge \bigcap_{j<i} E_{\sigma(j)})^{\lambda}}{\sum_{H' \in \mathcal{H}} \hat{p}(H') \prod_{i=1}^{n} \pi(E_{\sigma(i)}|H' \wedge \bigcap_{j<i} E_{\sigma(j)})^{\lambda}} \tag{15}$$

On $\mathcal{H}$-conditionally-independent evidence streams, this simplifies to:

$$\hat{U}(\hat{p}, \pi, \sigma(\mathbf{E}))(H) = \frac{\hat{p}(H) \prod_{i=1}^{n} \pi(E_i|H)^{\lambda}}{\sum_{H' \in \mathcal{H}} \hat{p}(H') \prod_{i=1}^{n} \pi(E_i|H')^{\lambda}} \tag{16}$$

Equation (16) is the same order-independent application of likelihoods familiar from classical Bayesian updating, with the addition of the temperature parameter.

6.2 Zero-likelihood principle

The second challenge was capturing the Zero-Likelihood Principle. The Zero-Likelihood Principle requires the probability of a hypothesis H to update to zero upon observing evidence with likelihood $\pi(E|H) = 0$ on H. In this context, this means:

> **(Zero-Likelihood Principle)** For all priors $\hat{p} \in \Delta(\mathcal{A}(\mathcal{H}))$, likelihoods $\pi \in \Pi(\mathcal{A})$, hypotheses $H \in \mathcal{H}$, and evidence $E \in \mathcal{A}$, if $\pi(E|H) = 0$ then $\hat{\mathcal{U}}(\hat{p}, \pi, E)(H) = 0$.

The Zero-Likelihood Principle follows immediately from Equation (16). Observing evidence E with likelihood $\pi(E|H) = 0$ multiplies the numerator of the posterior probability by zero, driving the posterior itself to zero.

6.3 Minimal likelihoodism

The final challenge was capturing Minimal Likelihoodism. Minimal Likelihoodism requires that evidence which is more likely on some hypothesis H than on another hypothesis H' should not decrease the relative probability of H against H'. In this context, this means:

(Minimal Likelihoodism) For all hypotheses $H, H' \in \mathcal{H}$ and evidence $E \in \mathcal{A}$, if $\mathcal{L}(H, H', E) > 1$, then $\frac{\hat{U}(\hat{p},\pi,E)(H)}{\hat{U}(\hat{p},\pi,E)(H')} \geq \frac{\hat{p}(H)}{\hat{p}(H')}$.

Tempered updating satisfies Minimal Likelihoodism by way of satisfying a stronger likelihoodist principle:

(Tempered Likelihoodism) $\frac{\mathcal{U}_B(\hat{p},\pi,E)(H)}{\mathcal{U}_B(\hat{p},\pi,E)(H')} = \mathcal{L}(H, H', E)^{\lambda} \frac{\hat{p}(H)}{\hat{p}(H')}$.

It follows immediately that when $\mathcal{L}(H, H', E) > 1$, $\frac{\mathcal{U}_B(\hat{p},\pi,E)(H)}{\mathcal{U}_B(\hat{p},\pi,E)(H')} > \frac{\hat{p}(H)}{\hat{p}(H')}$, with the proportional increase governed by a power of the likelihood. Tempered Likelihoodism comes apart from full-blown Likelihoodism (Equation (8)) only in being more eager to satisfy Minimal Likelihoodism because the likelihood $\mathcal{L}(H, H', E)$ is raised to an exponent $\lambda > 1$.

So far, we have seen that Tempered Updating avoids the challenges facing the Bump Rule. The next order of business is to show that Tempered Updating captures many of the same behaviors that attracted defenders of the Bump Rule.

7 Simulations: Learning speed

Defenders of the Bump Rule argue that it shows attractive behavior in simulations (Douven 2013; Douven and Wenmackers 2017; Douven 2020, 2022). Specifically, under many conditions, the Bump Rule learns the true hypothesis more quickly than Bayesian conditionalization does and achieves higher accuracy after a fixed number of steps.[11]

From this, defenders of the Bump Rule draw two conclusions. The weaker conclusion is that the Bump Rule may provide an attractive solution to the speed-accuracy tradeoff

[11] Following Pettigrew (2021), I will not discuss simulations involving social learning (Douven and Wenmackers 2017) as these introduce further assumptions that mainline Bayesians may reject.

in learning (Douven 2022; Heitz 2014; Wickelgren 1977), providing quicker mean learning rates with tolerable decreases in mean accuracy.

The stronger conclusion adopts an externalist, ecological reading of rationality (Douven 2020; Morton 2017; Schmidt 2019).[12] To say that rationality is ecological is to say that rationality is a matter of performing well in an environment. On externalist theories of ecological rationality, the relevant environment is not an expectation over various possible environments but rather the agent's actual environment. On this view, it is rational for agents to use strategies which perform well in their current environment.

In some environments, the Bump Rule learns more slowly or less accurately than Bayesian conditionalization does. But in many environments, the Bump Rule learns more quickly than Bayesian conditionalization does, and also shows greater accuracy. This allows defenders of externalist, ecological theories of rationality to draw the stronger conclusion that the Bump Rule is ecologically rational in many environments, no matter the relative importance of speed and accuracy in learning.

This section argues that tempered updating strikes an attractive balance between Bayesian conditionalization and the Bump Rule in learning speed, capturing a substantial portion of the Bump Rule's speed advantage in both friendly and unfriendly environments.

Section 8 shows that tempered updating strikes a similarly attractive balance in accuracy. Across two leading accuracy measures, the Bump Rule reduces mean accuracy when compared to Bayesian conditionalization. We will see that tempered updating makes up a sizable portion of this accuracy difference across accuracy measures in friendly environments. Tempered updating captures many of the performance advantages of the Bump Rule in some classes of difficult environments, and avoids its pitfalls in others.

Section 9 assesses implications for the choice between rules, focusing on the speed-accuracy tradeoff as well as the externalist ecological perspective.

[12] An alternative non-ecological reading of rationality (Thorstad forthcoming) might not support this stronger conclusion.

7.1 Speed: The good case

Continue our example from Section 2. An agent begins with a uniform prior over the hypotheses $\{H_0, H_1, \ldots, H_{10}\}$ about the bias of a fair coin. She observes coinflips and updates according to her favorite learning rule.

It is a familiar refrain that most learning rules eventually discover the truth in friendly problems. For this reason, a common question is how quickly a learning rule manages to discover the truth. To answer this question, Douven (2013) varies the true bias hypothesis H_i and runs 1,000 simulations for each true bias. Letting the learning speed of each rule be measured by the number of flips observed until it assigns at least 99% credence to the true hypothesis, Douven finds that the Bump Rule with $c = 0.1$ learns significantly faster than Bayesian conditionalization does in most environments.

Let us expand Douven's simulation to incorporate tempered updating with a moderate temperature $\lambda = 2$. Let us also compare both rules against the limiting case of empirical risk minimization (ERM), in which the agent puts full probability on the hypothesis or hypotheses that best fit current observed frequencies. Table 2 shows the mean learning speed for each rule across variations in the true bias.

Table 2 bears out Douven's claim that the Bump Rule learns more quickly on average than Bayesian updating does in this friendly problem. Tempered updating captures a substantial percentage of this speed improvement.

At the same time, speed in friendly learning problems is not everything. If we were only concerned with mean learning speed in friendly problems, we would adopt empirical risk minimization based on its superior speed in Table 2.

7.2 Speed: The bad case

What happens when we make the learning problem difficult? Sometimes, rules that perform well on friendly problems perform less well on difficult problems. In Section 8, we will see some performance reversals of this type when we are concerned with accuracy

True Bias	Bayes	Tempered ($\lambda = 2$)	Bump (c = 0.1)	ERM
0	44.0	22.0	19.0	1.0
0.1	127.4	66.0	46.5	18.5
0.2	228.0	132.2	81.3	18.3
0.3	307.3	176.4	101.1	16.9
0.4	355.7	203.4	105.8	20.6
0.5	362.4	209.3	106.9	13.8
0.6	348.0	201.1	108.1	23.5
0.7	302.9	176.9	96.5	17.3
0.8	232.1	133.1	81.8	18.4
0.9	128.4	65.8	44.1	18.1
1	44.0	22.0	19.0	1.0

Table 2: Mean learning speed (flips) by rule and true bias condition, good case, 1,000 simulations per true bias

rather than speed. However, the results of Section 7.1 are reasonably robust against two types of difficult problems.

One way to increase the difficulty of the learning problem is to fine-grain the hypothesis space. Specifically, instead of considering eleven bias hypotheses H_i that the true bias is $0.1i$, we consider one hundred and one bias hypotheses H'_i that the true bias is $0.01i$. Fine-graining could punish the Bump Rule and empirical risk minimization because of their winner-takes-all dynamic. Explanatory bumps are awarded to the best-fitting hypothesis, but not to close neighbors which may have nearly as strong explanatory credentials.

Table 3 reports the results of 1,000 simulations over this fine-grained hypothesis space, with randomized true bias. Although tempered updating learns considerably faster than Bayesian conditionalization, the Bump Rule shows a substantial speedup and ERM learns still more quickly.

Another way to increase the difficulty of the learning problem is to extremize priors. As we move away from a uniform prior and towards priors which concentrate probability mass on a small number of bias hypotheses, we increase the chance that the agent begins with a very low prior in the true bias hypothesis. This is a more difficult learning problem.

Table 3 reports the results of 1,000 simulations with extremized priors, drawn ran-

	Bayes	Tempered ($\lambda = 2$)	Bump (c = 0.1)	ERM
Fine-graining	23,784.62	13,645.85	1,434.73	239.27
Extremal priors	590.19	307.65	71.23	15.41

Table 3: Mean learning speed (flips) by rule, bad cases, 1,000 simulations

domly from a symmetric Dirichlet(0.1) distribution over the space of possible priors.[13] Again, tempered updating shows a speed advantage of Bayesian updating, with further improvements by the Bump Rule and Empirical Risk Minimization.

This discussion suggests that the ranking of rules by learning speeds is relatively robust. Even under difficult conditions, tempered updating learns more quickly than Bayes' rule does, and the Bump Rule learns more quickly still.

Because both rules learn faster than Bayesian conditionalization, even in expectation, there is no need to retreat to an externalist, ecological perspective to assess the hierarchy of learning speeds. If we are concerned only with learning speed, then even an expectational perspective gives both rules an advantage over Bayesian conditionalization.

8 Simulations: Accuracy

What happens when we study accuracy rather than learning speed? Both tempered updating and the Bump Rule sometimes pay accuracy costs for their increased speed, but there are important differences in the patterns of accuracy costs, and both rules sometimes reap accuracy benefits.

Following Douven (2022), let us consider two ways of measuring accuracy. The first is the Brier Score. The Brier Score penalizes agents for the squared distance of their beliefs from the truth. The Brier inaccuracy of a credence function p is then:

$$\textbf{(Brier Score)} \quad \frac{1}{|\mathcal{H}|}\sum_i (p(H_i) - T(H_i))^2 \tag{17}$$

[13] A Dirichlet(0.1) distribution assigns each prior $(x_0, \ldots, x_{10})$ on $(H_0, \ldots, H_{10})$ probability proportional to $\prod_{i=0}^{10} x_i^{-0.9}$ which favors extremized priors.

where $T(H_i)$ returns 1 if H_i is the true bias hypothesis and 0 otherwise. Under many conditions, Bayesian conditionalization minimizes expected Brier inaccuracy (Easwaran 2013; Greaves and Wallace 2006; Leitgeb and Pettigrew 2010a,b).

An alternative scoring rule is the log rule. The log rule differs from the Brier Score in two ways. First, it penalizes inaccuracy only on the true bias hypothesis H^*. Second, it measures distance logarithmically rather than by squared differences. The log-rule inaccuracy of credence function p is then:

$$\textbf{(Log Score)} \quad -ln(p(H^*)) \tag{18}$$

where H^* is the true bias hypothesis and ln is the natural logarithm.

Let us consider how each learning rule performs across candidate scoring rules in the same models considered above.

8.1 Accuracy: The good case

Tables 4-5 report the results of 1,000 simulations for each true bias in the good case. Results are reported in terms of the mean inaccuracy of rules after a fixed number of steps.

On the Brier Score (Table 4), inaccuracy of each rule decays in a qualitatively similar manner across steps, with some advantage for Bayes over tempered updating, and for tempered updating over the Bump Rule. By 250 steps, all rules have achieved an excellent Brier score of 0.01 or lower.

On the log score (Table 5), Bayes performs somewhat better than tempered updating, which performs significantly better than the Bump Rule. The worst performance of all is exhibited by empirical risk minimization, which regularly incurs infinite loss by giving probability zero to the true hypothesis.

This discussion suggests that in the good case, mean performance shows a speed-accuracy tradeoff, in which the ordering of rules by mean learning speed is precisely the reverse of their ordering by mean accuracy. Importantly, the existence of a speed-accuracy

Steps	Bayes	Tempered ($\lambda = 2$)	Bump (c = 0.1)	ERM
5	0.070	0.073	0.080	0.124
10	0.062	0.065	0.073	0.104
25	0.047	0.049	0.063	0.082
50	0.036	0.038	0.050	0.058
100	0.023	0.025	0.030	0.031
250	0.008	0.008	0.010	0.010
500	0.001	0.002	0.002	0.002

Table 4: Mean Brier inaccuracy by learning rule and steps elapsed, averaged over 1,000 trials for each true bias

Steps	Bayes	Tempered ($\lambda = 2$)	Bump (c = 0.1)	ERM
5	1.675	1.822	1.918	∞
10	1.376	1.523	1.762	∞
25	0.960	1.087	1.728	∞
50	0.671	0.777	1.810	∞
100	0.420	0.511	1.629	∞
250	0.147	0.200	1.090	∞
500	0.028	0.039	0.310	∞

Table 5: Mean log-rule inaccuracy by learning rule and steps elapsed, averaged over 1,000 trials for each true bias

tradeoff is invariant to the choice between Brier and log scoring rules. In fact, moving from the Brier score to the log score only heightens the speed-accuracy tradeoff.

Because these results concern mean inaccuracy, those persuaded by an externalist reading of ecological rationality may also want to know in how many specific environments each rule performs well. There are many ways to measure this. One question we can ask is in how many runs each rule performed best, in the sense of minimizing inaccuracy after a fixed number of runs.

Table 6 reports the results of 1,000 simulations for each true bias, averaged across true biases. The proportion of trials in which each rule performed best is reported across elapsed steps, focusing for brevity on the Brier score. Note that sums of rows may exceed one because in some cases, multiple rules tied for best performance.[14]

[14]For the purpose of computational simplification, rules were reported as tied when their scores differed by less than 10^{-12}. It could be productive to read this as a notion of 'near-ties' rather than strict ties.

Steps	Bayes	Tempered ($\lambda = 2$)	Bump (c = 0.1)	ERM
5	0.49	0.06	0.14	0.32
10	0.49	0.00	0.08	0.43
25	0.35	0.05	0.05	0.55
50	0.27	0.02	0.03	0.68
100	0.15	0.22	0.47	0.79
250	0.26	0.38	0.87	0.94
500	0.41	0.82	0.98	0.99

Table 6: Proportion of trials with lowest Brier inaccuracy after given number of steps, averaged over 1,000 trials for each true bias

With small data samples, both tempered updating and the Bump Rule produce the highest-accuracy results less often than Bayesian conditionalization does. This pattern reverses somewhat with large data samples. However, caution is needed before trading too much on this result, as empirical risk minimization robustly outperforms both tempered updating and the Bump Rule on this criterion.

8.2 Accuracy: The bad case

To study inaccuracy in difficult problems, a computational experiment was performed under the same fine-graining and extremal prior conditions used in Section 7.2. Tables 7-8 report mean Brier inaccuracy over 1,000 simulations of each condition.

With extremal priors (Table 7), both the Bump Rule and tempered updating produce lower Brier inaccuracy than Bayesian conditionalization does, with the advantage going to the Bump Rule. A natural way to understand this result is that breaking away from extremal priors takes time, so that rules which learn more quickly or aggressively are faster than Bayesian conditionalization in learning to pool probability mass on the correct hypothesis. This suggests that even in some classes of difficult problems, the Bump Rule and tempered updating may continue to outperform Bayesian conditionalization in both mean speed and mean accuracy.

With a fine-grained hypothesis space (Table 8) the story is different. Tempered updating

Steps	Bayes	Tempered ($\lambda = 2$)	Bump (c = 0.1)	ERM
5	0.13	0.13	0.09	0.12
10	0.13	0.12	0.07	0.10
25	0.12	0.11	0.06	0.08
50	0.11	0.10	0.05	0.06
100	0.09	0.08	0.03	0.03
250	0.07	0.05	0.01	0.01
500	0.05	0.03	0.00	0.00

Table 7: Mean Brier inaccuracy by learning rule and steps elapsed, averaged over 1,000 trials with extremal priors

scores nearly identically to Bayesian updating. Both the Bump Rule and empirical risk minimization lag noticeably behind. Even after 500 steps, the Bump Rule is not learning: the Bump Rule shows worsening Brier inaccuracy as compared with previous steps.

One natural way to interpret this result is that the winner-takes-all dynamic of empirical risk minimization and the Bump Rule can tend to increase inaccuracy in a fine-grained hypothesis space, because discontinuous explanatory bumps are awarded to the best-fitting hypothesis but not to neighbors which fit existing data nearly as well. With a large number of hypotheses, there is a high risk that explanatory bumps will be awarded to the wrong hypotheses. This risk is minimized by rules such as conditionalization and tempered updating, which use likelihoods to parcel out proportional explanatory credit among many hypotheses.

In a nutshell, the story of this section is that the Bump Rule shows good Brier accuracy in some types of difficult problems and less impressive Brier accuracy in others. Tempered updating does not always match the standout performance of the Bump Rule in some environments, such as learning with extremal priors, but neither does it show the same performance dips in others, such as learning with fine-grained hypothesis spaces.

Steps	Bayes	Tempered ($\lambda = 2$)	Bump (c = 0.1)	ERM
5	0.0097	0.0097	0.0118	0.0191
10	0.0096	0.0096	0.0119	0.0188
25	0.0094	0.0095	0.0110	0.0180
50	0.0092	0.0093	0.0112	0.0177
100	0.0089	0.0090	0.0121	0.0170
250	0.0083	0.0084	0.0132	0.0156
500	0.0079	0.0080	0.0133	0.0143

Table 8: Mean Brier inaccuracy by learning rule and steps elapsed, averaged over 1,000 trials over fine-grained hypotheses

9 Evaluating behavior

What should we make of these results? Where speed is concerned, the results are clear (Section 7). The Bump Rule shows a large speed improvement over Bayesian conditionalization. Tempered updating captures a sizable portion of this speed improvement. These results hold even if we are concerned with mean learning speed, rather than learning speed in specific cases. To the extent that we are concerned with speed, we should therefore see tempered updating as delivering a weakened version of the same benefits that the Bump Rule provides.

At the same time, speed is not everything. We also saw that empirical risk minimization learns more quickly than all other rules under the same conditions. For this reason, it is important to study accuracy as well as speed (Section 8).

In friendly problems, we saw that the Bump Rule shows lower mean accuracy than Bayesian conditionalization does, with the gap widening as we move from Brier scoring to log scoring. A substantial part of this gap is made up by tempered updating, especially under the log rule. This suggests that to the extent we are concerned with mean accuracy in the good case, tempered updating should be seen as a useful improvement over the Bump Rule.

From an externalist, ecological perspective, matters are more complex. With small data samples, both tempered updating and the Bump Rule produce the highest-accuracy results less often than Bayesian conditionalization does. This pattern reverses somewhat with large data samples. However, caution is needed before trading too much on this result, as empirical risk minimization often outperforms both tempered updating and the Bump Rule on this criterion.

In difficult problems, the story is likewise complex. Some ways of increasing problem difficulty, such as extremizing priors, gave a substantial advantage in Brier accuracy to the Bump Rule. This advantage was partly captured by tempered updating. Other ways of increasing problem difficulty, such as fine-graining the hypothesis space, gave a substantial disadvantage in Brier inaccuracy to the Bump Rule. This disadvantage was nearly erased by tempered updating.

All told, these results suggest that tempered updating captures many of the advantages of the Bump Rule, albeit sometimes less strongly. Tempered updating also sometimes avoids the disadvantages of the Bump Rule. This may make tempered updating a productive alternative to the Bump Rule for those attracted to the Bump Rule's performance in simulations, but seeking to avoid the challenges facing the Bump Rule in general (Section 4) or in unfriendly environments (Section 8.2).

10 Implications for incompatibilism

How do these results bear on the prospects for incompatibilism? One way to read this paper is as a novel way of supporting incompatibilism. We saw two ways to understand tempered updating as a principled alternative to Bayesian conditionalization (Section 5), which avoids the challenges raised for the Bump Rule (Section 6) and exhibits similarly attractive performance in simulations (Sections 7-9). We could read this paper as supporting incompatibilism by supporting tempered updating. That is a fine way to read this paper, though it is not the only reading.

I want to conclude by discussing four concerns that might be raised for tempered updating. These concerns will help to identify the audience to which tempered updating will be most attractive.

10.1 Arguments for conditionalization

Many Bayesians offer arguments aiming to show that Bayesian conditionalization is the uniquely rational updating rule. One popular type of argument is the *Dutch book argument* (Armendt 1980; Lewis 1999; Pettigrew 2021) which aims to show that agents who fail to conditionalize are vulnerable to sure loss. Another class of arguments are *accuracy-based*, aiming to show that conditionalization uniquely maximizes expected accuracy (Easwaran 2013; Greaves and Wallace 2006; Pettigrew 2021) or is accuracy-dominant (Briggs and Pettigrew 2020; Nielsen 2021).

To a large extent, these debates will need to be settled elsewhere. Incompatibilist explanationists have offered some responses to both types of arguments for conditionalization (Douven 2022). Others have also responded to the Dutch book (Christensen 1991; Mahtani 2012) and accuracy-based arguments for conditionalization (Schoenfield 2017; Wedgwood 2018).

This paper does contribute a new perspective through the loss-minimization interpretation of tempered updating in Section 5.2. The loss-minimization interpretation takes a page from generalized Bayesian statistics (Bissiri et al. 2016; Grünwald and van Ommen 2017) in reminding us that conditionalization maximizes expected accuracy only once we fix measures of data-fitting, regularization rules, and the relative importance of fitting data against regularization. Even those who think that rationality fixes a single approach to data-fitting and regularization need not be committed to the view that there is a single permissible way to balance them. Under some conditions, it may be permissible for agents to exhibit a risk-seeking emphasis on fitting data over regularizing towards the prior, an emphasis that leads to tempered updating under the loss-minimization interpretation.

10.2 Safe Bayesian updating

I said in Section 5 that many statisticians have studied tempered updating. What I neglected to mention is that statisticians primarily study and recommend cool temperatures $\lambda < 1$. Tempered updating became popular as a way to be more conservative than standard Bayesian conditionalization, a safety-seeking move which prevents overfitting in difficult or mis-specified learning problems (Grünwald and van Ommen 2017; Holmes and Walker 2017; Wu and Martin 2023).

In this perspective, the explanationist's choice of warm temperatures $\lambda > 1$ is, while not unprecedented, an uncommon application of tempered updating. Even if we accept a risk-based argument for the permissibility of varying risk attitudes λ to match the agent's risk attitudes, in many contexts agents may look more fondly on risk-averse rules than on the risk-seeking explanationist rule studied here.

From an ecological perspective, externalist or not, the criticism may be softened. Here, the choice of temperature serves to demarcate the environments in which tempered updating may be rational. The finding is that although tempered updating with $\lambda > 1$ may not be appropriate in environments where risk-seeking should be avoided, it may be an excellent alternative to the Bump Rule in environments where risk-seeking is desired. We may, however, be left with the consequence that inference to the best explanation is less-often appropriate than some advocates may have hoped.

10.3 Hypothesis-dependence

Section 3 introduced the *ambitious Bayesian* project of characterizing the static and dynamic behavior of a probability function over all propositions of interest to an agent. We saw that the Bump Rule does not pursue the ambitious Bayesian project. The Bump Rule tells agents only how to update their views over the sub-algebra generated by a single hypothesis partition. We saw in Section 5 that tempered updating has the same problem. Tempered updating is a hypothesis-dependent update rule.

Not all theorists have the ambitious Bayesian project. Many critics of Bayesian theorizing see its generality not as a feature, but rather as a bug, because even simple learning problems must be solved relative to the agent's global worldview in a way that is often difficult given an agent's limited information (Belot 2013; Efron 2013; Gelman 2008). Moreover, traditional arguments for conditionalization including Dutch Book (Armendt 1980; Lewis 1999; Pettigrew 2021) and accuracy-based arguments (Easwaran 2013; Greaves and Wallace 2006; Pettigrew 2021) are usually formulated in the special case of updating over a hypothesis partition. This restriction may be essential (Gallow 2019; Hild 1998), though some efforts have been made to overcome it (Easwaran and Nielsen 2025; Schultheis forthcoming). Hence, while conditionalization has the advantage of being well-defined over an arbitrary algebra, we may not recover some traditional motivations for conditionalization in this setting.

However, many philosophers do have the ambitious Bayesian project. For these philosophers, both the Bump Rule and tempered updating will be nonstarters. This fact, perhaps more than any other, helps to mark out the differences between Bayesian conditionalization and leading incompatibilist alternatives.

10.4 Differences from the Bump Rule

Some explanationists may be concerned by two differences between tempered updating and the Bump Rule. First, whereas the Bump Rule provides a boost only to the best explanation, tempered updating provides a proportional boost to other competing explanations. If strong emphasis is put on the desire to boost only the *best* explanation, then tempered updating will not be an appropriate gloss of inference to the best explanation.

Second, whereas the Bump Rule evaluates explanatory considerations globally through a bump based on fit to the total evidence, tempered updating evaluates explanatory considerations locally through the likelihoods of each successive item of evidence. This may be interpreted as a way of showing insufficient concern for the global nature of explanation.

Some of the results in this paper suggest that explanationists may wish to be cautious in insisting on rules with these features. We saw in Section 2 that the winner-takes-all nature of the Bump Rule provides a discontinuous boost to some hypotheses above their neighbors. We saw in Section 8 that this boost causes the Bump Rule to struggle in fine-grained hypothesis spaces, where tempered updating is often more accurate because of its ability to award proportional explanatory credit to neighbors. These results may speak in favor of reading explanationism as compatible with proportional explanatory credit.

Likewise, we saw in Section 4 that the global nature of the Bump Rule leads it to violate a range of likelihoodist principles including Minimal Likelihoodism and the Zero-Likelihood Principle. Because the Bump Rule provides a global explanatory boost that is independent of likelihoods, it can update against the direction of current likelihoods, even to the extent of assigning nonzero probability to hypotheses which give zero likelihood to incoming evidence. We saw in Section 6 that tempered updating avoids these challenges by removing the global explanatory bump. These results may speak in favor of reading explanationism as compatible with local rather than global explanatory credit.

10.5 Taking stock

In this paper, we saw that tempered updating provides an attractive alternative to the Bump Rule. It does not resolve all challenges for incompatibilist explanationism, including Dutch Book and accuracy-arguments for conditionalization (Section 10.1), and concerns about risk-seeking (Section 10.2) or hypothesis-dependence (Section 10.3). Nor does it give all explanationists all of what they want (Section 10.4). But tempered updating does avoid the challenges facing the Bump Rule (Section 6) and deliver many of the benefits that attracted incompatibilists to the Bump Rule (Sections 7-9).

At a minimum, I hope that these results serve to reinforce the idea that IBE can be given a rigorous articulation as a well-pedigreed probabilistic update rule that is both distinct from Bayesian updating and perhaps defensible. Further studies of the performance of tempered updating against Bayesian inference may provide a better articulation of the strengths and weaknesses of each, as well as perhaps the situations in which each rule may be more appropriate.

References

Armendt, Brad. 1980. "Is there a Dutch Book argument for probability kinematics?" *Philosophy of Science* 47:583–8.

Belot, Gordon. 2013. "Bayesian orgulity." *Philosophy of Science* 80:483–503.

Bird, Alexander. 2017. "Inference to the best explanation, Bayesiansm, and knowledge." In Kevin McCain and Ted Poston (eds.), *Best explanations: New essays on inference to the best explanation*, 97–120. Oxford University Press.

Bissiri, Pier Giovanni, Holmes, Chris, and Walker, Stephen. 2016. "A general framework for updating belief distributions." *Journal of the Royal Statistical Society Series B: Statistical Methodology* 78:1103–1130.

Briggs, R.A. and Pettigrew, Richard. 2020. "An accuracy-dominance argument for conditionalization." *Noûs* 54:162–81.

Buchak, Lara. 2013. *Risk and rationality*. Oxford University Press.

Carnap, Rudolf. 1950. *Logical foundations of probability*. University of Chicago Press.

Christensen, David. 1991. "Clever bookies and coherent beliefs." *Philosophical Review* 100:229–47.

Dellsén, Finnur. 2018. "The heuristic conception of inference to the best explanation." *Philosophical Studies* 175:1745–66.

—. forthcoming. "Inferring to the best explanation from uncertain evidence." *Philosophy of Science* .

Douven, Igor. 1999. “Inference to the best explanation made coherent.” *Philosophy of Science* 66:S424–35.

—. 2013. “Inference to the best explanation, Dutch books, and inaccuracy minimisation.” *Philosophical Quarterly* 63:428–44.

—. 2020. “The ecological rationality of explanatory reasoning.” *Studies in History and Philosophy of Science Part A* 79:1–14. doi:10.1016/j.shpsa.2019.06.004.

—. 2022. *The art of abduction*. MIT Press.

Douven, Igor and Wenmackers, Sylvia. 2017. “Inference to the best explanation versus Bayes’s rule in a social setting.” *British Journal for the Philosophy of Science* 58:535–70.

Easwaran, Kenny. 2013. “Expected accuracy supports conditionalization—and conglomerability and reflection.” *Philosophy of Science* 80:119–142.

Easwaran, Kenny and Nielsen, Michael. 2025. “Updating by maximizing expected accuracy in non-partitional settings.” *Journal of Philosophical Logic* 54:1095–1134.

Efron, Bradley. 2013. “Bayes’ theorem in the 21st century.” *Science* 340:1177–8.

Fitelson, Branden. 2007. “Likelihoodism, Bayesianism, and relational confirmation.” *Synthese* 156:473–89.

Friel, Nial and Pettit, Anthony. 2008. “Marginal likelihood estimation via power posteriors.” *Journal of the Royal Statistical Society Series B: Statistical Methodology* 70:589–607.

Gallow, J. Dmitri. 2019. “Diachronic Dutch books and evidential import.” *Philosophy and Phenomenological Research* 99:49–80.

Gelman, Andrew. 2008. “Objections to Bayesian statistics.” *Bayesian Analysis* 3:445–50.

Greaves, Hilary and Wallace, David. 2006. “Justifying conditionalization: Conditionalization maximizes expected epistemic utility.” *Mind* 115:607–32.

Grünwald, Peter and van Ommen, Thijs. 2017. "Inconsistency of Bayesian inference for misspecified linear models, and a proposal for repairing it." *Bayesian Analysis* 12:1069–1103.

Hájek, Alan. 2003. "What conditional probability could not be." *Synthese* 137:273–323.

Harman, Gilbert. 1986. *Change in view*. MIT Press.

Heitz, Richard. 2014. "The speed-accuracy tradeoff: History, physiology, methodology, and behavior." *Frontiers in Neuroscience* 8:1–19. doi:10.3389/fnins.2014.00150.

Hild, Matthias. 1998. "The coherence argument against conditionalization." *Synthese* 115:229–58.

Holmes, Chris and Walker, Stephen. 2017. "Assigning a value to a power likelihood in a general Bayesian model." *Biometrika* 104:497–503.

Huemer, Michael. 2009. "Explanationist aid for the theory of inductive logic." *British Journal for the Philosophy of Science* 60:345–75.

Ibrahim, Joseph and Chen, Ming-Hui. 2000. "Power prior distributions for regression models." *Statistical Science* 15:46–60.

Jeffrey, Richard. 1965. *The logic of decision*. University of Chicago Press.

Kitcher, Philip. 1981. "Explanatory unification." *Philosophy of Science* 48:507–31.

Kullback, Solomon and Leibler, Richard. 1951. "On information and sufficiency." *Annals of Mathematical Statistics* 22:79–86.

Lange, Marc. 2014. "Aspects of mathematical explanation: Symmetry, unity and salience." *Philosophical Review* 123:485–531.

—. 2022. "Putting explanation back into "inference to the best explanation"." *Noûs* 56:84–109.

Leitgeb, Hannes and Pettigrew, Richard. 2010a. "An objective justification of Bayesianism I: Measuring inaccuracy." *Philosophy of Science* 77:201–35.

—. 2010b. "An objective justification of Bayesianism II: The consequences of minimizing inaccuracy." *Philosophy of Science* 77:236–72.

Lewis, David. 1999. "Why conditionalize?" In *Papers in metaphysics and epistemology*, 403–7. Cambridge University Press.

Lipton, Peter. 2004. *Inference to the best explanation*. Routledge.

Lombrozo, Tania. 2007. "Simplicity and probability in causal explanation." *Cognitive Psychology* 55:232–57.

Mahtani, Anna. 2012. "Diachronic Dutch book arguments." *Philosophical Review* 121:443–50.

McGrew, Timothy. 2003. "Confirmation, heuristics and explanatory reasoning." *British Journal for the Philosophy of Science* 54:553–67.

Morton, Jennifer. 2017. "Reasoning under scarcity." *Australasian Journal of Philosophy* 95:543–59. doi:10.1080/00048402.2016.1236139.

Musgrave, Alan. 1988. "The ultimate argument for scientific realism." In Robert Nola (ed.), *Relativism and realism in science*, 229–52. Kluwer Academic Publishers.

Nielsen, Michael. 2021. "Accuracy dominance and conditionalization." *Philosophical Studies* 178:3217–36.

Okasha, Samir. 2000. "van Fraassen's critique of inference to the best explanation." *Studies in History and Philosophy of Science Part A* 31:691–710.

Pettigrew, Richard. 2021. "On the pragmatic and epistemic virtues of inference to the best explanation." *Synthese* 199:12407–38.

Popper, Karl. 1959. *The logic of scientific discovery*. Routledge and Kegan Paul.

Poston, Ted. 2014. *Reason and explanation: A defense of explanatory coherentism*. Palgrave Macmillan.

Psillos, Stathis. 2000. "Abduction: Between conceptual richness and computational complexity." In Peter Flach and Antonis Kakas (eds.), *Abduction and induction: Essays on their relation and integration*, 59–74. Springer.

Schmidt, Andreas. 2019. "Getting real on rationality – behavioral science, nudging, and public policy." *Ethics* 129:511–543. doi:10.1086/702970.

Schoenfield, Miriam. 2017. "Conditionalization does not (in general) maximize expected accuracy." *Mind* 126:1155–87.

Schultheis, Ginger. forthcoming. "Accurate updating." *Philosophy of Science* forthcoming.

Sober, Elliott. 2008. *Evidence and evolution: The logic behind the science*. Cambridge University Press.

—. 2015. *Ockham's razors: A user's manual*. Cambridge University Press.

Staffel, Julia. 2020. *Unsettled thoughts: A theory of degrees of rationality*. Oxford University Press.

Steel, Daniel. 2007. "Bayesian confirmation theory and the likelihood principle." *Synthese* 156:53–77.

Sutton, Richard and Barto, Andrew. 2018. *Reinforcement learning: An introduction*. MIT Press, second edition.

Thorstad, David. forthcoming. "Ecological rationality without externalism." *Philosophy of Science* forthcoming.

Trpin, Borut and Pellert, Max. 2019. "Inference to the best explanation in uncertain evidential situations." *British Journal for the Philosophy of Science* 70:977–1001.

van Fraassen, Bas. 1980. *The scientific image*. Oxford University Press.

—. 1989. *Laws and symmetry*. Oxford University Press.

Wedgwood, Ralph. 2018. "Epistemic teleology: Synchronic and diachronic." In Kristoffer Ahlström-Vij and Jeffrey Dunn (eds.), *Epistemic consequentialism*, 85–112. Oxford University Press.

Weisberg, Jonathan. 2009. "Locating IBE in the Bayesian framework." *Synthese* 167:125–43.

Wickelgren, Wayne. 1977. "Speed-accuracy tradeoff in information processing dynamics." *Acta Psychologica* 41:67–85.

Wu, Pei-Shien and Martin, Ryan. 2023. "A comparison of learning rate selection methods in generalized Bayesian inference." *Bayesian Analysis* 18:105–32.