Blackwell-Monotone Updating Rules

Mark Whitmeyer*

October 25, 2024

Abstract

An updating rule specifies how an agent reacts to information. An updating rule is Blackwell monotone if more information is always better for an agent in a decision problem and strictly Blackwell monotone if, in addition, there is always a decision problem in which more information is strictly better for an agent. Bayes' law is strictly Blackwell monotone, and I show that within a broad class of updating rules–those that distort the Bayesian posteriors in a signal-independent manner–it is the only strictly Blackwell-monotone updating rule. Moreover, when the state is non-binary, I show that Bayes' law and the trivial updating rule in which an agent dogmatically holds a single belief are the only continuous Blackwell-monotone updating rules.

*Arizona State University. Email: mark.whitmeyer@gmail.com. I thank the editor, Emir Kamenica, and four anonymous referees for their comments. Faruk Gul, Vasudha Jain, Ehud Lehrer, Benson Tsz Kin Leung, Paula Onuchic, Eddie Schlee, Ludvig Sinander, Shoshana Vasserman, Joseph Whitmeyer, Tom Wiseman, Kun Zhang and a number of seminar and conference audiences provided invaluable feedback. I am also grateful to Tong Wang, whose question inspired this paper. Nour Chalhoub and Konstantin von Beringe provided excellent research assistance. This paper was previously titled "Bayes = Blackwell, Almost."

1 Introduction

The Blackwell order over statistical experiments (Blackwell (1951) and Blackwell (1953)) provides an elegant way of comparing information structures. To a Bayesian expected-utility maximizing agent, if one experiment ranks above another in the Blackwell order, it yields a higher welfare to the agent than the other, no matter the agent's decision problem. We term this property possessed by Bayesian updating–that Blackwell-more-informative experiments yield higher payoffs to an agent no matter her decision problem–*Blackwell monotonicity*. If an agent violates Blackwell monotonicity, she cannot be Bayesian.

This paper asks the converse question: is it true that if an agent is not Bayesian, she must contravene Blackwell monotonicity? One the one hand, many of the leading non-Bayesian updating rules formulated to explain peoples' (non-Bayesian) behavior violate Blackwell monotonicity. For example, the conservative Bayesians of Edwards (1968) may strictly prefer less information to more. Another example is updating rules that exhibit confirmatory bias (Rabin and Schrag (1999)). Divisible updating (Cripps (2022)) also may yield a negative value for information; likewise the $\alpha - \beta$ model of Grether (1980).

On the other hand, a related question does not single out Bayes' law. Take some (at least partially informative) experiment. Is it true that if the agent does not update this information according to Bayes' law, there is always a decision problem in which she prefers to observe nothing than this experiment; *viz.*, for which this experiment has strictly negative value to the agent? As Morris and Shin (1997) and Braghieri (2023) reveal, the answer is no. A broad variety of updating rules are such that some information is always better than none.[1]

For the first half of this paper, we restrict attention to a special class of updating

[1]Notably, one such family of rules corresponds to underreaction, in which the agent's posterior from any piece of information (signal realization) is some average of the Bayesian posterior and the prior. Even these rules are not Blackwell monotone, and we elaborate upon why shortly.

rules–those that systematically distort beliefs, i.e., for which the agent's posterior is the image of some experiment-independent distortion of the Bayesian posterior. We also refine Blackwell monotonicity by introducing a "strict" modifier: an updating rule is strictly Blackwell monotone if it is not only Blackwell monotone but preserves the strict superiority of an experiment over another.

In the main result of this paper, Theorem 3.1, we show that in the class of updating rules that systematically distort beliefs, Bayes' law is the unique strictly Blackwell monotone updating rule. That is, if an agent always prefers more information, and there is always a decision problem in which strictly more information can be exploited, she must be a Bayesian. We then push the limits of this result by moving beyond systematic distortions and find that there are two broader classes of updating rules in which Bayesianism is equivalent to strict Blackwell monotonicity.

An updating rule is *convex* if an agent always (weakly) prefers to observe experiment π with probability λ and π' with its complement–that is, randomize over two experiments–than observe their convex combination, experiment $\lambda\pi + (1-\lambda)\pi'$, for sure. An updating rule is *grounded* if the absence of information (a completely uninformative experiment) is updated correctly. Proposition 4.3 reveals that a grounded, convex updating rule is strictly Blackwell monotone if and only if it is Bayes' law. *Focused* updating rules generalize those that systematically distort beliefs: such rules are those for which identical columns must produce identical posteriors, irrespective of the experiments in which they lie. In Proposition 4.7, we find that a grounded, focused updating rule is strictly Blackwell monotone if and only if it is Bayes' law.

As we discuss in §1.1, there is an asymmetry between updating rules corresponding to overreactions to information by the agent, versus those in which the agent underreacts. In particular, underreaction is a milder mistake: false positives–wrongly abandoning the prior-optimal action–can be very costly; false negatives, less so. Underreaction can only lead to false negatives and so the value of some information (versus none) can never be negative. But Blackwell monotonicity requires something stronger; namely, that the agent prefer any marginal increase in information. The central feature of these environments (with systematic distortions, convex, and focused) is that they all enable us to make underreaction costly in the sense that we can find decision problems in which a local improvement in informativeness leads to a more likely false negative.

The various additional structures we place on the class of updating rules are meaningful. In §4.1, we construct a strictly Blackwell monotone and grounded rule other than Bayes' law. Importantly, however, by our equivalence results, it does not satisfy the other reasonable desiderata of convexity or focus. Thus, one perspective on our findings is that they provide an impossibility result: strict Blackwell-monotonicity and one of these other criteria cannot jointly be satisfied by any updating rule other than Bayes' law. Even within the class of systematic distortions (and even if we require that the distortion be continuous), the "strict" modifier is also important. In particular, when there are just two states, a updating rule that systematically distorts posteriors according to a continuous distortion is Blackwell monotone if and only if it corresponds to extreme-belief aversion (Appendix C).

These results are also useful in that they remove the need to check whether a particular specification of a (systematically distorting) non-Bayesian updating rule is Blackwell monotone. Unless there are two states and it is some form of extreme-belief aversion, any non-Bayesian updating rule that systematically distorts posteriors is such that there exist two Blackwell-ranked experiments with the property that the less-informative of the two is strictly better for the agent. These findings also motivate the construction of orders over experiments for non-Bayesians. How might an underreacter rank experiments, for instance?

In light of the non-exclusiveness of Bayes' law under the criterion introduced by Morris and Shin (1997) and Braghieri (2023), it is surprising that Bayes' law can be singled out by Blackwell monotonicity. Why, in particular, aren't simple updating rules in which an agent underreacts to information Blackwell monotone?

1.1 Why Not Underreact?

Although departures from Bayesian updating cannot strictly benefit a DM, it is known that not all errors are equally harmful. Namely, there is a strong sense in which *underreaction*, where a DM's posterior is a convex combination of the Bayesian posterior and the prior, is superior to *overreaction*, where the Bayesian posterior is a convex combination of the DM's posterior and the prior. As Morris and Shin (1997) and Braghieri (2023) show, a DM who underreacts can never be hurt by information: some information is always better than none.[2] In stark contrast, for an agent who overreacts, there always exists a decision problem in which her losses from information (versus none) are unboundedly large.

It is easy to see why overreaction can be so harmful: the DM can be tricked into taking an extreme action that she should not take, and if incentives are steep enough, this mistake will be very costly. The innocuity of underreaction is more subtle: essentially, the martingality of the Bayesian posteriors means that a decision problem in which a DM makes mistakes at some beliefs must be one in which she benefits at others, and this benefit must outweigh the cost of any mistakes.

Crucially, these papers are making statements about which updating rules can and cannot render a negative value for information. In contrast, here, we study a different question. We are interested in characterizing the updating rules for which the *marginal* value of information is always positive, in which more is always

[2] Morris and Shin (1997) and Braghieri (2023) provide a full characterization of updating rules that lead to a positive value for information (versus none). von Beringe and Whitmeyer (2024) have a simpler and closely related thesis–underreacters can never be exploited by a malevolent principal, overreacters always can–and their main result is implied by one in Braghieri (2023).

preferred to less. That is, the other papers are “some versus none,” ours is “more versus less.” It is also notable that the class of updating rules shown by Morris and Shin (1997) and Braghieri (2023) to render a positive value for information is quite large, whereas requiring a positive marginal value for information (strict Blackwell monotonicity) singles out Bayes’ law.

We now illustrate these ideas in the context of a simple two-state example. The state space, Θ, is binary, $\Theta = \{0,1\}$ and the DM’s prior is $\mu_0 = \mathbb{P}(1) = 1/2$. The DM has a binary decision problem with action set $A = \{a_0, a_1\}$. Action a_0 yields a state-independent payoff of 0 (the DM’s utility is $u(a_0, \theta) = 0$ for all $\theta \in \Theta$), whereas action a_1 yields payoff -30 in state 0 and 20 in state 1 ($u(a_1, 0) = -30$ and $u(a_1, 1) = 20$). The DM is indifferent between the two actions when her belief $\mu \equiv \mathbb{P}(1)$ equals 3/5.

Now take a binary experiment with realizations L and H, where the Bayesian posterior upon observing L is 0 and the Bayesian posterior upon observing H is some $\mu_H > 1/2$. If the DM overreacts, this signal can be worse for her than no information. In particular, suppose $\mu_H < 3/5$, in which case a Bayesian DM will never take action a_1, which yields her a payoff of 0, the same as her payoff to no information. However, if she overreacts and instead updates the high signal realization to some posterior that is strictly larger than 3/5, she will mistakenly take the high action, producing a strictly negative *ex ante* expected payoff.

This is depicted in Figure 1. The black line segment is the DM’s payoff (in her belief μ) to action a_0. The purple line segment is her payoff to a_1. The DM’s prior (1/2) is the large purple circle, and the low and high Bayesian posteriors, $\mu_L = 0$ and $\mu_H = 9/16$, are the small solid green and blue dots, respectively. Her expected payoff at belief μ_H from taking action a_1 is specified by the red x, and is strictly negative. On the other hand, the large hollow blue dot is her non-Bayesian posterior $\hat{\mu}_H = 5/8$. Her expected payoff at this belief corresponds to the blue x, which leads her to incorrectly take action a_1 following realization H.

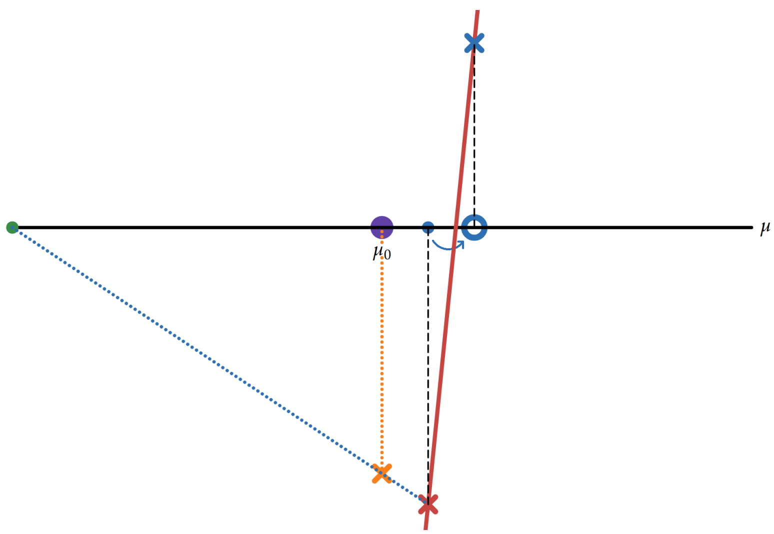

Figure 1: A Pernicious Overreaction

With underreaction, the worst that can happen is a false negative: it may be that the Bayesian posterior following the high signal realization, μ_H, is strictly larger than 3/5, in which case the DM ought to take action a_1. If, instead, the DM's posterior following the high realization is less than 3/5, she will persist with the *status quo*, taking the prior-optimal action of a_0. But this hurts her only in the sense that she is failing to take advantage of an opportunity–she is still no worse off than if she had had observed nothing.

This failure to take advantage of information is what makes underreaction violate Blackwell monotonicity, as we will now see. Suppose the DM underreacts in the sense that her updating rule produces posteriors that are 50/50 averages of the Bayesian posteriors with the prior. Take a signal that produces a distribution over Bayesian posteriors that is supported on $\{0, 5/8, 7/8\}$ (with respective probabilities 1/3, 1/3, and 1/3). The DM's updating produces posterior 1/4 instead of 0, 9/16 instead of 5/8, and 11/16 instead of 7/8. Consequently, the DM takes action a_1 if and only if she observes the highest signal realization. She fails to fully take advantage of information because she should also take a_1 following the intermediate signal realization but does not (as $9/16 < 3/5 < 5/8$).

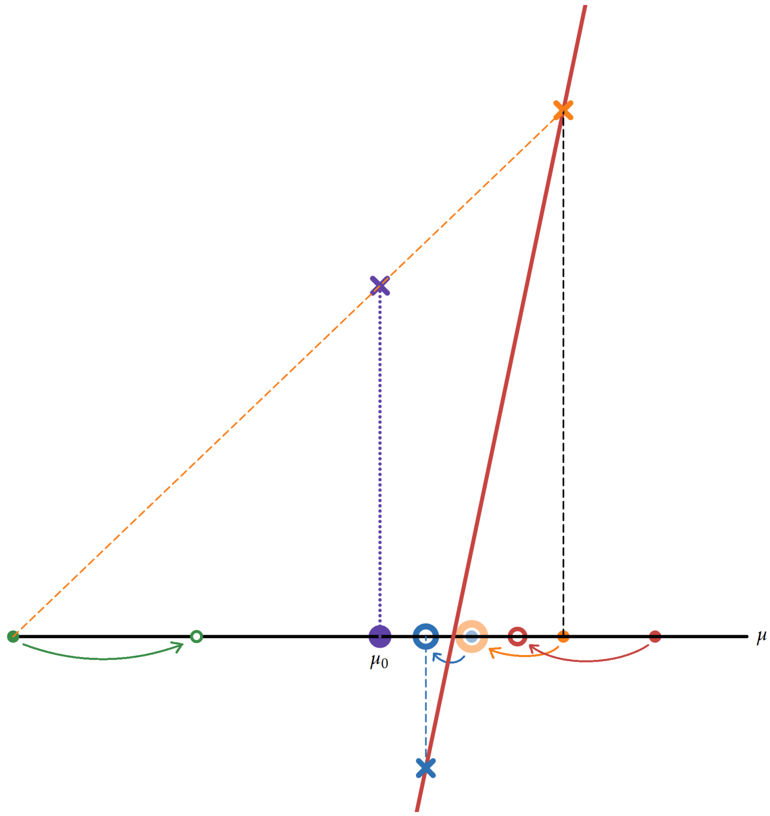

Figure 2: Underreaction Leads to Type-II Errors

The specified ternary distribution corresponds to a strictly more informative experiment than a binary one with support $\{0, 3/4\}$, yet yields a strictly lower *ex ante* expected payoff to the DM. This is because the binary experiment does not produce any consequential mistakes: although the high realization in the binary is underreacted to, the DM's posterior (5/8) is still high enough for her to pick the correct action at the interim point.

This is depicted in Figure 2. As before, the prior is the large purple circle. The solid dots are the (common) low posterior 0, $\mu_M := 5/8$, 3/4, and 7/8; in green, blue, orange, and red, respectively. The underreacting *DM's* posteriors are the hollow dots instead; in green, blue, orange, and red, respectively. The purple *x* is the Bayesian DM's *ex ante* expected payoff from each experiment–the extra information is superfluous. The non-Bayesian DM, however, makes a mistake: her expected payoff following realization M is the strictly negative blue x, so she will mistakenly choose a_0.

1.2 Related Work

By now, many papers explore non-Bayesian updating. One vein of the literature formulates axioms that produce updating rules other than Bayes' law. Epstein (2006) studies an agent whose behavior is in the spirit of the prone-to-temptation agent of Gul and Pesendorfer (2001). Ortoleva (2012) axiomatizes a model in which an agent does not behave as a perfect Bayesian when confronted with unexpected news. Of special note is Jakobsen (2019), who introduces a model of coarse Bayesian updating in which a decision-maker (DM) partitions the belief simplex into a collection of convex sets. Jakobsen (2019) presents an example (Example 4) of a coarse Bayesian updater who, nevertheless, assigns a higher value to more information. This is a particular case of extreme-belief aversion, the unique family of rules that are Blackwell monotone when there are two states (Appendix C). His Proposition 7 states precisely when a regular (for which all cells of the partition have full dimension) coarse Bayesian updating rule is Blackwell monotone.

This work is also related to the work on dynamically consistent beliefs–see, e.g., Gul and Lantto (1990); Machina and Schmeidler (1992); Border and Segal (1994); Siniscalchi (2011); and the survey, Machina (1989). Another seminal paper in that area is Epstein and Le Breton (1993), who show that "dynamically consistent beliefs must be Bayesian," thereby establishing an equivalence (as Bayesian beliefs are obviously dynamically consistent). Accordingly, we find that within our specified classes, dynamic consistency and strict Blackwell monotonicity are also equivalent.

Closely tied to the notion of dynamic consistency is the value of information for DMs with non-expected-utility preferences. That some experiments may be harmful to a DM is illustrated in Wakker (1988), Hilton (1990), Safra and Sulganik

(1995), and Hill (2020). Li and Zhou (2016) show that the Blackwell order holds for almost all DMs with uncertainty-averse preferences provided they can commit *ex ante* to actions, and Çelen (2012) establishes that it holds for an MEU DM.

2 Setup

There is a finite set of states of nature, Θ, with $|\Theta| = n \geq 2$. $\Delta \equiv \Delta(\Theta)$ is the $(n-1)$-simplex, the set of probabilities on Θ, understood as a subset of $\mathbb{R}^{n-1}$, topologized by the Euclidean metric. $\mu_0 \in \Delta^\circ$ denotes our decision-maker's (DM's) full-support prior, where Δ° denotes the (topological) interior of Δ. For any belief $\mu \in \Delta$, we write $\mu(\theta) \equiv \mathbb{P}_\mu(\theta)$.

A (statistical) experiment is a pair (π, S), consisting of a stochastic map $\pi\colon \Theta \to \Delta(S)$ and a finite set of signal realizations S for which i. the unconditional probability of each $s \in S$ is strictly positive; and ii. the Bayesian posterior after observing each s, μ_s, is unique (for $s, s' \in S$, if $s \neq s'$, $\mu_s \neq \mu_{s'}$).[3] We note the formula for the Bayesian posterior: for all $\theta \in \Theta$,

$$\mu_s(\theta;\pi) = \frac{\mu_0(\theta)\pi(s|\theta)}{\sum_{\theta' \in \Theta} \mu_0(\theta')\pi(s|\theta')}.$$

Π denotes the set of experiments. In turn, from the *ex ante* perspective a statistical experiment induces a distribution over Bayesian posteriors $\rho \in \Delta\Delta(\Theta)$, where $\mathbb{P}_\rho(\mu_s) = \rho(\mu_s) = \sum_{\theta' \in \Theta} \mu_0(\theta')\pi(s|\theta')$.

We are interested in studying departures from Bayes' law. To that end, we denote the DM's posterior upon observing signal realization s, given experiment (π, S), by $\hat{\mu}_s(\theta;\pi)$.[4] An Updating Rule, U, is a map

$$U\colon \Pi \to \Delta^S$$
$$(\pi, S) \mapsto (\hat{\mu}_s(\theta;\pi))_{s \in S},$$

[3]Neither assumption affects the results of this paper and both help simplify the analysis.

[4]Throughout, beliefs with hats will be those produced by the (potentially non-Bayesian) updating rule, and those without hats will be those produced by Bayes' law.

i.e., it maps a experiment to a collection of posteriors, one for each signal realization.

Following de Clippel and Zhang (2022), we say that an updating rule Systematically Distorts Beliefs if there exists a function $\varphi\colon \Delta \to \Delta$ such that for all statistical experiments $(\pi, S) \in \Pi$, and signal realizations $s \in S$,

$$\hat{\mu}_s(\theta;\pi) = \varphi(\mu_s(\theta;\pi)).$$

For most of the paper, we restrict attention to updating rules that systematically distort beliefs, which from now on we term Updating Rules. This allows us to drop subscripts and arguments: given distortion φ and a Bayesian posterior μ, the DM's posterior is $\hat{\mu} \equiv \varphi(\mu)$.

A Decision Problem (A, u) consists of a compact set of actions A and a continuous utility function $u\colon A \times \Theta \to \mathbb{R}$. For a Bayesian DM, any decision problem induces a value function in the posterior μ

$$V(\mu) \equiv \max_{a \in A} \mathbb{E}_{\mu} u(a, \theta) \ .$$

V is convex, which implies a positive value of information for a Bayesian DM. In this paper, we are interested in evaluating the value of information for non-Bayesians. For simplicity, we specify that at every $\hat{\mu} \in \Delta$ the DM's choice of action is Consistent: her choice depends only on the realized posterior, i.e., selection $a^*(\hat{\mu}) \in \arg\max_{a \in A} \mathbb{E}_{\hat{\mu}} u(a, \theta)$ is a function of $\hat{\mu}$.[5]

Now we state the two central definitions of the paper.

Definition 2.1. An updating rule is Blackwell Monotone if for any decision problem (A, u) and consistent choice of action $a^*\colon \Delta \to A$, a DM's *ex ante* expected utility from observing experiment (π, S) is higher than from observing (π', S') if $\pi \succeq \pi'$, where $\succeq$ is the (Blackwell) order over experiments.

[5]Naturally, we are also assuming *interim* optimality: a^* is a selection from the arg max correspondence at belief $\hat{\mu}$.

Definition 2.2. An updating rule is Strictly Blackwell Monotone if it is Blackwell monotone and for any two experiments (π, S) and (π', S') for which $\pi \succ \pi'$ there exists a decision problem and consistent choice of action $a^*: \Delta \to A$ such that the DM's *ex ante* expected utility is strictly higher from observing the draw from π than from π'.

For a fixed decision problem and consistent choice of action, the function $W(\mu) := \mathbb{E}_\mu u(a^*(\varphi(\mu)), \theta)$ is a well-defined function of the Bayesian posterior μ. Thus, letting ρ' and ρ be the Bayesian distributions over posteriors corresponding to π' and π, an updating rule is Blackwell monotone if for any decision problem (A, u), consistent $a^*: \Delta \to A$, and pair $\pi \succeq \pi'$,

$$\mathbb{E}_\rho W(\mu) \geq \mathbb{E}_{\rho'} W(\mu).$$

Note that Blackwell monotonicity is is equivalent to W's convexity in μ. An updating rule is strictly Blackwell monotone if (in addition to Blackwell monotonicity), for any pair $\pi \succ \pi'$ there exists a decision problem and a consistent decision rule such that

$$\mathbb{E}_\rho W(\mu) > \mathbb{E}_{\rho'} W(\mu).$$

2.1 Blackwell Monotonicity

We are taking a stance on how to compute the value of information. There are several reasonable ways to evaluate a non-Bayesian DM's value of information. In all cases, the DM's action at any belief must be interim-optimal; *viz.*, must maximize her expected utility given her posterior $\hat{\mu} = \varphi(\mu)$. Thus, from a Bayesian DM's perspective, the DM with non-Bayesian belief $\hat{\mu} \neq \mu$ may make a mistake. But how should we evaluate this expected payoff, with respect to the Bayesian belief μ or to $\hat{\mu}$? Moreover, when we compute the DM's payoff from the *ex ante* perspective, should we use the correct distribution over Bayesian posteriors, or should we allow for forecasting mistakes and, thus, use an alternative (Bayes-plausible) distribution over posteriors?

Ours is the perspective of a sophisticated DM (or observer), who forecasts the distribution over signal realizations–and, thus, Bayesian posteriors–correctly, yet understands that the DM is not a perfect Bayesian but will potentially make mistakes at interim beliefs. She may, therefore, refuse free information, in anticipation of her likely errors. In Appendix B.1, we show that when the DM instead evaluates expected payoffs using $\varphi(\mu)$ (but still forecasts signal realizations correctly), U is Blackwell monotone (in this new sense) if and only if φ is affine.

In the latter case, the DM's expected utility at the Bayesian posterior μ is $V(\varphi(\mu))$, and, *ex ante*, we evaluate $\mathbb{E}_{\rho} V(\varphi(\mu))$. The DM's mistake is made twice, both when the DM chooses her action at belief $\varphi(\mu)$ but also when computing the expectation of the action's payoff. A third case is that in which the forecast distribution over signal realizations is also incorrect–some Bayes-plausible $\hat{\rho}$ instead of ρ–and so we evaluate $\mathbb{E}_{\hat{\rho}} V(\hat{\mu})$. We provide a full characterization of Blackwell monotonicity for this conception of it in Appendix B.2.

2.2 Why Systematic Distortions?

The primary focus of this paper is on updating rules that systematically distort beliefs. There are several justifications for this. First, systematic distortions are those for which counterfactual events do not affect decision-making either directly or indirectly: only the column of the experiment corresponding to the realized signal matters for the DM's posterior and the agent can simplify fractions, which renders the other columns totally irrelevant *ex interim*. This is consistent with our method of evaluating the value of information in that the DM encounters the signal realization and takes an interim-optimal action at her resulting posterior. Why should counterfactual events matter in her calculus?

Second, the bulk of the existing literature on updating rules focuses on those that systematically distort beliefs. Thus, our results concern a particularly relevant and useful class of rules. A third argument in favor of our emphasis on systematic distortions is that the subsequent results concerning focused (and grounded) rules and convex (and grounded) rules also single out Bayes' law as equivalent to strict Blackwell monotonicity. In this light, the systematic-distortion specification can be seen as a slightly simplified environment that is extremely tractable and yet whose properties extend more broadly.

3 Blackwell-Monotone Updating Rules

In this section, we state and prove the main result of the paper. It is important to keep in mind that in this section, by updating rule, we mean "updating rule that systematically distorts beliefs."

Theorem 3.1. *An updating rule is strictly Blackwell monotone if and only if it is Bayes' law.*

That is, an updating rule is strictly Blackwell monotone if and only if $\varphi(\mu) = \mu$ for all $\mu \in \Delta$. The theorem follows from a number of smaller results, the proofs to which may be found in Appendix A. We begin by arguing that if an updating rule is Blackwell monotone, the distortion φ must be continuous on Δ°.

Lemma 3.2. *If updating rule U is Blackwell monotone, φ is continuous on Δ°.*

The intuition behind this lemma is fairly straightforward. U being Blackwell monotone is equivalent to the induced value function W's convexity in the Bayesian posterior W. But convex functions are continuous except on the boundaries of their domain. Thus, if φ is not continuous, and has a "jump" somewhere, we can find a decision problem in which the "jump" is inherited by W, rendering it non-convex.

We then build upon this lemma to discipline the behavior of φ for all full-support beliefs. We say that a distortion φ is Trivial on a set $Y \subseteq \Delta$ if $\varphi(\mu) = \bar{\mu}$ for some $\bar{\mu} \in \Delta$ for all $\mu \in Y$. φ is the Identity map on $Y \subseteq \Delta$ if $\varphi(\mu) = \mu$ for all $\mu \in Y$. We also highlight one particular departure from Bayesian updating:

Definition 3.3. Let there be two states ($n = 2$). A DM with updating rule U displays Extreme-Belief Aversion if there exist two intervals $C_1 := (0, c)$ and $C_2 := (d, 1)$ (with $c \leq d$) such that $\varphi(\mu) = c$ for all $\mu \in C_1$, $\varphi(\mu) = d$ for all $\mu \in C_2$ and $\varphi(\mu) = \mu$ for all $\mu \in [c, d]$.

We discuss this pattern of behavior in Appendix C, where we specialize to the two-state environment.

Proposition 3.4. *(i) Let there be two states ($n = 2$). If U is Blackwell monotone, the DM displays extreme-belief aversion.*

(ii) Let there be three or more states ($n \geq 3$). If U is Blackwell monotone, φ is either trivial or the identity map on Δ°.

Lemma 3.2 is key to establishing this proposition, whose proof is purely topological. Thanks to the lemma, we know that Blackwell monotonicity implies continuity of the distortion φ on the interior of the simplex. Using this continuity, we argue that if φ is not the identity map on Δ° it must be locally trivial at some belief $\mu \in \Delta^\circ$ for which $\varphi(\mu) \neq \mu$. That is, every belief in an open ball around that μ must be mapped to the same point. This local triviality plus the facts that φ is continuous and Δ° is connected allow us to conclude the global structure stated in the proposition.[6] Moreover, if we add the additional requirement that φ is continuous on Δ, we arrive at the following stark conclusion when there are three or more states.

Corollary 3.5. *Let there be three or more states ($n \geq 3$). If U is Blackwell monotone and φ is continuous on Δ, φ is either trivial or the identity map on Δ.*

[6] I am grateful to the anonymous reviewer who suggested this approach.

But U cannot be Blackwell monotone and such that the DM gets every full-support posterior "right," yet gets some non-full-support posterior "wrong:"

Lemma 3.6. *If U is Blackwell monotone and φ is the identity map on Δ°, φ is the identity map on Δ.*

The reasoning behind this lemma is simple. If φ distorts some belief, x, on the boundary of Δ but none on the interior, we can take a binary experiment π whose induced Bayesian distribution over posteriors has one of its two support points on that x. But then we just take an ever-so-slightly less informative experiment, π', in which x is replaced with x' that is on the line segment between x and the prior. No matter how close it is to x, x' lies, nevertheless, in the interior of Δ, and so $\varphi(x') = x'$. Consequently we can always find a decision problem in which x leads to a harmful mistake but x' does not, or a decision problem in which the DM fails to take advantage of her information at x but does at x'.

On the other hand, if φ is trivial on some open ball in Δ°, the updating rule cannot be strictly Blackwell monotone:

Lemma 3.7. *If there exists $\mu' \in \Delta^\circ$ and $\varepsilon > 0$ such that $\varphi(\mu) = \bar{\mu} \in \Delta$ for all $\mu \in B_\varepsilon(\mu')$, U is not strictly Blackwell monotone.*

Finally, we combine the intermediate results to produce the theorem.

Proof of Theorem 3.1. Proposition 3.4 and Lemma 3.7 tell us that if U is strictly Blackwell monotone, φ must be the identity map on Δ°. Lemma 3.6 states that this property must extend to the boundary as well, which yields Theorem 3.1. ■

Even without the additional strictness refinement, if we impose that φ is continuous and nontrivial, Blackwell monotonicity already singles out Bayes' law when there are more than three states (Corollary 3.5). Moreover, the continuity assumption only has bite for extreme posteriors–the distortion is necessarily continuous for full-support posteriors. And the two-state environment? It turns out that a DM's updating rule is Blackwell monotone and continuous if and only if she displays extreme-belief aversion. We discuss this in Appendix C.

3.1 Sketching a Special Setting

To gain intuition for Theorem 3.1, let us sketch an alternative proof for a restricted class of updating rules. We say a distortion, φ, is Grounded if $\varphi(\mu_0) = \mu_0$.

The result to be argued is

Corollary 3.8. *An updating rule whose distortion is grounded is strictly Blackwell monotone if and only if it is Bayes' law.*

Naturally, we need only establish the necessity of Bayes' law. The assumption that φ is grounded is extremely helpful, as it allows us to immediately rule out all but a special class of errors. To elaborate, we note the following result, which follows from discussion in Braghieri (2023). We state a more general version of this remark (when we allow for non-systematic distortions) in the next section.

Remark 3.9. If φ is grounded, U is Blackwell monotone only if $\varphi(\mu) \in \text{conv}\{\mu_0, \mu\}$ for any $\mu \in \Delta$.

The idea is simple: if φ does not pull each Bayesian posterior (weakly) "back" along the line segment toward the prior–if the DM's mistake for *any* binary experiment is anything other than underreaction–there is a binary experiment whose Bayesian distribution over posteriors is supported on at least one belief μ such that we can find a hyperplane that strictly separates μ from $\varphi(\mu)$ and μ_0. Thus, there is a (binary) decision problem where the DM strictly is worse off than if she obtained no information, as upon obtaining $\varphi(\mu)$ she will make a big mistake. Figure 1 illustrates this in the two-state environment.

What this revealed necessity of underreaction to all binary information does is to allow us to draw conclusions about any binary experiment merely by studying binary experiments in the two-state environment. Let us do so. Take the low (strictly less than μ_0) Bayesian posterior, μ_L, produced by a binary experiment, π, (where the other Bayesian posterior is some $\mu_H > \mu_0$) and suppose $\hat{\mu}_L > \mu_L$. Take two additional binary experiments, π' and π'', whose (respective) Bayesian distributions over posteriors are supported on $\left\{\mu_L', \mu_H\right\}$ and $\left\{\mu_L'', \mu_H\right\}$, where $\mu_L < \mu_L' < \mu_L'' < \hat{\mu}_L$.

By construction, $\pi \succ \pi' \succ \pi''$, and so if U is Blackwell monotone, it must be that the monotonicity of the posteriors is weakly preserved by the distortion: $\hat{\mu}_L \leq \hat{\mu}_L' \leq \hat{\mu}_L''$. Otherwise, if $\hat{\mu}_L > \hat{\mu}_L'$ for instance, we could find a binary decision problem for which one action is strictly optimal at μ_L, μ_L', and $\hat{\mu}_L'$ but strictly sub-optimal at $\hat{\mu}_L$, so that the DM does not take advantage of a profitable opportunity when the experiment is π but does with less information (π', violating Blackwell monotonicity.

Slightly more subtle is the fact that if U is Blackwell monotone, $\hat{\mu}_L' = \hat{\mu}_L''$. If that were not the case, $\hat{\mu}_L' < \hat{\mu}_L''$, we could generate a violation of Blackwell monotonicity as in §1.1's example–more specifically, Figure 2. Namely, we take the ternary experiment π_3 whose Bayesian distribution over posteriors is supported on $\left\{\mu_L, \mu_L'', \mu_H\right\}$ and that is strictly more informative than π'. Remember, however, that we know that $\hat{\mu}_L'' > \hat{\mu}_L'$ so there exists a binary decision problem for which the DM fails to take the correct action upon obtaining the interim posterior $\hat{\mu}_L''$ but does at interim posteriors $\hat{\mu}_L'$ and $\hat{\mu}_L$. But in this decision problem, there is no value to a Bayesian DM to the extra information afforded by π over π', and a DM distorting the probabilities in the specified way does not make a mistake when she observes experiment π' but does when she observes π, another violation of Blackwell monotonicity.

So, if U is Blackwell monotone, μ_L' and μ_L'' must be mapped to some common belief by φ. But then so too must any $\mu \in \left[\mu_L', \mu_L''\right]$, in which case there is no difference to the DM between a binary experiment whose corresponding Bayesian distribution over posteriors, π^* has some support point $\mu^* \in \left(\mu_L', \mu_L''\right)$ and another experiment whose Bayesian distribution over posteriors is obtained by a local "spread" of μ^*, i.e., in which the other support point is the same as under π^* and realizing with the same probability, but where μ^* is spread to some collection of posteriors all lying within $\left(\mu_L', \mu_L''\right)$ and whose barycenter is μ^*. Any experiment of the latter form is strictly more informative than π^*, yet for no decision problem will this additional information ever be useful. Consequently, U is not strictly Blackwell monotone. Moreover, although this discussion was couched in the binary-state setting, this logic also applies to any binary distribution in the n-state environment. And, as any $\mu \in \Delta \setminus \{\mu_0\}$ can be produced by a binary experiment, we conclude that if U is strictly Blackwell monotone, $\varphi(\mu) = \mu$ for all $\mu \in \Delta$.

4 Non-Systematic Distortions

Now let us remove the requirement that the updating rule systematically distorts beliefs. That is, given experiment $(\pi, S) \in \Pi$, and upon observing realization $s \in S$, the DM's posterior is

$$\hat{\mu}_s(\theta; \pi) = \varphi_\pi(\mu_s(\theta; \pi)),$$

i.e., the distortion of the Bayesian posterior depends on what the experiment was itself.[7]

In this section, we consider three additional classes of updating rules, which we term grounded, convex, and focused. Our definition of grounded updating rules is a simple modification of the previous term concerning distortions. Namely, we say an updating rule is Grounded if the DM's posterior upon observing the trivial experiment with a single signal realization is the prior. That is, the DM correctly

[7] We do not permit the distortion to depend on the *realization* s. This is without loss of generality as we have assumed that each realization induces a unique Bayesian posterior. Allowing for duplicate Bayesian beliefs and realization-contingent distortions would not affect our results.

updates a completely uninformative signal. We maintain the standing consistency requirement on the DM's interim choice of action. Then, for a fixed statistical experiment (π, S), decision problem, and consistent choice of action, the function $W_\pi(\mu) := \mathbb{E}_\mu u(a^*(\varphi_\pi(\mu)), \theta)$ is a well-defined function of the Bayesian posterior μ. Note that as the distortion of the Bayesian posterior φ_π is experiment-dependent, so too is the *ex ante* Value function W_π.

4.1 A Strictly Monotone Updating Rule Other Than Bayes' Law

Absent the assumption of systematic distortions, groundedness, by itself, is not enough to single out Bayes' law as the unique Blackwell-monotone updating rule, nor even the unique strictly Blackwell-monotone rule. Let us construct an alternative one. Let $\Theta = \{0,1\}$. For any experiment (π, S), with corresponding ρ, define

$$\sigma(\pi) := \mathbb{E}_\rho\left[\frac{(\mu-\mu_0)^2}{\mu_0(1-\mu_0)}\right],$$

$$\underline{\mu}_\pi := \mu_0(1-\sigma(\pi)), \quad \text{and} \quad \bar{\mu}_\pi := \mu_0 + \sigma(\pi)(1-\mu_0).$$

$\sigma \colon \Pi \to [0,1]$ is strictly monotone in the Blackwell order: if $\pi \succ \pi'$, $\sigma(\pi) > \sigma(\pi')$. $\sigma(\pi)$ also ranges from 0 (a fully uninformative experiment) to 1 (full information).

We define

$$\varphi_\pi(\mu) = \begin{cases} \underline{\mu}_\pi, & \text{if} \quad \mu \leq \underline{\mu}_\pi \\ \mu, & \text{if} \quad \underline{\mu}_\pi < \mu < \bar{\mu}_\pi \\ \bar{\mu}_\pi, & \text{if} \quad \mu \geq \bar{\mu}_\pi. \end{cases}$$

That is, φ_π corresponds to experiment-dependent extreme-belief aversion, where the degree to which the DM eschews obtaining extreme beliefs is decreasing in the informativeness of the experiment. Observe that this updating rule is grounded. Moreover,

Remark 4.1. This updating rule is strictly Blackwell monotone.

Please visit Appendix A.5 for a proof of this remark. This rule alters the standard extreme-belief-aversion rule, which is Blackwell monotone (proved in Appendix C)–but not strictly Blackwell monotone–to make it strictly so via the experiment-dependent weight $\sigma(\cdot)$. But can there be strictly Blackwell-monotone rules other than Bayes' law when there are three or more states? After all, we already know that the two-state environment is quite special.

Yes, there can: for example, suppose that for each line that goes through μ_0 the DM updates any experiment whose Bayesian posteriors all lie on that line in the manner defined in Remark 4.1, treating the intersection of the simplex and the specified line as the two-state 1-simplex. Any experiment whose Bayesian posteriors are not all collinear is updated according to Bayes' law. This is a strictly Blackwell-monotone (and grounded) updating rule.

4.2 Convex Updating

For any pair of experiments (π, S) and (π', S'), a Convex Combination of the two experiments is, for some $\lambda \in [0,1]$, experiment $(\pi_\lambda, S \cup S')$, where $\pi_\lambda(s|\theta) = \lambda\pi(s|\theta)$ for all $s \in S$, for all $\theta \in \Theta$ and $\pi_\lambda(s'|\theta) = \lambda\pi(s'|\theta)$ for all $s' \in S'$, for all $\theta \in \Theta$.

Definition 4.2. An Updating Rule is Convex if for any triple of experiments (π, S), (π', S'), and $(\pi_\lambda, S \cup S')$, where $(\pi_\lambda, S \cup S')$ is a convex combination of the first two experiments, and consistent choice of action $a^*\colon \Delta \to A$,

$$\lambda \mathbb{E}_\rho W_\pi(\mu) + (1-\lambda)\mathbb{E}_{\rho'} W_{\pi'}(\mu) \geq \mathbb{E}_{\rho_\lambda} W_{\pi_\lambda}(\mu),$$

where ρ, ρ' and $\rho_\lambda = \lambda\rho + (1-\lambda)\rho'$ are the corresponding distributions over Bayesian posteriors (respectively).

Any updating rule that systematically distorts beliefs is convex. An updating rule that does not systematically distort beliefs, yet is convex, is "no learning without full disclosure," where for any experiment other than the fully informative one, each realization is updated to the prior, and the DM updates the fully-informative experiment according to Bayes law. In fact, this updating rule is both Blackwell monotone and grounded. However, it is not strictly Blackwell monotone.

Proposition 4.3. *A grounded, convex updating rule is strictly Blackwell monotone if and only if it is Bayes' law.*

Only the necessity direction needs proving, as Bayes' law is grounded, strictly Blackwell monotone, and convex. Our approach is to use the following generalization of Remark 3.9 (beyond the environment of systematic distortions) to establish a pair of lemmas, which together imply the proposition. We defer the detailed proofs of these lemmas to Appendices A.6 and A.7.

Remark 4.4. Let binary experiment (π, S) induce the Bayesian distribution over posteriors ρ, supported on $\{\mu_1, \mu_2\}$ $(\mu_1 \neq \mu_2)$. If U is grounded, U is Blackwell monotone only if $\hat{\mu}_2^\pi \in \operatorname{conv}\{\mu_0, \mu_2\}$ and $\hat{\mu}_1^\pi \in \operatorname{conv}\{\mu_1, \mu_0\}$.[8]

In the first lemma, we prove the necessity of Bayesian updating for groundedness, convexity, and strict Blackwell monotonicity within the class of binary experiments.

Lemma 4.5. *If U is grounded, convex, and strictly Blackwell monotone, U is Bayes' law for any binary experiment.*

The reasoning behind this lemma is essentially the same as that in §3.1, although the fact that the DM's distortion of the Bayesian posterior may be experiment-specific complicates matters. In particular, we need the additional convexity assumption to carry out the trick where we use a judiciously-constructed ternary experiment to discipline the errors made for posteriors produced by the specified binary experiments. Eventually, we arrive at a pair of experiments as in the last

[8] We adapt the earlier notation so that $\hat{\mu}^\pi := \varphi_\pi(\mu)$.

paragraph of §3.1, where i. one produces three (collinear) posteriors and the other two; ii. the latter is obtained from the former by collapsing the low two posteriors to their barycenter; and iii. the extra information contained in the ternary distribution is lost by way of the updating rule.

Lemma 4.6. *If U is grounded, convex, and strictly Blackwell monotone, U is Bayes' law for any experiment.*

Compared to the previous lemma, this one is easy to prove. Only the necessity direction warrants a proof and we suppose for the sake of contradiction that the updating rule is not Bayes' law for some non-binary experiment π (with corresponding Bayesian distribution over posteriors ρ). This means that we can find a hyperplane that strictly separates the offending Bayesian belief, μ, from $\hat{\mu}^{\pi}$. This hyperplane partitions the simplex, and we use it to generate a binary Bayesian distribution over posteriors by collapsing the collections of support points of ρ on each of the partition elements to their respective barycenters. Being binary, by the previous lemma, the DM gets this distribution "right." Moreover, by construction, the corresponding binary experiment is strictly better for the DM than π for some decision problem, yet is strictly less informative than π. We conclude that U is not Blackwell monotone.

4.3 Focused Updating

We say an updating rule is Focused if for any pair of experiments (π, S) and (π', S'), if $\pi(s|\theta) = \pi(s'|\theta)$ for all $\theta \in \Theta$ for some pair $s \in S$, $s' \in S'$, then $\hat{\mu}_s(\theta;\pi) = \hat{\mu}_{s'}(\theta;\pi')$.

Proposition 4.7. *A grounded, focused updating rule is strictly Blackwell monotone if and only if it is Bayes' law.*

The proof of this proposition is extremely similar to that of Proposition 4.3. We start by using Remark 4.4 to treat binary experiments as if they are all in the two-state environment. Then, we argue that all binary experiments must be updated correctly. Finally, we do the same collapsing-to-the-respective-barycenters trick to extend the necessity of Bayes' law to arbitrary experiments. For the details, please visit Appendix A.8.

Focused updating rules possess a weaker form of column-independence than those that systematically distort beliefs. The latter class are those for which the DM must be able to understand when two ratios are equivalent, irrespective of the precise numerators and denominators of the fractions. *Viz.*, a systematic distorter understands that 1/2 and 9/18 are equivalent. Focused updating rules are those for which the DM must be able to understand when two *identical fractions* are the same–i.e., that 1 equals 1, 2 = 2, 9 = 9 and 18 = 18–but may not recognize that 1/2 = 9/18. Accordingly focused updating rules are a much broader class, and assume a much lower baseline of cognition.[9] For such updaters, focus merely imposes that upon seeing the exact same thing *ex interim*, a DM's beliefs must be the same.

5 Motivated Reasoning

Up until now, we have fixed the DM's full-support prior, but not the decision problem–so that a Blackwell-monotone updating rule is such that more information is more valuable in every decision problem. What if we instead start with a fixed decision problem? Our search, therefore, is for updating rules that generate a positive marginal value of information *for the particular decision problem under consideration*.

Fix a full-support prior $\mu_0 \in \Delta^\circ$ and a finite-action decision problem with at least two actions in which no action is weakly dominated. Let the set of actions be $A = \{a_1, \dots, a_m\}$ (with $m \in \mathbb{N}$, $m \geq 2$). Recalling the definition of value function

[9] To be fair, it *is* more difficult to recognize that .28/.95 and 7/19 are the same (are they?) than 7/19 and 7/19.

$V := \max_{a\in A} \mathbb{E}_\mu u(a,\theta)$, a Problem-Specific Updating Rule (that systematically distorts beliefs), U_V, is an updating rule for which the DM's posterior is a posterior-separable function of the Bayesian posterior $\varphi_V : \Delta \to \Delta$.

We maintain the specification that the DM's choice at every belief $\mu \in \Delta$ is consistent (and interim-optimal); *viz.*, depends only on the realized posterior.

Definition 5.1. A problem-specific updating rule, U_V, is Blackwell Monotone if for any consistent choice of action $a^* : \Delta \to A$, a DM's *ex ante* expected utility from observing experiment (π, S) is higher than from observing (π', S') if $\pi \succeq \pi'$, where $\succeq$ is the (Blackwell) order over experiments.

A Menu of actions is some nonempty subset $\bar{A} \subseteq A$. Given a menu of actions $\bar{A}$, let $V_{\bar{A}}(\mu) := \max_{a\in\bar{A}} \mathbb{E}_\mu u(a,\theta)$ be the DM's value function in Bayesian-belief μ. We say a problem-specific updating rule generates an effective menu on $Y \subseteq \Delta$ if for any consistent choice, there exists some $\bar{A} \subseteq A$ such that $W = V_{\bar{A}}$ on Y.

Theorem 5.2. *A problem-specific updating rule is Blackwell monotone if it generates an effective menu on* Δ*. It is Blackwell monotone only if it generates an effective menu on* Δ°.

The essence of this theorem is that for a fixed decision problem, updating rules that respect the Blackwell order must be behaviorally equivalent to a DM updating to full-support posteriors correctly but ignoring particular actions except possibly on a measure-zero subset of beliefs. To elaborate, any menu $\bar{A} \subseteq A$ generates a polyhedral subdivision of Δ

$$C^{\bar{A}} = \left\{C_a^{\bar{A}}\right\}_{a\in\bar{A}}.$$

For each $a \in \bar{A}$, $C_a^{\bar{A}}$ is the set of beliefs at which action a is optimal when the DM has access to menu $\bar{A}$. Evidently,

Remark 5.3. A problem-specific updating rule generates an effective menu on Δ° if and only if there exists some $\bar{A} \subseteq A$ such that for all $\mu \in \Delta^\circ$, $\mu \notin C_a^{\bar{A}}$ implies $\varphi_V(\mu) \notin C_a^A$.

As this remark and Theorem 5.2 reveal, many different distortions can produce Blackwell monotonicity. However, the behavior induced by such rules must take a particular form; namely, the almost-total ignorance of some subset of actions.

A Omitted Proofs

A.1 Lemma 3.2 Proof

Proof. We prove the contrapositive: if φ is not continuous on Δ°, there is a decision problem in which the induced W is not convex on Δ°. By Theorem 10.1 in Rockafellar (1970) it suffices to show that W is not continuous on Δ°.

We enumerate the states $\Theta = \{0,1,\dots,n-1\}$ and for $\mu \in \Delta$ let μ^i denote $\mathbb{P}(i)$ $(i = 1,\dots,n-1)$. Suppose φ is not continuous at some $\mu \in \Delta^\circ$. That is, there exists some sequence $\{\mu_n\} \subseteq \Delta^\circ$ such that $\mu_n \to \mu$ but $\varphi(\mu_n) \not\to \varphi(\mu)$. Without loss of generality (as we could just relabel) we assume that $\varphi^1(\mu_n) \not\to \varphi^1(\mu)$.

Now, let us construct a decision problem as follows. The set of actions is $A \equiv [0,1]$, and the DM's utility function is

$$u(a,\theta) = -\mathbf{1}_{\{\theta=1\}}(a-1)^2 - \left(1-\mathbf{1}_{\{\theta=1\}}\right)a^2.$$

Consequently, given belief $\mu \in \Delta$, the DM's interim expected payoff (which guides her decision) is

$$-\mu^1(a-1)^2 - \left(1-\mu^1\right)a^2.$$

By construction, $a^*(\varphi(\mu)) = \varphi^1(\mu)$, so

$$W(\mu) = -\mu^1\left(\varphi^1(\mu)-1\right)^2 - \left(1-\mu^1\right)\varphi^1(\mu)^2 = -\mu^1 + 2\mu^1\varphi^1(\mu) - \varphi^1(\mu)^2$$

Observe that $W(\mu_n) \to W(\mu)$ if and only if

$$\varphi^1(\mu_n)(2\mu_n^1 - \varphi^1(\mu_n)) - \varphi^1(\mu)(2\mu^1 - \varphi^1(\mu)) \to 0,$$

which holds if and only if $\varphi^1(\mu_n) \to \varphi^1(\mu)$ or $\varphi^1(\mu_n) \to 2\mu^1 - \varphi^1(\mu)$.

We have assumed away the first possibility, so if $\varphi^1(\mu_n) \not\to 2\mu^1 - \varphi^1(\mu)$, we are done. Accordingly, let

$$\varphi^1(\mu_n) \to \tilde{\mu}^1 := 2\mu^1 - \varphi^1(\mu) \neq \varphi^1(\mu).$$

Observe that this implies $\varphi^1(\mu) \neq \mu^1 \neq \tilde{\mu}^1$. Because of this, there exist scalars $\alpha, \beta \in \mathbb{R}$ such that

$$\alpha\varphi^1(\mu) + \beta > 0, \quad \alpha\mu^1 + \beta > 0, \quad \text{and} \quad 0 > \alpha\tilde{\mu}^1 + \beta,$$

or there exist scalars $\alpha, \beta \in \mathbb{R}$ such that

$$\alpha\varphi^1(\mu) + \beta > 0, \quad \alpha\mu^1 + \beta < 0, \quad \text{and} \quad 0 > \alpha\tilde{\mu}^1 + \beta.$$

Suppose the first case is feasible. We define a new two-action decision problem with action set $\{a_1, a_2\}$, where the payoff to action a_2 is the state-independent 0 and the expected payoff to action a_1 (in belief μ) is $\alpha\mu^1 + \beta$. Thus, recalling that $W(\mu) := \mathbb{E}_\mu u(a^*(\varphi(\mu)), \theta)$,

$$W(\mu_n) \to 0 < \alpha\mu^1 + \beta = W(\mu),$$

and so W is not continuous at μ. We omit the second case as it mirrors the first. ■

A.2 Proposition 3.4 Proof

I'm grateful to an anonymous referee for sketching this approach.

Proof. Let U be Blackwell monotone. By Lemma 3.2, this implies that φ is continuous on Δ°. Suppose there exists some $x \in \Delta^\circ$ for which $\varphi(x) = y \neq x$.

Claim A.1. *There exists an $\varepsilon > 0$ such that for all $x' \in B_\varepsilon(x)$, $\varphi(x') = y$.*

Proof. Suppose for the sake of contradiction not. Then there exists a sequence $\{x_n\}$ in Δ that converges to x and such that $y_n := \varphi(x_n) \neq y$ for all $n \in \mathbb{N}$.

Consider a two-action decision problem in which the payoff (in belief μ) to action 1 is 0 and the payoff to action 2 is $\alpha\mu - \beta$, where α and β are such that

$\alpha y - \beta = 0$ and $\alpha x - \beta < 0$. We must have i. $\alpha y_n - \beta \geq 0$ for infinitely many members of the sequence $\{y_n\}$; or ii. $\alpha y_n - \beta < 0$ for infinitely many members of the sequence $\{y_n\}$.

i. In the first case, by construction, there is a subsequence $\left\{y_{n_k}\right\}$ such that $\alpha y_{n_k} - \beta \geq 0$ for all y_{n_k}. We impose for all such beliefs, the DM takes action 2 and for belief y, the DM takes action 1. As $x_n \to x$, $x_{n_k} \to x$, so

$$\lim_{n_k \to \infty} W\left(x_{n_k}\right) = \lim_{n_k \to \infty} \alpha x_{n_k} - \beta = \alpha x - \beta < 0 = W(x).$$

ii. In the second case, by construction, there is a subsequence $\left\{y_{n_k}\right\}$ such that $\alpha y_{n_k} - \beta < 0$ for all y_{n_k}, so for all such beliefs, the DM takes action 1. Accordingly, $W(x_{n_k}) = 0$ for all x_{n_k}. We impose for belief y, the DM takes action 2, so

$$\lim_{n_k \to \infty} W\left(x_{n_k}\right) = 0 > \alpha x - \beta = W(x).$$

In both cases, W is discontinuous at x and, therefore, non-convex, contradicting that U is Blackwell monotone. ■

This claim implies that either $\varphi^{-1}(y)$ is open in $(\Delta^\circ, \|\cdot\|_E)$ or is the union of an open set with y itself (in the event where $\varphi(y) = y$ as well). By the continuity of φ, $\varphi^{-1}(y)$ is closed in $(\Delta^\circ, \|\cdot\|_E)$. As Δ° is connected, if there are three or more states $\varphi^{-1}(y) = \Delta^\circ$, i.e., φ is trivial on Δ°. If $n = 2$, either $\varphi^{-1}(y) = (0,1)$, $\varphi^{-1}(y) = (0,a]$, or $\varphi^{-1}(y) = [b,1)$. ■

A.3 Lemma 3.6 Proof

Proof. Suppose for the sake of contraposition that there exists some $x \in \Delta^\circ \setminus \Delta$ for which $\varphi(x) \neq x$. If $\hat{x}$ does not lie on the line segment between x and μ_0, we are done, as there exists a decision problem for which the DM strictly prefers no information to any binary experiment whose distribution over Bayesian posteriors has support on x (see Remark 3.9).

Thus, for some $\lambda \in [0,1)$, $\hat{x} = \lambda x + (1-\lambda)\mu_0$. But then there exists a decision problem for which a DM prefers the binary experiment that induces ρ' with support $\{y, \eta\hat{x} + (1-\eta)x\}$ to that inducing ρ with support $\{y, x\}$. Accordingly, U is not Blackwell monotone. ■

A.4 Lemma 3.7 Proof

Proof. Suppose that φ is locally trivial, i.e., assume the stated condition. Take a binary experiment, π, whose induced Bayesian distribution over posteriors, ρ, is supported on $\{\mu', \mu''\}$ for some $\mu'' \in \Delta$. Then, take another experiment, $\tilde{\pi}$, whose Bayesian distribution over posteriors, $\tilde{\rho}$ has support on $\{\mu_1, \mu_2, \mu''\}$, where $\mu_1, \mu_2 \in B_\varepsilon(\mu')$ and $\mathbb{P}_\rho(\mu'') = \mathbb{P}_{\tilde{\rho}}(\mu'')$. By construction, $\tilde{\pi} \succ \pi$, yet for any decision problem and any consistent choice of action, π and $\tilde{\pi}$ yield the DM the same *ex ante* expected utility, violating strict Blackwell monotonicity. ■

A.5 Remark 4.1 Proof

Proof. Take two experiments π and π' with $\pi \succ \pi'$. For any decision problem and consistent decision rule, and for any $\nu \in \left\{\underline{\mu}_\pi, \bar{\mu}_\pi, \underline{\mu}_{\pi'}, \bar{\mu}_{\pi'}\right\}$, let

$$f_\nu(\mu) = \mathbb{E}_\mu u(a^*(\nu), \theta) = \alpha_\nu \mu + \beta_\nu,$$

for $\alpha_\nu, \beta_\nu \in \mathbb{R}$. For each $\tilde{\pi} \in \{\pi, \pi'\}$, define $\bar{A}_{\tilde{\pi}} \subseteq A$ as

$$\bar{A}_{\tilde{\pi}} := \left\{a \in A\colon \mathbb{E}_\mu u(a, \theta) = V(\mu) \text{ for some } \mu \in \left(\underline{\mu}_{\tilde{\pi}}, \bar{\mu}_{\tilde{\pi}}\right)\right\}.$$

By construction,

$$W_\pi(\mu) = \max\left\{f_{\underline{\mu}_\pi}(\mu), f_{\bar{\mu}_\pi}(\mu), \sup_{a \in \bar{A}_\pi} \mathbb{E}_\mu u(a, \theta)\right\},$$

and

$$W_{\pi'}(\mu) = \max\left\{f_{\underline{\mu}_{\pi'}}(\mu), f_{\bar{\mu}_{\pi'}}(\mu), \sup_{a \in \bar{A}_{\pi'}} \mathbb{E}_\mu u(a, \theta)\right\}.$$

Evidently, W_π, $W_{\pi'}$, and $W_\pi - W_{\pi'}$ are convex, so by Theorem 3.1 in Whitmeyer (2023), $\mathbb{E}_\rho W_\pi(\mu) \geq \mathbb{E}_{\rho'} W_\pi(\mu) \geq \mathbb{E}_{\rho'} W_{\pi'}(\mu)$.

Strictness is easy: the standard quadratic-loss utility and unit-interval action set decision problem yields the strictly convex $V(\mu) = -\mu(1-\mu)$ on $[0,1]$ and, therefore, the conclusion. ■

A.6 Lemma 4.5 Proof

Proof. Let U be grounded. Suppose for the sake of contraposition that U is not Bayes' law for some binary experiment (π, S). Appealing to Remark 4.4, we may specify without loss of generality that the state is binary and that the Bayesian distribution ρ has support $\{0,1\}$ with $\varphi_\pi(0) \equiv \gamma \in (0, \mu_0]$ and $\varphi_\pi(1) \equiv \delta \in [\mu_0, 1]$.

Take five (**five!**)[10] additional experiments (π_1, S_1), (π_2, S_2), (π_3, S_3), (π_4, S_4), and (π_5, S_5), defined as follows. Let ρ^i ($i \in \{1, \ldots, 5\}$) denote the distribution over Bayesian posteriors induced by each π_i. Then, for $0 < 4\eta < \gamma$,

$$\operatorname{supp}\rho = \{0,1\},\ \operatorname{supp}\rho^1 = \{0, 4\eta, 1\},\ \operatorname{supp}\rho^2 = \{2\eta, 1\},\ \operatorname{supp}\rho^3 = \{2\eta, 4\eta, 1\},$$
$$\operatorname{supp}\rho^4 = \{3\eta, 1\},\ \text{and } \operatorname{supp}\rho^5 = \{4\eta, 1\},$$

with

$$p := \mathbb{P}_{\rho^2}(2\eta),\ \mathbb{P}_{\rho^1}(4\eta) = \mathbb{P}_{\rho^1}(0) = \frac{p}{2}, q := \mathbb{P}_{\rho^4}(3\eta), \text{ and } \mathbb{P}_{\rho^3}(4\eta) = \mathbb{P}_{\rho^3}(2\eta) = \frac{q}{2}.$$

Note that

$$p = \frac{1-\mu_0}{1-2\eta}$$

by Bayes-plausibility. By construction,

$$\pi \succ \pi_1 \succ \pi_2 \succ \pi_3 \succ \pi_4 \succ \pi_5.$$

Claim A.2. *If U is convex and Blackwell monotone,*

$$\delta \geq \varphi_{\pi_1}(1) = \varphi_{\pi_2}(1) = \varphi_{\pi_3}(1) = \varphi_{\pi_4}(1) = \varphi_{\pi_5}(1) \geq \mu_0.$$

[10] This quantity seems excessive, but it seems like we need all five.

Proof. Throughout the proof of this claim, we can ignore the other support points (the ones less than μ_0) of each ρ^i, as Remark 4.4 tells us that each is mapped by the respective distortions to beliefs less than μ_0 if U is Blackwell monotone and grounded (which we specified at the start of the proof of the lemma). Moreover, by Remark 4.4, $\varphi_{\pi_i}(1) \geq \mu_0$ for each i is necessary for Blackwell monotonicity, so we also assume this.

Next, we argue that

$$\delta \geq \varphi_{\pi_1}(1) \geq \varphi_{\pi_2}(1) \geq \varphi_{\pi_3}(1) \geq \varphi_{\pi_4}(1) \geq \varphi_{\pi_5}(1)$$

if U is Blackwell monotone. Suppose for the sake of contraposition that $\varphi_{\pi_4}(1) < \varphi_{\pi_5}(1)$. Take a binary decision problem in which action a_2 yields a state-independent payoff of 0 and action a_1 yields a payoff, in belief μ, of

$$\mu - \frac{\varphi_{\pi_4}(1) + \varphi_{\pi_5}(1)}{2}.$$

The DM's *ex ante* expected payoff from π_4 is 0 and is strictly positive from π_5, so U is not Blackwell monotone. The analogous arguments produce the rest of the chain, pair by pair.

Finally, we argue that $\varphi_{\pi_1}(1) \leq \varphi_{\pi_5}(1)$ if U is convex and Blackwell monotone. Suppose for the sake of contraposition not, i.e., that $\varphi_{\pi_1}(1) > \varphi_{\pi_5}(1)$. By construction,

$$\rho^1 = \frac{1}{2(1-2\eta)}\rho + \frac{1-4\eta}{2(1-2\eta)}\rho^5.$$

Take a binary decision problem in which action a_2 yields a state-independent payoff of 0 and action a_1 yields a payoff, in belief μ, of

$$\mu - \frac{\varphi_{\pi_1}(1) + \varphi_{\pi_5}(1)}{2}.$$

The DM's *ex ante* expected payoff under π_5 is $v_5 := 0$. The DM's *ex ante* expected payoff under π is

$$v := \mu_0\left(1 - \frac{\varphi_{\pi_1}(1) + \varphi_{\pi_5}(1)}{2}\right),$$

and it is

$$v_1 := \frac{\mu_0 - 2\eta}{1-2\eta}\left(1 - \frac{\varphi_{\pi_1}(1) + \varphi_{\pi_5}(1)}{2}\right)$$

under π_1.

Then,

$$v_1 - \frac{1}{2(1-2\eta)}v - \frac{1-4\eta}{2(1-2\eta)}v_5 = \left(\frac{\mu_0 - 2\eta}{1-2\eta} - \frac{\mu_0}{2(1-2\eta)}\right)\left(1 - \frac{\varphi_{\pi_2}(1) + \varphi_{\pi_3}(1)}{2}\right) > 0,$$

so U is not convex. ■

Claim A.3. *If U is Blackwell monotone, for all $\xi \in \{0, 4\eta\}$ and $\nu \in \{2\eta, 4\eta\}$,*

$$\gamma \le \varphi_{\pi_1}(\xi) \le \varphi_{\pi_2}(2\eta) \le \varphi_{\pi_3}(\nu) \le \varphi_{\pi_4}(3\eta) \le \varphi_{\pi_5}(4\eta) \le \mu_0.$$

Proof. First, we argue that

$$\gamma \le \varphi_{\pi_2}(2\eta) \le \varphi_{\pi_4}(3\eta) \le \varphi_{\pi_5}(4\eta) \le \mu_0$$

is necessary for Blackwell monotonicity. By Remark 4.4, $\varphi_{\pi_2}(2\eta) \in [2\eta, \mu_0]$, $\varphi_{\pi_4}(3\eta) \in [3\eta, \mu_0]$ and $\varphi_{\pi_5}(4\eta) \in [4\eta, \mu_0]$ are necessary. Now suppose for the sake of contraposition that $\varphi_{\pi_2}(2\eta) \in [2\eta, \gamma)$.

Take a binary decision problem in which action a_2 yields a state-independent payoff of 0 and action a_1 a payoff (in belief μ) of

$$-\left(\mu - \frac{\varphi_{\pi_2}(2\eta) + \gamma}{2}\right).$$

Under π, the DM gets 0, as she'll always take action a_2; whereas, under π_2, the DM gets a strictly positive expected payoff, a violation of Blackwell monotonicity. The analogous arguments pair-by-pair yields the necessity of the rest of the chain of inequalities.

Second, we argue that if U is Blackwell monotone, $\max\left\{\varphi_{\pi_1}(0), \varphi_{\pi_1}(4\eta)\right\} \le \varphi_{\pi_2}(2\eta)$ and $\max\left\{\varphi_{\pi_3}(2\eta), \varphi_{\pi_3}(4\eta)\right\} \le \varphi_{\pi_4}(3\eta)$. Let us tackle the first of these. Suppose for the sake of contraposition that $\varphi_{\pi_1}(4\eta) > \varphi_{\pi_2}(2\eta)$. Take a binary decision problem in which action a_2 yields a state-independent payoff of 0 and action a_1 a payoff (in belief μ) of

$$-\left(\mu - \frac{\varphi_{\pi_2}(2\eta) + \varphi_{\pi_1}(4\eta)}{2}\right).$$

The DM's expected payoff under π_1 is bounded above by

$$v_1 := \frac{p}{2}\left(\frac{\varphi_{\pi_2}(2\eta) + \varphi_{\pi_1}(4\eta)}{2}\right);$$

whereas her expected payoff under π_2 is

$$v_2 := -p\left(2\eta - \frac{\varphi_{\pi_2}(2\eta) + \varphi_{\pi_1}(4\eta)}{2}\right).$$

Then, $v_2 - v_1$ equals

$$\frac{p}{2}\left(\frac{\varphi_{\pi_2}(2\eta) + \varphi_{\pi_1}(4\eta)}{2} - 4\eta\right) > \frac{p}{2}\left(\varphi_{\pi_2}(2\eta) - 4\eta\right) > 0,$$

a violation of Blackwell monotonicity. The analogous argument yields the necessity of $\varphi_{\pi_1}(0) \le \varphi_{\pi_2}(2\eta)$. Likewise, the proof can repeated to argue $\max\left\{\varphi_{\pi_3}(2\eta), \varphi_{\pi_3}(4\eta)\right\} \le \varphi_{\pi_4}(3\eta)$.

Third, we argue that if U is Blackwell monotone, $\gamma \le \min\left\{\varphi_{\pi_1}(0), \varphi_{\pi_1}(4\eta)\right\}$ and $\varphi_{\pi_2}(2\eta) \le \min\left\{\varphi_{\pi_3}(2\eta), \varphi_{\pi_3}(4\eta)\right\}$. Let us tackle the first of these. Suppose for the sake of contraposition that $\varphi_{\pi_1}(0) < \gamma$. Take a binary decision problem in which action a_2 yields a state-independent payoff of 0 and action a_1 a payoff (in belief μ) of

$$-\left(\mu - \frac{\max\left\{\varphi_{\pi_1}(0), 4\eta\right\} + \gamma}{2}\right).$$

The DM's *ex ante* expected payoff under π_1 is strictly positive, whereas it is 0 under π, a violation of Blackwell monotonicity. Note that we have assumed without loss of generality (or else U would not be Blackwell monotone) that $\varphi_{\pi_1}(4\eta) \in [0, \mu_0]$.

Now suppose for the sake of contraposition that $\varphi_{\pi_1}(4\eta) < \gamma \le \varphi_{\pi_1}(0)$. If $\varphi_{\pi_1}(4\eta) < 4\eta$, U is not Blackwell monotone so we specify $\varphi_{\pi_1}(4\eta) \ge 4\eta$. Take a binary decision problem in which action a_2 yields a state-independent payoff of 0 and action a_1 a payoff (in belief μ) of

$$-\left(\mu - \frac{\varphi_{\pi_1}(4\eta) + \gamma}{2}\right).$$

The DM's *ex ante* expected payoff under π_1 is strictly positive, whereas it is 0 under π, a violation of Blackwell monotonicity.

This proof can be repeated to argue $\varphi_{\pi_2}(2\eta) \leq \min\left\{\varphi_{\pi_3}(2\eta), \varphi_{\pi_3}(4\eta)\right\}$. ■

Claim A.4. *If U is convex and Blackwell monotone,*

$$\varphi_{\pi_2}(2\eta) = \varphi_{\pi_3}(2\eta) = \varphi_{\pi_3}(4\eta) = \varphi_{\pi_4}(3\eta) = \varphi_{\pi_5}(4\eta).$$

Proof. From Claim A.3, it suffices to show that $\max\left\{\varphi_{\pi_1}(0), \varphi_{\pi_1}(4\eta)\right\} \geq \varphi_{\pi_5}(4\eta)$. Suppose for the sake of contraposition not, i.e., that $\max\left\{\varphi_{\pi_1}(0), \varphi_{\pi_1}(4\eta)\right\} < \varphi_{\pi_5}(4\eta)$. Take a binary decision problem in which action a_2 yields a state-independent payoff of 0 and action a_1 a payoff (in belief μ) of

$$-\left(\mu - \frac{\max\left\{\varphi_{\pi_1}(0), \varphi_{\pi_1}(4\eta)\right\} + \varphi_{\pi_5}(4\eta)}{2}\right).$$

The DM's *ex ante* expected payoff under π_5 is $v_5 := 0$. Her *ex ante* expected payoff under π is

$$v := (1-\mu_0)\frac{\max\left\{\varphi_{\pi_1}(0), \varphi_{\pi_1}(4\eta)\right\} + \varphi_{\pi_5}(4\eta)}{2},$$

and it is

$$v_1 := -\frac{1-\mu_0}{1-2\eta}\left(2\eta - \frac{\max\left\{\varphi_{\pi_1}(0), \varphi_{\pi_1}(4\eta)\right\} + \varphi_{\pi_5}(4\eta)}{2}\right) = \frac{1}{1-2\eta}v - 2\eta\frac{1-\mu_0}{1-2\eta},$$

under π_1. Then,

$$v_1 - \frac{1}{2(1-2\eta)}v - \frac{1-4\eta}{2(1-2\eta)}v_5 = \frac{1}{2(1-2\eta)}\left(v - 4\eta(1-\mu_0)\right) > 0,$$

so U is not convex. ■

Claim A.5. *If U is convex, U is not strictly Blackwell monotone.*

Proof. By Claims A.2 and A.4, if U is convex and Blackwell monotone, $\varphi_{\pi_4}(1) = \varphi_{\pi_3}(1)$ and $\varphi_{\pi_3}(2\eta) = \varphi_{\pi_3}(4\eta) = \varphi_{\pi_4}(3\eta)$. By construction $\pi_3 \succ \pi_4$, yet for any decision problem and consistent choice of action, π_3 and π_4 yield the DM the same *ex ante* expected payoff, so U is not strictly Blackwell monotone. ■

This concludes the proof of the lemma. ■

A.7 Lemma 4.6 Proof

Proof. Let U be grounded and convex. Take an arbitrary non-binary experiment (π, S) and let ρ be the induced Bayesian distribution over posteriors. Suppose for the sake of contraposition that there exists some $x \in \operatorname{supp}\rho$ for which $\varphi_\pi(x) \equiv \hat{x} \neq x$. Consequently, there exists $\alpha \in \mathbb{R}^{n-1}$ and $\beta \in \mathbb{R}$ such that

$$\alpha x + \beta > 0 > \alpha \hat{x} + \beta.$$

Without loss of generality we assume $\varphi_\pi(\mu) = \mu$ for all $\mu \in \operatorname{supp}\rho \setminus \{x\}$. We define

$$M := \{\mu \in \operatorname{supp}\rho \colon \alpha\mu + \beta \geq 0\}.$$

If $\operatorname{supp}\rho \setminus M = \emptyset$, we are done, as Remark 4.4 tell us that U is not Blackwell monotone.

Suppose that $\operatorname{supp}\rho \setminus M \neq \emptyset$, and take the binary Bayesian distribution over posteriors obtained by collapsing the posteriors in M and the posteriors in $\operatorname{supp}\rho \setminus M$ to their respective barycenters (under ρ). Call this Bayesian distribution over posteriors ρ' and denote the experiment that induces it π'. Naturally $\pi \succ \pi'$. Moreover, if $\varphi_{\pi'}(\mu) \neq \mu$ for any of the two posteriors in $\operatorname{supp}\rho'$, Lemma 4.5 tells us that U is not strictly Blackwell monotone.

Suppose, therefore, that $\varphi_{\pi'}(\mu) = \mu$ for all $\mu \in \operatorname{supp}\rho'$. But then the DM must strictly prefer π' to π, as she makes an error with strictly positive probability under π. We conclude that U is not Blackwell monotone. ■

A.8 Proposition 4.7 Proof

The proposition is the product of Lemmas A.6 and A.11. As ever, we need only to establish the necessity of Bayes' law.

Lemma A.6. *If U is grounded, focused, and strictly Blackwell monotone, U is Bayes' law for any binary experiment.*

Proof. We begin just as in the proof of Lemma 4.5 (see Appendix A.6). Let U be grounded and suppose for the sake of contraposition that U is not Bayes' law for some binary experiment (π, S). As before, Remark 4.4 enables us to specify without loss of generality that the state is binary and that the Bayesian distribution ρ has support $\{0,1\}$ with $\varphi_\pi(0) \equiv \gamma \in (0,\mu_0]$ and $\varphi_\pi(1) \equiv \delta \in [\mu_0, 1]$.

We take the five additional experiments constructed in Lemma 4.5's proof and add a sixth to the mix; (π_6, S_6), where π_6 is binary and obtained by combining the columns of π_1 that produce Bayesian posteriors 0 and 1 (adding the two elements in each row of the two columns together). Consequently, the Bayesian distribution over posteriors corresponding to π_6 has support $\{4\eta, \tau\}$, where $\mathbb{P}_{\rho^1}(4\eta) = \mathbb{P}_{\rho^6}(4\eta)$. Moreover, if U is focused, $\varphi_{\pi_6}(4\eta) = \varphi_{\pi_1}(4\eta)$. Note also that

$$\pi \succ \pi_1 \succ \pi_2 \succ \pi_3 \succ \pi_4 \succ \pi_5 \succ \pi_6.$$

Next,

Claim A.7. *If U is Blackwell monotone, for all $\xi \in \{0, 4\eta\}$ and $\nu \in \{2\eta, 4\eta\}$,*

$$\gamma \leq \varphi_{\pi_1}(\xi) \leq \varphi_{\pi_2}(2\eta) \leq \varphi_{\pi_3}(\nu) \leq \varphi_{\pi_4}(3\eta) \leq \varphi_{\pi_5}(4\eta) \leq \varphi_{\pi_6}(4\eta) \leq \mu_0.$$

Proof. Claim A.3 states the entire chain other than $\varphi_{\pi_5}(4\eta) \leq \varphi_{\pi_6}(4\eta)$, but the logic for that is identical, so we can conclude it as well. ■

Claim A.8. *If U is focused and Blackwell monotone,*

$$\varphi_{\pi_2}(2\eta) = \varphi_{\pi_3}(2\eta) = \varphi_{\pi_3}(4\eta) = \varphi_{\pi_4}(3\eta) = \varphi_{\pi_5}(4\eta) = \varphi_{\pi_6}(4\eta).$$

Proof. From Claim A.7, it suffices to show that $\varphi_{\pi_1}(4\eta) = \varphi_{\pi_6}(4\eta)$. We have already observed that if this is not true then U is not focused, so we are done. ∎

Claim A.9. *If U is focused, $\varphi_{\pi_3}(1) = \varphi_{\pi_4}(1)$.*

Proof. This follows immediately from the fact that π_4 is obtained from π_3 by combining two columns, leaving one (the one that yields posterior 1) untouched. ∎

Claim A.10. *If U is focused, U is not strictly Blackwell monotone.*

Proof. This proof mimics that of Claim A.5. By Claims A.8 and A.9, if U is focused and Blackwell monotone, $\varphi_{\pi_4}(1) = \varphi_{\pi_3}(1)$ and $\varphi_{\pi_3}(2\eta) = \varphi_{\pi_3}(4\eta) = \varphi_{\pi_4}(3\eta)$. $\pi_3 \succ \pi_4$, yet for any decision problem and consistent choice of action, π_3 and π_4 produce the same *ex ante* expected payoff, so U is not strictly Blackwell monotone. ∎

We conclude the lemma. ∎

Lemma A.11. *If U is focused, convex, and strictly Blackwell monotone, U is Bayes' law for any experiment.*

Proof. The proof is identical to Lemma 4.6's proof and so we leave it out. ∎

A.9 Theorem 5.2 Proof

Proof. (⇐) The convexity of $V_{\bar{A}}$ gives us Blackwell monotonicity.

(⇒) If U_V is Blackwell monotone, W is convex on Δ. This implies that on Δ°,

$$W(\mu) = \sup_{j \in J} \left\{\alpha_j \mu + \beta_j\right\},$$

for some index set J, where $\alpha_j \in \mathbb{R}^{n-1}$ and $\beta_j \in \mathbb{R}$ for all $j \in J$. For each $a \in A$, let $f_a(\mu) := \alpha_a \mu + \beta_a$ denote its payoff in belief μ (where $\alpha_a \in \mathbb{R}^{n-1}$ and $\beta_a \in \mathbb{R}$ for all $a \in A$). As the set of beliefs at which $f_a(\mu) \neq f_{a'}(\mu)$ for any distinct pair $a, a' \in A$ is dense in Δ°,

$$W(\mu) = \sup_{j \in J} \left\{\alpha_j \mu + \beta_j\right\} = \max_{a \in \bar{A}} f_a(\mu) = V_{\bar{A}}(\mu),$$

for some $\bar{A} \subseteq A$ on a dense subset of Δ°, implying that $W = V_{\bar{A}}$ on Δ°. ■

B Alternative Notions of Blackwell Monotonicity

B.1 When Do Two Wrongs Make a Right?

One other way of thinking about the DM's value for information is to evaluate the decision made *ex interim* according to belief $\varphi(\mu)$ rather than μ. In this case, more information is more valuable in a Blackwell sense if and only if $V(\varphi(\mu))$ is convex, as the *ex ante* expected payoff is $\mathbb{E}_\rho V(\varphi(\mu))$.

We say that map $\varphi\colon \Delta \to \Delta$ is affine if $\varphi(x) = Ax + b$ for some $(n-1)\times(n-1)$ matrix A and $b \in \mathbb{R}^{n-1}$.

Lemma B.1. *$V(\varphi(x))$ is convex for all convex V if and only if φ is affine.*

Proof. ($\Rightarrow$) If $\varphi(x) = Ax + b$, then for all $x, x' \in \Delta$ and $\lambda \in (0,1)$

$$\begin{aligned} V(\varphi(\lambda x + (1-\lambda)x')) &= V(\lambda(Ax+b) + (1-\lambda)(Ax'+b)) \\ &\leq \lambda V(Ax+b) + (1-\lambda)V(Ax'+b) \\ &= \lambda V(\varphi(x)) + (1-\lambda)V(\varphi(x'))\ , \end{aligned}$$

so $V \circ \varphi$ is convex.

($\Leftarrow$) Suppose for the sake of contraposition that there exist distinct $x, x' \in \Delta$ and $\lambda \in (0,1)$ such that $\varphi(x) = Ax + b$ and $\varphi(x') = Ax' + b$ but

$$\varphi(\lambda x + (1-\lambda)x') \neq A(\lambda x + (1-\lambda)x') + b. \tag{B.1}$$

Let $V(x) = \alpha x$, where $\alpha \in \mathbb{R}^{n-1}$, so

$$\begin{aligned} \lambda V(\varphi(x)) + (1-\lambda)V(\varphi(x')) &= \lambda\alpha\varphi(x) + (1-\lambda)\alpha\varphi(x') \\ &= \lambda\alpha Ax + (1-\lambda)\alpha Ax' + \alpha b \qquad (B.2) \\ &= \alpha A(\lambda x + (1-\lambda)x') + \alpha b\ . \end{aligned}$$

Appealing to Expression B.1, WLOG we assume

$$\alpha\left(A\left(\lambda x+(1-\lambda)x'\right)+b\right)\neq\alpha\varphi\left(\lambda x+(1-\lambda)x'\right)$$

(as otherwise we could just modify α). If

$$\alpha\left(A\left(\lambda x+(1-\lambda)x'\right)+b\right)<\alpha\varphi\left(\lambda x+(1-\lambda)x'\right),$$

we have, from Equation B.2,

$$\begin{aligned}\lambda V\left(\varphi\left(x\right)\right)+(1-\lambda)V\left(\varphi\left(x'\right)\right)&=\alpha\left(A\left(\lambda x+(1-\lambda)x'\right)+b\right)\\&<\alpha\varphi\left(\lambda x+(1-\lambda)x'\right)=V\left(\varphi\left(\lambda x+(1-\lambda)x'\right)\right).\end{aligned}$$

so $V\circ\varphi$ is not convex. If

$$\alpha\left(A\left(\lambda x+(1-\lambda)x'\right)+b\right)>\alpha\varphi\left(\lambda x+(1-\lambda)x'\right),$$

we simply define $V(x)=-\alpha x$, in which case, again, we have

$$\lambda V\left(\varphi\left(x\right)\right)+(1-\lambda)V\left(\varphi\left(x'\right)\right)<V\left(\varphi\left(\lambda x+(1-\lambda)x'\right)\right),$$

so $V\circ\varphi$ is not convex. ■

B.2 Forecasting Errors

A second alternative way of evaluating the value of information is one in which we allow for forecast errors from the *ex ante* perspective. In this case, we specify that the anticipated distribution over posteriors is $\hat{\rho}$ and so the *ex ante* expected payoff is $\mathbb{E}_{\hat{\rho}}V(\hat{\mu})$. We also assume that the DM understands the martingality of beliefs, so that the anticipated distribution over posteriors, $\hat{\rho}$, has mean μ_0.

Consequently, an updating rule is Blackwell monotone (in this sense) if $\pi\succeq\pi'$ implies $\mathbb{E}_{\hat{\rho}}\mathbb{E}_{\hat{\mu}}u(a,\theta)\geq\mathbb{E}_{\hat{\rho}'}\mathbb{E}_{\hat{\mu}}u(a,\theta)$ for any decision problem. Then, we have the following easy result:

Remark B.2. An updating rule is Blackwell monotone if and only if $\pi\succeq\pi'$ implies $\hat{\rho}$ is a mean-preserving spread of $\hat{\rho}'$.

C Two States and Extreme-Belief Aversion

Let there be two states ($n = 2$). Recall that a DM with updating rule U displays extreme-belief aversion if there exist two intervals $C_1 := (0, c)$ and $C_2 := (d, 1)$ (with $0 \le c \le d \le 1$) such that $\varphi(\mu) = c$ for all $\mu \in C_1$, $\varphi(\mu) = d$ for all $\mu \in C_2$ and $\varphi(\mu) = \mu$ for all $\mu \in [c, d]$.

Proposition 3.4 tells us that extreme-belief aversion is implied by Blackwell monotonicity. Now let us show that it is sufficient when φ is continuous.

Proposition C.1. *Let $n = 2$. An updating rule U that systematically distorts beliefs according to a continuous distortion φ is Blackwell monotone if and only if the DM displays extreme-belief aversion.*

Proof. Suppose U displays extreme-belief aversion. Because φ is continuous, $\varphi(0) = c$ and $\varphi(d) = 1$. Define $\bar{A} \subseteq A$ as

$$\bar{A} := \left\{a \in A \colon \mathbb{E}_\mu u(a, \theta) = V(\mu) \text{ for some } \mu \in (c, d)\right\}.$$

Given the DM's consistent decision rule, let $f_c(\mu) := \alpha_c \mu + \beta_c$ and $f_d(\mu) := \alpha_d \mu + \beta_d$ be the expected payoffs (in belief μ) to the actions the DM takes at beliefs c and d, respectively, where $\alpha_i, \beta_i \in \mathbb{R}$ for $i \in \{c, d\}$. By interim optimality and the definition of $\bar{A}$, $f_c(\mu) \ge \sup_{a \in \bar{A}} \mathbb{E}_\mu u(a, \theta)$ for all $\mu \le c$. Likewise, $f_d(\mu) \ge \sup_{a \in \bar{A}} \mathbb{E}_\mu u(a, \theta)$ for all $\mu \ge d$. Thus, $W(\mu) = \max\left\{f_c(\mu), f_d(\mu), \sup_{a \in \bar{A}} \mathbb{E}_\mu u(a, \theta)\right\}$, which is convex. ∎

Here is an example of such a rule.

Example C.2. There are two states, $\Theta = \{0, 1\}$, and the set of actions is the unit interval, $A = [0, 1]$. The DM's utility function is the standard "quadratic loss" utility, translated up by .3 (to make the graph prettier): $u(a, \theta) = -(a - \theta)^2 + .3$. Accordingly, $V(x) = -x(1 - x) + .3$. Here is an Interactive Link, where one can adjust the parameters, c and d, that specify the continuous extreme-belief averse rule by moving the corresponding sliders.

References

David Blackwell. Comparison of experiments. In *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*, pages 93–102, Berkeley, Calif., 1951. University of California Press.

David Blackwell. Equivalent comparisons of experiments. *The Annals of Mathematical Statistics*, 24(2):265–272, 1953.

Kim C. Border and Uzi Segal. Dynamic consistency implies approximately expected utility preferences. *Journal of Economic Theory*, 63(2):170–188, 1994.

Luca Braghieri. Biased decoding and the foundations of communication. *Mimeo*, 2023.

Boğaçhan Çelen. Informativeness of experiments for meu. *Journal of Mathematical Economics*, 48(6):404–406, 2012.

Martin Cripps. Divisible updating. *Mimeo*, 2022.

Geoffroy de Clippel and Xu Zhang. Non-bayesian persuasion. *Journal of Political Economy*, 130(10):2594–2642, 2022.

Ward Edwards. Conservatism in human information processing. In *B. Kleinmuntz (Ed.)*, *Formal Representation of Human Judgment*, pages 17–52. Wiley, New York, 1968.

Larry G. Epstein. An axiomatic model of non-bayesian updating. *The Review of Economic Studies*, 73(2):413–436, 04 2006.

Larry G. Epstein and Michel Le Breton. Dynamically consistent beliefs must be bayesian. *Journal of Economic Theory*, 61(1):1–22, 1993.

David M. Grether. Bayes rule as a descriptive model: The representativeness heuristic. *The Quarterly Journal of Economics*, 95(3):537–557, 1980.

Faruk Gul and Outi Lantto. Betweenness satisfying preferences and dynamic choice. *Journal of Economic Theory*, 52(1):162–177, 1990.

Faruk Gul and Wolfgang Pesendorfer. Temptation and self-control. *Econometrica*, 69(6):1403–1435, 2001.

Brian Hill. Dynamic consistency and ambiguity: A reappraisal. *Games and Economic Behavior*, 120:289–310, 2020.

Ronald W. Hilton. Failure of blackwell's theorem under machina's generalization of expected-utility analysis without the independence axiom. *Journal of Economic Behavior & Organization*, 13(2):233–244, 1990.

Alexander M. Jakobsen. Coarse bayesian updating. *Mimeo*, August 2019.

Jian Li and Junjie Zhou. Blackwell's informativeness ranking with uncertainty-averse preferences. *Games and Economic Behavior*, 96:18–29, 2016.

Mark J. Machina. Dynamic consistency and non-expected utility models of choice under uncertainty. *Journal of Economic Literature*, 27(4):1622–1668, 1989.

Mark J. Machina and David Schmeidler. A more robust definition of subjective probability. *Econometrica*, 60(4):745–780, 1992.

Stephen Morris and Hyun Song Shin. The rationality and efficacy of decisions under uncertainty and the value of an experiment. *Economic Theory*, 9:309–324, 1997.

Pietro Ortoleva. Modeling the change of paradigm: Non-bayesian reactions to unexpected news. *American Economic Review*, 102(6):2410–36, May 2012.

Matthew Rabin and Joel L. Schrag. First impressions matter: A model of confirmatory bias. *The Quarterly Journal of Economics*, 114(1):37–82, 1999.

R. Tyrrell Rockafellar. *Convex Analysis*, volume 28 of *Princeton Mathematical Series*. Princeton University Press, Princeton, NJ, 1970.

Zvi Safra and Eyal Sulganik. On the nonexistence of blackwell's theorem-type results with general preference relations. *Journal of Risk and Uncertainty*, 10: 187–201, 1995.

Marciano Siniscalchi. Dynamic choice under ambiguity. *Theoretical Economics*, 6 (3):379–421, 2011.

Konstantin von Beringe and Mark Whitmeyer. The perils of overreaction. *Mimeo*, 2024.

Peter Wakker. Nonexpected utility as aversion of information. *Journal of Behavioral Decision Making*, 1:169–175, 1988.

Mark Whitmeyer. Making information more valuable. *Mimeo*, 2023.