Bayes = Blackwell, Almost

Mark Whitmeyer*

July 20, 2023

Abstract

There are updating rules other than Bayes' law that render the value of information positive. I find all of them.

*Arizona State University. Email: mark.whitmeyer@gmail.com. I thank 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 for their comments. I am also grateful to Tong Wang, whose question inspired this paper. Nour Chalhoub and Konstantin von Beringe provided excellent research assistance.

1 Introduction

When economists write models, they typically endow agents with utility functions that are increasing functions of the agents' wealths. This is justified by appealing to free disposal: more wealth should always be weakly preferred to less, as an agent could always burn any perceived excess. The notion of "more-is-better" also seems reasonable with respect to information in a decision problem. For a Bayesian decision-maker this is, indeed, the case; and, as with wealth, a free-disposal argument provides the justification.

Just as a utility function can be written that is not an increasing function of the agent's wealth; however, so too an updating rule–a rule for how to react to new information–need not satisfy the condition that more information is preferred to less. With the conservative Bayesianism of Edwards (1968), agents may strictly prefer less information. 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).

What we discover in this paper is that Bayes' law is the unique (nontrivial, continuously distorting) updating rule that satisfies the desideratum of more information being preferred to less. To an expected-utility maximizer faced with a decision problem, information is valuable. It is known that to such an agent, Bayes' law is the optimal way to react to new information, that is, the updating rule that maximizes the decision-maker's *ex ante* expected payoff. We show here that the rationale for Bayes' law is even stronger. This paper, thus, provides additional justification for imposing Bayesianism in models.

To go into specifics, the exceptional updating rules noted above are examples of updating rules that (following de Clippel and Zhang (2022)) **systematically distort beliefs**, for which there exists a distortion function from the correct Bayesian posterior to that produced by the updating rule. Restricting attention to such rules, we show that if an updating rule is such that any experiment is more valuable to a decision maker than any garbling (i.e., respects the Blackwell order) when there are three or more states, corresponds to a continuous distortion function, and is non-trivial (does not map every Bayesian belief to the same belief), the updating rule is Bayes' law. That is, for three or more states, Bayes' law is the unique non-trivial updating rule obtained by a continuous distortion that respects the Blackwell order.

On the other hand, a continuous, nontrivial updating rule that respects the Blackwell order has more freedom when there are only two states. Such updating rules must divide the interior of the belief space into at most three regions. In the central region, the updating rule is Bayes' law, whereas on the outer two regions, the updating rule is a coarse rule that maps all beliefs in the region to the same belief. However, if we further impose that the distortion function is differentiable, Bayes' law is the lone survivor. All in all, Bayes' law is the unique non-trivial updating rule obtained by a continuously differentiable distortion that respects the Blackwell order.

1.1 Sketch of Approach

In establishing these results, it is helpful to introduce two categories of error: expansive, where the updating rule produces beliefs outside of the convex hull of the correct (Bayesian) posterior distribution's convex; and contractive, where the updating rule produces beliefs within the convex hull Bayesian posterior distribution's support. From there, we start with an arbitrary posterior for which the updating rule produces an expansive (contractive) error. For any distribution with such a belief in its (convexly-independent) support, we show that any belief that is a strict convex combination of the initial error posterior and some collection of the other support points that is interpreted correctly by the updating rule must yield a violation of the Blackwell order.

Next, we show how the remainder of the specified face of the simplex can be "filled in." We show that any point in the relative interior of a face containing an error must be in the relative interior of the convex hull of the support of some distribution for which the updating rule has an error. We proceed by contraposition; assuming not, then illustrating how we can then find a decision problem for which the value of information is negative. Beyond this, we need merely take care of a few loose ends, as well as show that the updating rules our approach failed to eliminate do, in fact, respect the Blackwell order, closing the gap between necessity and sufficiency.

1.2 Related Work

By now there are many papers that 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. Every Bayesian belief, then, is understood by the DM as the representative belief corresponding to the partition element in which the (correct) Bayesian belief lies.

Jakobsen (2019) presents an example (Example 4) of a coarse Bayesian updater who, nevertheless, assigns a higher value to more information. We show that this is a particular case of what we term occasionally coarse updating rules, the unique family of rules that respect the Blackwell order when there are two states (Theorem 5.1). His Proposition 7 states precisely when a regular (for which all cells of the partition have full dimension) coarse Bayesian updating rule respects the Blackwell order. One implication of our main theorem (Theorem 5.1) is that Bayes' law and the regular coarse Bayesian updating rule with a single element (the entire simplex) are the only (regular) coarse updating rules that respect the Blackwell order when there are three or more states.

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). Naturally, although dynamically consistent beliefs must mean that the DM's updating rule respects the Blackwell order, our Theorem 5.1 demonstrates that beliefs that respect the Blackwell order need not be dynamically consistent.

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 $|\Theta| = n$. $\Delta \equiv \Delta(\Theta)$ is the $(n-1)$-simplex, the set of probabilities on Θ. Let $\mu \in \operatorname{int}\Delta$ denote our decision-maker's (DM's) full-support prior. A statistical experiment, or signal, is a map $\pi\colon \Theta \to \Delta(S)$, where S is a finite set of signal realizations. Denote the set of all experiments with finite support Π.

$\Delta^2 \equiv \Delta(\Delta(\Theta))$ denotes the set of distributions over posterior probabilities (posteriors) $x \in \Delta$. An **Updating Rule**, U is a map

$$\begin{aligned} U\colon \Delta \times \Pi &\to \Delta^2 \\ (\mu, \pi) &\mapsto \rho_U \end{aligned},$$

where $\rho \in \Delta^2$ is a distribution over posteriors whose support is a subset of Δ. One notable updating rule is the Bayesian updating rule, U_B, which produces the Bayesian distribution over posteriors, ρ_B, i.e., $(\mu, \pi) \mapsto_{U_B} \rho_B$.

Corresponding to an updating rule, U, is a mapping from the Bayesian distribution over posteriors to the distribution over posteriors produced by the updating rule, Φ.

$$\begin{aligned} \Phi\colon \Delta^2 &\to \Delta^2 \\ \rho_B &\mapsto \rho \end{aligned}.$$

We define this mapping to be such that the following diagram commutes:

$$\begin{tikzcd} \Delta \times \Pi \arrow[r, "U_B"] \arrow[dr, "U"'] & \Delta^2 \arrow[d, "\Phi"] \\ & \Delta^2 \end{tikzcd}$$

Following de Clippel and Zhang (2022), we say that an updating rule U **Systematically Distorts Beliefs** if its corresponding Φ is such that there exists a well-defined function

$\varphi\colon \Delta \to \Delta$, where

$$\rho_B \mapsto_\Phi \rho$$
$$\{x_1, \ldots, x_k\} \mapsto \{\varphi(x_1), \ldots, \varphi(x_k)\},$$

(assuming that each x_i is distinct) and $\mathbb{P}_{\rho_B}\left(x_j\right) = \mathbb{P}_\rho\left(\varphi\left(x_j\right)\right)$ for all j. Throughout, we will restrict attention to updating rules that systematically distort beliefs–which from now we term **Updating Rules**–which allows us to conduct our analysis within Δ. An updating rule, U, **Respects the Blackwell Order** if for any prior $\mu \in \Delta$, compact action set A, and continuous utility function $u\colon A \times \Theta \to \mathbb{R}$, a decision maker's *ex ante* expected utility from observing experiment π is higher than from observing π' if $\pi \succeq \pi'$, where $\succeq$ is the (Blackwell) partial order over experiments. π' is said to be a garbling of π.

We note the following result. Letting experiment π correspond to the (Bayesian) distribution over posteriors with affinely independent support ρ_B and experiment π' correspond to the distribution over posteriors with affinely independent support ρ'_B,

Theorem 2.1 (Wu (2022))**.** *$\pi \succeq \pi'$ if and only if* $\operatorname{conv}\operatorname{supp}\rho'_B \subseteq \operatorname{conv}\operatorname{supp}\rho_B$.

Given a set of actions A, we define the value function to be

$$V(x) := \max_{a \in A} \mathbb{E}_{x(\theta)} u(a, \theta) \ ,$$

where $x = x(\theta) \in \Delta$ is the DM's belief. Value function V is continuous and convex. Moreover, when A is finite it is piecewise affine, and its graph is a polyhedral surface in $\mathbb{R}^n$. Associated with any such value function is the projection of V onto Δ, which yields a finite collection C of polytopes C_i ($i = 1, \ldots, m$, where $|A| = m$). By construction, action a_i is optimal for any belief $x \in C_i$ and uniquely optimal for any belief $x \in \operatorname{int} C_i$. C is a mathematical object known as a **Power Diagram**.

2.1 Additional Preliminaries

Restricting attention to updating rules that systematically distort beliefs allows us to not only place our analysis squarely in belief space but also focus for the most part on experiments that produce Bayesian distributions over posteriors with affinely independent

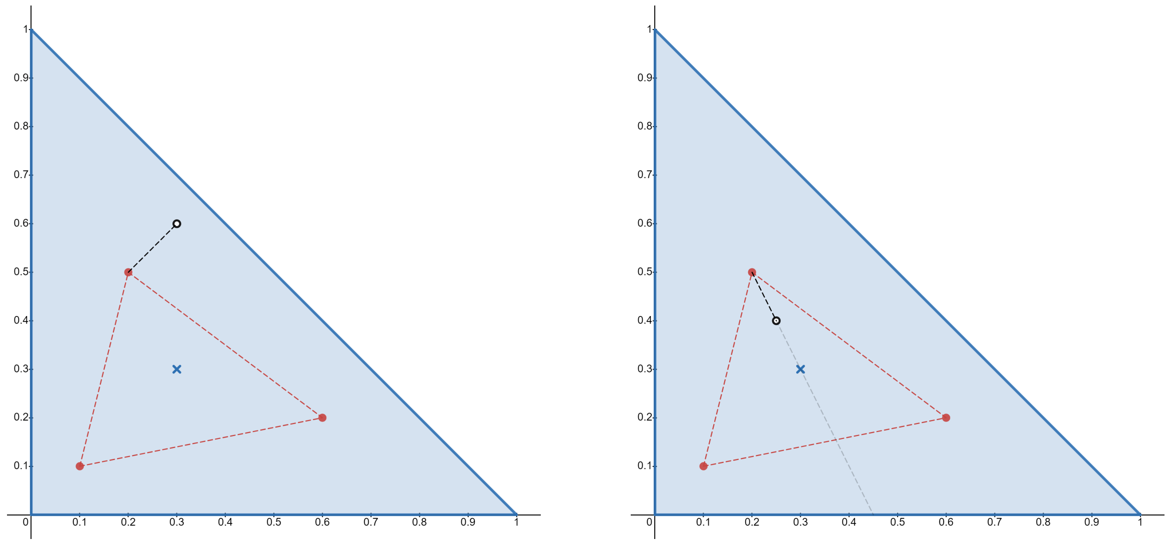

(a) An expansive error

(b) A contractive error

Figure 1: The two varieties of error. $\operatorname{supp}\rho_B$ the red dots; μ, the blue x; $\hat{x}_1$, hollow black.

support. Given such a ρ_B, there are two broad classes of errors that an updating rule can produce.

We say that an error is **Expansive** for ρ_B if for some ρ_B with support on k affinely independent points of support $\{x_1, \ldots, x_k\}$ ($k \leq n$), there exists an $x_i \in \operatorname{supp}\rho_B$ such that $\varphi(x_i) \notin \operatorname{conv}\operatorname{supp}\rho_B$. Alternatively, we say that an error is **Expansive** for some (correct) $x \in \Delta$ if $\varphi(x)$ does not lie on the line segment between x and μ. An updating rule that is such that there is an expansive error for some ρ_B or some x **Produces an Expansive Error**.

In contrast, we say that an error is **Contractive** for ρ_B if for some ρ_B with support on k affinely independent points of support $\{x_1, \ldots, x_k\}$ ($k \leq n$), there exists an $x_i \in \operatorname{supp}\rho_B$ such that $\varphi(x_i) \in \operatorname{conv}\operatorname{supp}\rho_B$ and $x_i \neq \varphi(x_i)$. Alternatively, we say that an error is **Contractive** for some (correct) $x \in \Delta$ if $\varphi(x)$ lies on the line segment between x and μ (and $\varphi(x) \neq x$). An updating rule that is such that there is a contractive error for some ρ_B or some x **Produces a Contractive Error**. An updating rule may produce both expansive and contractive errors. Figure 1 illustrates the two types of error.

We say that an updating rule is **Trivial** on a subset $S \subseteq \Delta$ if it is such that φ is a constant function on S, i.e., for all $x \in S$, $\varphi(x) = x^*$ for some $x^* \in \Delta$. That is, a trivial updating rule on S is one that maps every correct Bayesian posterior, $x \in S$, to a common posterior, x^*.

Let $\ell(x,y)$ denote the line segment between x and y for any $x,y \in \Delta$ and $\ell^\circ(x,y)$ denote the line segment between x and y for any $x,y \in \Delta$, not including its endpoints:

$$\ell(x,y) := \{x' \in \Delta | \exists \lambda \in [0,1] : \lambda x + (1-\lambda) y = x'\},$$

and

$$\ell^\circ(x,y) := \ell(x,y) / \{x,y\}.$$

One special case is when the endpoints of the line segment are μ and some $y \in \Delta$. We denote this $\ell_y \equiv \ell(\mu,y)$. Denote the set of m-faces of the $(n-1)$-simplex by $\mathcal{S}_m$ $(0 \le m \le n-1)$. For each $S_m \in \mathcal{S}_m$ let $\mathcal{S}_t(S_m)$ denote the set of t-faces of Δ for which S_m is also a face, and $\mathcal{S}(S_m) = \cup_{t=m}^{n-1} \mathcal{S}_t(S_m)$. Formally, for $t \ge m$,

$$\mathcal{S}_m = \{S_m | S_m \text{ is an } m\text{-face of } \Delta\}, \quad \text{and} \quad \mathcal{S}_t(S_m) = \{S_t | S_m \text{ is an } m\text{-face of } S_t\}.$$

Note that we understand the $n-1$-face of a simplex to be the simplex itself, and we denote by E the set of all vertices in a simplex, i.e., $E = \mathcal{S}_0$.

3 Expansive Errors

In this section, we focus on updating rules that produce expansive errors. Define the sets

$$\hat{\mathcal{E}} := \{e_i \in E | \ \varphi(e_i) \neq e_i\},$$

the (possibly empty) set of vertices for which U produces an error; and let

$$\hat{\mathcal{S}}_m := \{S_m \in \mathcal{S}_m | \ \exists \ x \in S_m : \ \varphi(x) \neq x\},$$

$$\hat{S} := \cup_m \hat{\mathcal{S}}_m, \quad \hat{\mathcal{S}}^1 := \cup_{m \ge 1} \hat{\mathcal{S}}_m \quad \text{and} \quad \hat{\mathcal{S}}^2 := \cup_{m \ge 2} \hat{\mathcal{S}}_m$$

be the (possibly empty) set of faces of Δ, faces of Δ other than vertices, and faces of Δ other than vertices or edges, respectively, containing beliefs for which U produces an error.

We say an updating rule is **Occasionally Stubborn** if

1. $S_m \in \hat{\mathcal{S}}$ implies $S_t \in \hat{\mathcal{S}}$ for all $S_t \in \mathcal{S}(S_m)$;
2. For all $S_m \in \hat{\mathcal{S}}^2$, for all m, there exists an $x^* \in \Delta$ such that $\varphi(x) = x^*$ for all $x \in \operatorname{int} S_m$;
3. If $x^* \in \operatorname{int} S_1^{'}$ for some $S_1^{'} \in \hat{\mathcal{S}}_1$ then either $\varphi(x) = x^*$ for all $x \in \operatorname{int} S_1^{'}$ or there exists a vertex of $S_1^{'}$, $e_i^{'}$ such that $\varphi(x) = x^*$ for all $x \in \ell^\circ(e_i{'}, x^*)$ and $\varphi(x) = x$ for all $x \in S_1{'} \setminus \ell^\circ(e_i{'}, x^*)$. Moreover, for all $S_1 \in \hat{\mathcal{S}}^1 \setminus \left\{S_1^{'}\right\}$, $\varphi(x) = x^*$ for all $x \in \operatorname{int} S_1$; and
4. $\varphi(e_i) \in \ell(x^*, e_i)$ for all $e_i \in \hat{\mathcal{E}}$.

That is, occasionally stubborn updating rules must either get "every belief correct" on the relative interior of a (non-vertex, non-edge) face; or must be such that every belief on the relative interior of that face, and any other face with the original face as one of its own faces must be updated to the same belief, x^* (1 and 2). Moreover, this incorrect, stubborn, belief is unique: all incorrectly updated beliefs other than the vertices and possibly a subset of beliefs on one edge must be updated to x^* (2). If a vertex is updated incorrectly, there is more freedom: it can be updated to any belief that is "more extreme" with respect to that vertex than x^* (4). Finally, there may be a special case in which the updating rule produces an error for beliefs on an edge and the incorrect belief x^* also lies on that edge (3). In this event, either every belief on the interior of that edge is mapped to the incorrect belief, or only the portion of beliefs between x^* and a vertex are, with the remainder updated correctly (3).

The central result of this section is the following theorem.

Theorem 3.1. *When there are three or more states, any updating rule that respects the Blackwell order and produces an expansive error must be occasionally stubborn.*

A corollary of this is

Corollary 3.2. *When there are three or more states, if an updating rule respects the Blackwell order and produces an expansive error, it must be trivial on the interior of the probability simplex,* $\operatorname{int} \Delta$.

In proving Theorem 3.1, we start with the following proposition. We specify that ρ_B is a distribution over posteriors with affinely-independent support on k points: $\operatorname{supp}\rho_B = \{x_1,\ldots,x_k\}$.

Proposition 3.3. *If updating rule U is expansive for some Bayesian posterior $x_i \in \operatorname{supp}\rho_B$ and respects the Blackwell order, it is expansive for all Bayesian posteriors x for which $x = \lambda \cdot \operatorname{supp}\rho_B$, where $\sum_j^k \lambda_j = 1$, $\lambda_j \in [0,1]$ for all $j = 1,\ldots,k$ and $\lambda_i > 0$. Moreover, for all such x, $\varphi(x) = x^*$.*

Another way to put this proposition is as follows: if updating rule U is expansive for a point x_i in support of a Bayesian distribution over posteriors, ρ_B, and respects the Blackwell order, then the updating rule must be expansive for any point within the convex hull of the support of ρ_B, other than those obtained by taking convex combinations of points other than x_i. Moreover, these points must all be mapped to the same x^* by φ.

Proof. We sketch the proof here, leaving the details to Appendix A.1.

Step 1 ("Bringing the Error Point In"): We start with some belief x_1, for which U has an expansive error, and some ρ_B (with affinely independent support) of which x_1 is a support point. We construct two new distributions ρ_B' and ρ_B'', by keeping all of the support points of ρ_B the same except for the 1st. That support point, x_1', for ρ_B', instead, lies within the convex hull of ρ_B's support (and is not x_1), i.e., it is "brought in." Accordingly, by construction ρ_B' is a strict MPC of ρ_B. We do likewise with ρ_B'', "bringing the 1st point in support of ρ_B' in" *in the same direction on which x_1 was brought in*, i.e., x_1, x_1' and x_1'' are collinear.

Step 2 ("Banishing the New Points"): Next, we argue (in Claim A.1) that if U respects the Blackwell order, it must also produce an expansive error for the new points x_1' and x_1''. In fact, $\varphi\left(x_1'\right) \equiv \hat{x}_1'$ and $\varphi\left(x_1''\right) \equiv \hat{x}_1''$ must lie outside of the convex hull of ρ_B's support. We argue by contraposition: if one did not, we could find a decision problem for which the value of information is strictly negative.

Step 3 ("If Distinct then More Contracted $\Rightarrow$ More Extreme"): Our third step is to argue (in Claim A.2) that if U respects the Blackwell order and $\hat{x}_1' \neq \hat{x}_1''$, $\hat{x}_1''$ must be "more extreme" than $\hat{x}_1'$ in the sense that it must lie outside of the convex hull of $\operatorname{supp}\rho_B' \cup \left\{\hat{x}_1'\right\}$. Again we argue by contraposition: we construct a decision problem in which a strictly less informative experiment is strictly better.

Figure 2 illustrates steps 1 through 3. In 2a, x_1 is in red, x_1' in purple, x_1'' green, and the prior is the blue cross. $\hat{x}_1$ is the hollow black point. 2b depicts how if $\hat{x}_1'$ (hollow purple) lies within ρ_B's support, we can find a value function (whose power diagram consists of the blue and black shaded regions) for which the DM strictly prefers ρ_B' to ρ_B. Similarly, 2c shows how $\hat{x}_1''$ (hollow green) must be "more extreme" than $\hat{x}_1'$.

Step 4 ("All Meet the Same End"): Next, we argue (in Claim A.3) that, in fact, if U respects the Blackwell order, $\hat{x}_1'$ and $\hat{x}_1''$ are not distinct. They are same point $\hat{x}_1^*$. Contraposition is again our approach, but now the proof is a little more subtle: we compare a convex combination of ρ_B and ρ_B'' with its strict MPC, ρ_B', and show that unless the destinations of x_1' and x_1'' are the same, we can find a decision problem for which ρ_B' is strictly preferred.

Step 5 ("Repeating Steps 1-4"): We now repeat the first four steps, with the modification that now we "bring in" some support point of ρ_B, x_t, other than x_1, i.e., with respect to which U may not produce an expansive error. As in Step 1, we construct two new MPCs, keeping all points except for the tth unchanged, and such that the new points, x_t^{Δ} and $x_t^{\Delta\Delta}$, and x_t, itself, are collinear. Then, we argue that if U respects the Blackwell order, it must produce expansive errors for these new points, and that they must be mapped to the same point by φ.

Step 6 ("Filling in a Small Gap"): The proof is almost done, but there is a small gap left to be filled. We need to show that i) the points obtained by "bringing in the error point in" and "bringing non-error points in" are mapped to the same point; and ii) the points at which the "bringing in the error point in" and "bringing non-error points in" processes meet are also mapped to the same point, x^*. ∎

Next, the following lemma addresses the special case in which the error corresponds to an uninformative experiment, i.e., $\rho_B = \delta_\mu$, which allows us to move from δ_μ to the environment of Proposition 3.3.

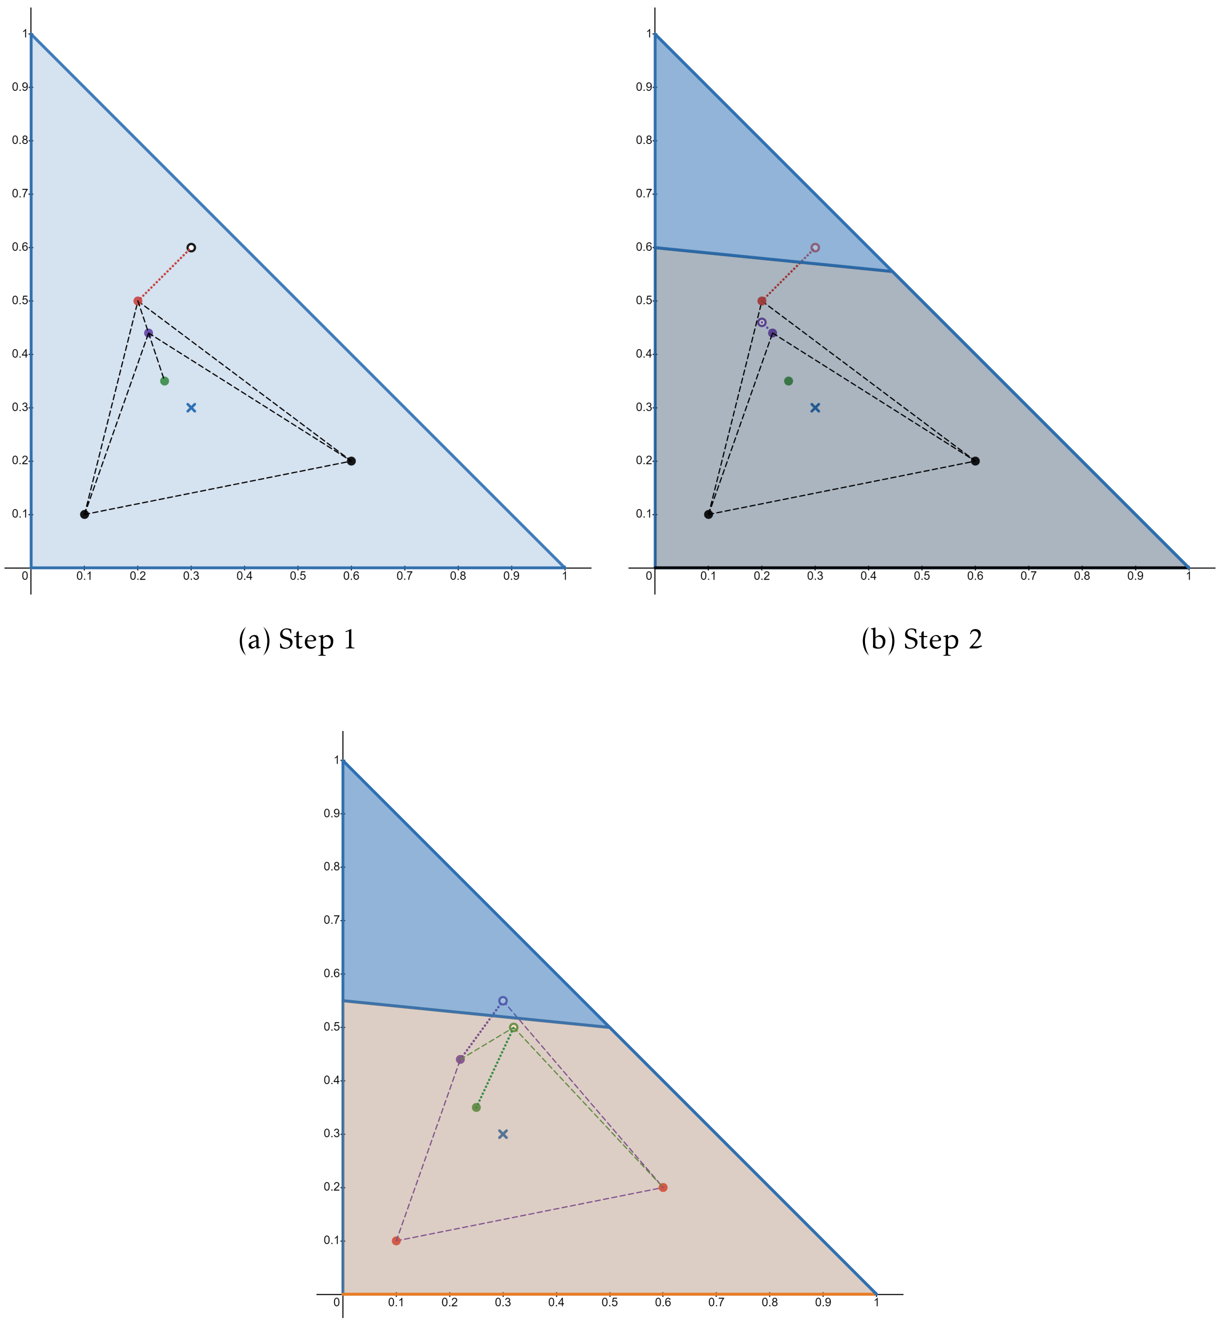

(a) Step 1

(b) Step 2

(c) Step 3

Figure 2: First steps of Proposition 3.3's proof

Lemma 3.4. *If U respects the Blackwell order and $\varphi(\mu) \neq \mu$, then the updating rule produces an expansive error for some other distribution ρ_B with affinely independent support and $\left|\operatorname{supp}\rho_B\right| \geq 2$.*

Before proving the theorem, let us study the two-state special case, as we will appeal to our analysis there in the proof of Theorem 4.1.

3.1 Two States

Our first result of this subsection anticipates our full treatment of contractive errors in the next section. It is useful to have this lemma here, as we appeal to it in the proof of Proposition 3.6.

Lemma 3.5. *Let $n = 2$, and suppose an updating rule U produces a contractive error for some $x > \mu$ and respects the Blackwell order. Then, there exists $x^* \in [\mu, x)$ such that for all $y \in [x^*, 1)$, $\varphi(y) = x^*$. The mirrored statement holds if $x < \mu$.*

This lemma is similar to Steps 2, 3 and 4, above. In short, we argue that if U respects the Blackwell order, if one moves inward (towards the prior) from a posterior, x, for which U has a contractive error, U must also have a contractive error for these new points, which must all be mapped to the same incorrect belief, x^*, by φ. In fact, all points that are more extreme than x^*, except for 1, including those that are more extreme than x we started with, must be mapped to x^* by φ.

With this in hand, we state the main result of this subsection:

Proposition 3.6. *Let the number of states $n = 2$. If U respects the Blackwell order and produces an expansive error for some posterior $x' \in (\mu, 1)$, $\varphi(x) = x^* > x'$ for all $x \in (0, x^*)$. Moreover, there exist two intervals (one of which is possibly empty) $I_1 := [x^*, \bar{x})$ and $I_2 := [\bar{x}, 1)$ ($x^* \leq \bar{x} < 1$), with $\varphi(x) = x$ for all $x \in I_1$ and $\varphi(x) = \bar{x}$ for all $x \in I_2$.*

The mirrored statement holds if $x' \in (0, \mu)$, instead.

This proposition establishes that any updating rule that respects the Blackwell order and produces an expansive error must be extremely coarse. Unless the signal realization leaves the decision maker certain that the state is low, say, he must interpret all signal realizations favoring the low state as, instead, favoring the high state.

4 Contractive Errors

In this section, we turn our attention to updating rules that produce contractive errors. An example of such an updating rule from the literature is "conservative" Bayesianism (Edwards (1968)), in which case $\varphi(x) = \xi\mu + (1-\xi)x$ for all $x \in \Delta$, for some $\xi \in (0,1)$.

Theorem 4.1. *When there are three or more states, any updating rule that respects the Blackwell order, produces a contractive error, and does not produce an expansive error, must be occasionally stubborn.*

A corollary of this is

Corollary 4.2. *When there are three or more states, if an updating rule respects the Blackwell order, produces a contractive error, and does not produce an expansive error, it must be trivial on the interior of the probability simplex,* $\operatorname{int}\Delta$.

We will prove this by showing that an updating rule that respects the Blackwell order and produces a contractive error must either produce an expansive error on that face, in which case Theorem 3.1 holds; or be such that every point in $\operatorname{int} S_m$ is mapped to μ by φ. To that end, our first observation is the easy fact that updating rules that do not produce expansive errors must be such that $\varphi(\mu) = \mu$.

Our next result points out a key difference between contractive and expansive errors. Namely, updating rules that produce (and only produce) contractive errors must map every point "toward" the prior.

Lemma 4.3. *If an updating rule U is contractive for some Bayesian distribution over posteriors ρ_B, does not produce an expansive error, and respects the Blackwell order then $\varphi(x) \in \ell_x$ for all $x \in \operatorname{supp}\rho_B$.*

This lemma allows us to restrict some of the analysis to a single-dimensional environment. The next lemma is a consequence of this and Lemma 3.5.

Lemma 4.4. *Suppose an updating rule U produces a contractive error for some $x_1 \in \Delta$, does not produce an expansive error, and respects the Blackwell order. Writing $\varphi(x_1) =: \hat{x}_1 \neq x_1$, there exists an $x_1^* \in \ell_{\hat{x}_1}$ such that for all $x' \in \ell^\circ\left(x_1^*, x_1\right)$, $\varphi(x') = x_1^*$.*

We follow this result with Proposition 4.5, which is an analog of Proposition 3.3.

Proposition 4.5. *Suppose an updating rule U produces a contractive error for some x_1 in support of some ρ_B with affinely independent support on n points, does not produce an expansive error, and respects the Blackwell order. Then U produces a contractive error for all x for which $x = \lambda \cdot \operatorname{supp}\rho_B$, where $\sum_j^n \lambda_j = 1$, $\lambda_j \in [0,1]$ for all $j = 1,\dots,n$ and $\lambda_1 > 0$. Moreover, for all such x, $\varphi(x) = \mu$.*

Proof. We sketch the proof here, leaving the proof to Appendix A.7.

Step 1 ("Moving Along the Edge"): We start with some belief x_1, for which U has an contractive error, and some ρ_B (with affinely independent support) of which x_1 is a support point. We construct a new distribution ρ_B', by keeping all of the support points of ρ_B the same except for the 1st, which we "bring in" along an arbitrary one of the edges. We show that $\varphi\left(x_1'\right) = \mu$ for any such new x_1' or else we could find a decision problem where ρ_B' is strictly preferred to ρ_B. We then do a similar procedure for new distributions $\rho_B^\dagger$, which are constructed from ρ_B by keeping all but the sth point ($s \neq 1$) the same, and "bringing" the sth point in along the edge connecting it and x_1. All such points must be mapped to μ by φ.

Step 2 ("Face Points Mapped to the Prior"): Our final step is to fill in the remaining faces of the simplex that is the convex hull of $\operatorname{supp}\rho_B$, Δ_{ρ_B}. We do this starting with the edges tackled in step 1, taking convex combinations of points on those edges to obtain points on the interior of the 2-dimensional faces of Δ_{ρ_B} that have x_1 as a vertex. All such points must be mapped to μ by φ. Then we take convex combinations of points on those 2-d faces of Δ_{ρ_B} to obtain points on the interior of the 3-d faces of Δ_{ρ_B} that have x_1 as a vertex. This process continues until we fill in $\operatorname{int}\Delta_{\rho_B}$ itself, completing the proof. ■

We can now prove the theorem.

Proof of Theorem 4.1. Let U have a contractive error for some $x_a \in S_m$ but not an expansive error for any $x \in S_m$. Fix an arbitrary $x \in \operatorname{int} S_t$ where S_t ($t \geq m$) is a face of Δ and where S_m is a face of S_t, i.e., $S_t \in \mathcal{S}_t(S_m)$.

Case 1: $x \notin \operatorname{int}\Delta$. Any such x lies in the relative interior of the convex hull of k ($\leq t$) affinely independent points, $\{x_1,\dots,x_k\}$ that lie in S_t, one of which is x_a. Pick some $x' \in$

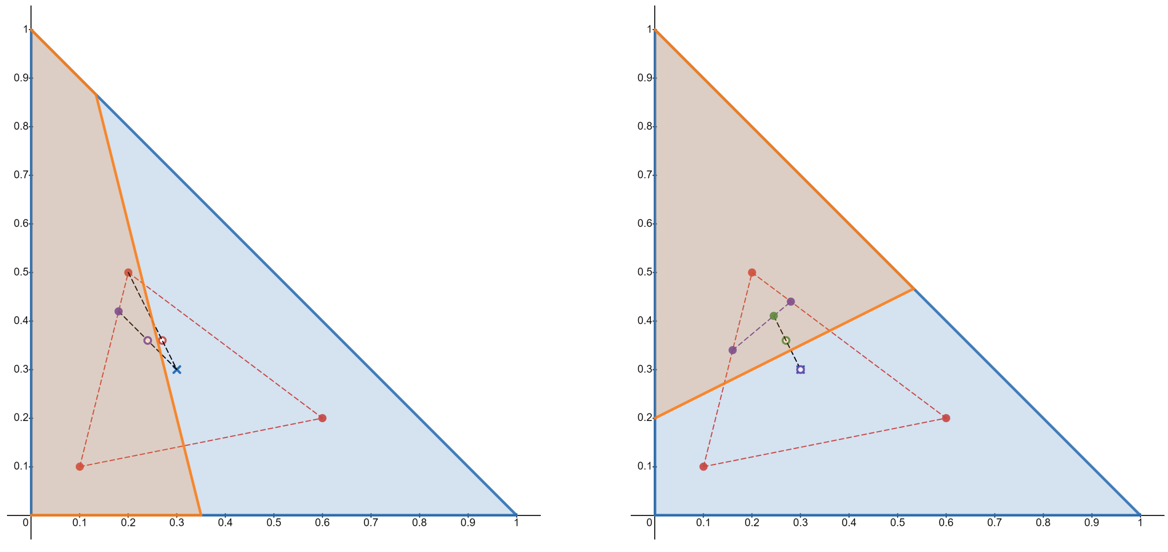

(a) Step 1 (b) Step 2

Figure 3: Proposition 4.5 Proof

$\operatorname{int}\operatorname{conv}\{x_1,\ldots,x_k\}$ and some $y\in\partial\Delta$ with $\mu=\lambda y+(1-\lambda)x'$ for some $\lambda\in(0,1)$. Accordingly, $\mu\in\operatorname{int}\operatorname{conv}\{x_1,\ldots,x_k,y\}$. Setting $\operatorname{supp}\rho_B=\{x_1,\ldots,x_k,y\}$, by Proposition 4.5, $\varphi(x)=\mu$.

Case 2: $x\in\operatorname{int}\Delta$**.** If there exists a collection of $3\leq k\leq n$ affinely independent points, one of which is x_a such that x and μ lie in the relative interior of their convex hull, we are done. Suppose instead there does not exist such a collection. However, we can then pick a distribution ρ_B with support on $3\leq k\leq n$ affinely independent points, one of which is x_a, then a distribution ρ_B' with support with $3\leq k\leq n$ affinely independent points with some $x'\in\operatorname{int}\rho_B$ in support and $x\in\operatorname{int}\operatorname{conv}\operatorname{supp}\rho_B'$. By Proposition 4.5, $\varphi(x)=\mu$.

Claim A.13 disciplines where φ can map the vertices, completing the proof. ■

4.1 Two States

Having already proved Lemma 3.5–which specifies that when there are just two states, updating rules that produce contractive errors for some belief must map all beliefs (other than the most extreme belief, the vertex) more extreme than the incorrectly mapped one to the same incorrect x^* (in a contractive manner)–there is nothing left to do. An immediate consequence of the lemma is

Proposition 4.6. *Let the number of states $n = 2$. If U respects the Blackwell order, produces a contractive error for some posterior $x' \in (\mu, 1]$, and does not produce an expansive error, then $\varphi(x) = x^* \in [\mu, x')$ for all $x \in [x^*, 1)$. Moreover, $\varphi(x) = x$ for all $x \in [\mu, x^*]$ and $\varphi(1) \geq x^*$.*

The mirrored statement holds for contractive errors for some $x' \in [0, \mu)$.

When there are just two states, like when there are expansive errors, updating rules that respect the Blackwell order must be coarse. In contrast to the expansive-error case, errors cannot be extreme: signal realizations that lead to Bayesian beliefs more confident about one state (than under the prior), must still yield beliefs more confident about that state under the non-Bayesian rule.

5 Updating Rules that Respect the Blackwell Order, and Those That Do Not

At last we arrive at the main result of the paper, a full characterization of precisely which updating rules (that systematically distort beliefs) respect the Blackwell order.

If there are just two states, we say an updating rule is **Occasionally Coarse** if there exist two (one or both of which are possibly empty) intervals $C_1 := (0, a)$ and $C_2 := (b, 1)$, with $0 \leq a \leq b \leq 1$ such that

1. $\varphi(x) = a$ for all $x \in C_1$,
2. $\varphi(x) = b$ for all $x \in C_2$,
3. $\varphi(x) = x$ for all $x \in [a, b]$, and
4. $\varphi(0) \leq a$ and $\varphi(1) \geq b$.

An updating rule that is occasionally coarse has at most two convex regions (intervals) of beliefs on which it collapses any belief to a single belief. Moreover, there is possibly another convex region, in between these two coarse regions, in which the updating rule is Bayes' law. The DM may also make mistakes about beliefs corresponding to certainty (the vertices) but they cannot be too severe. This is exactly the "more extreme" requirement for occasionally stubborn updating rules.

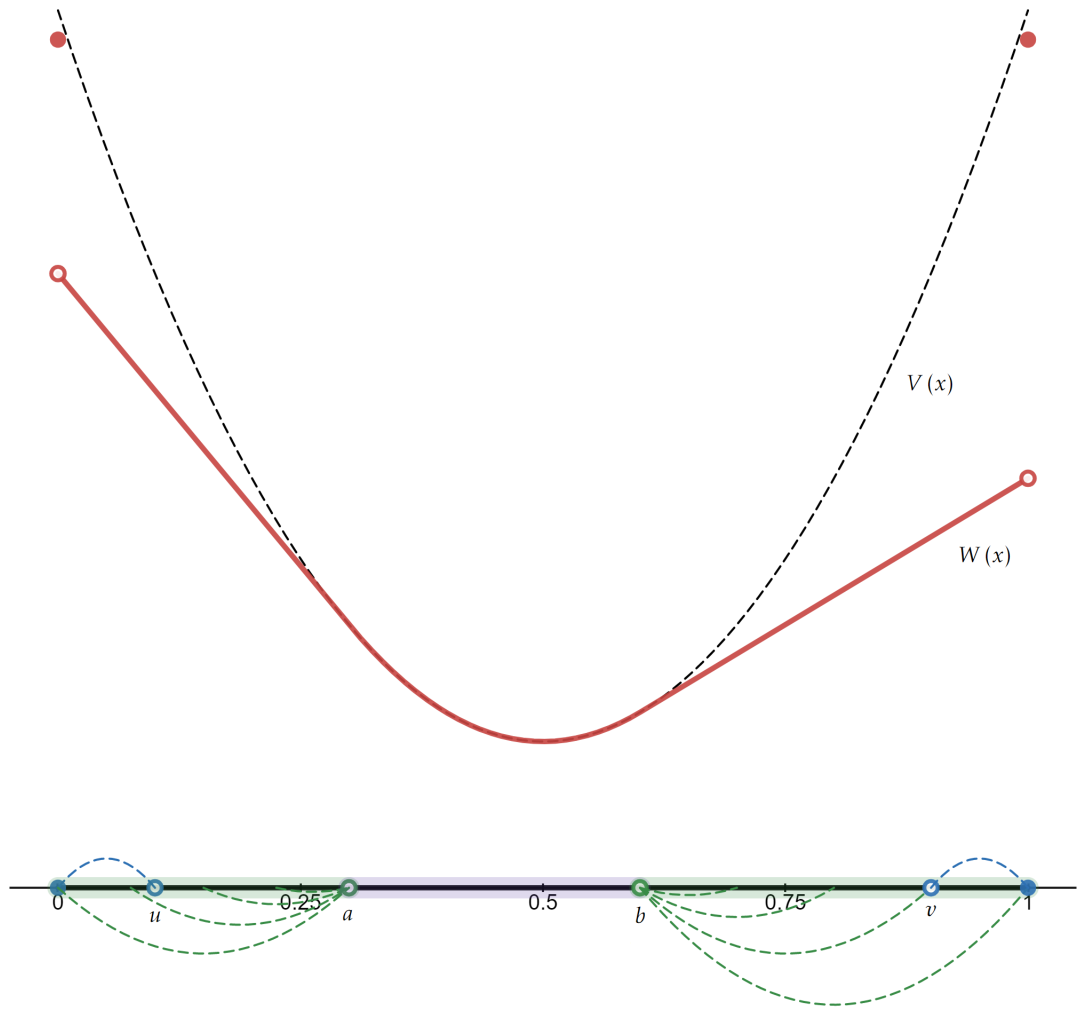

Figure 4: Occasionally Coarse Updating

Theorem 5.1. *If there are two states ($n = 2$), an updating rule respects the Blackwell order if and only if it is occasionally coarse. If there are three or more states ($n \geq 3$), an updating rule respects the Blackwell order if and only if it is occasionally stubborn.*

Example 5.2 (Occasionally Coarse Rules)**.** 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, $V(x) = -x(1-x) + .3$. Figure 4 illustrates the updating rule on the simplex, the value function V, and the corresponding indirect value function W. Here is an Interactive Link, where one can adjust the parameters–u, a, b and v that determine the family of

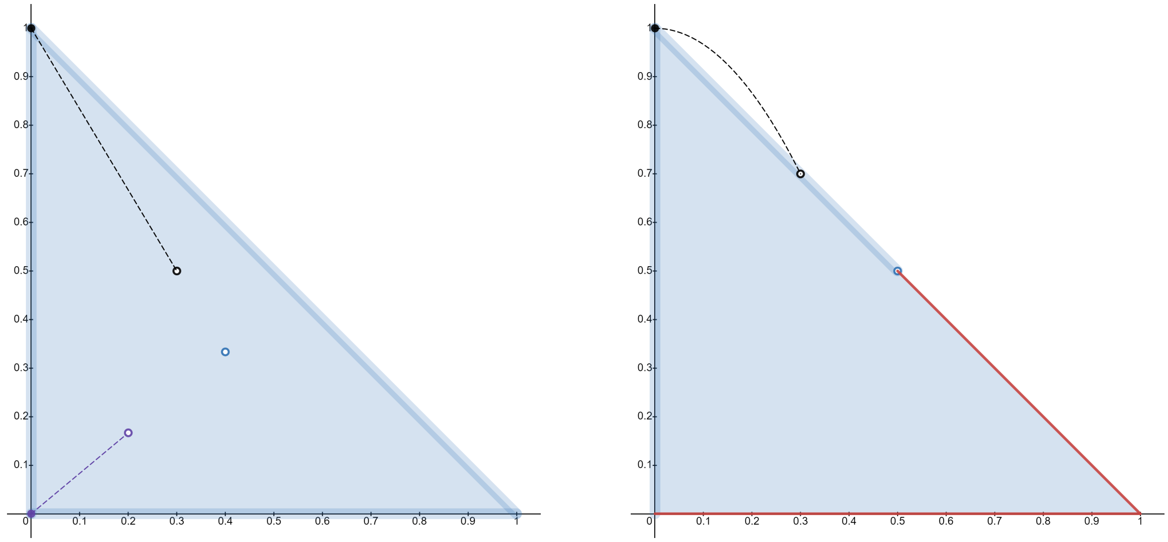

(a) Trivial on Edges

(b) Trivial on 1.5 Edges

Figure 5: Occasionally Stubborn Updating

occasionally coarse rules–by moving the corresponding sliders.

Example 5.3 (Occasionally Stubborn Rules)**.** Figure 5 illustrates two occasionally stubborn updating rules when there are 3 states. In the first (5a) $\varphi((0,0)) = \left(\frac{1}{5}, \frac{1}{6}\right)$, $\varphi((0,1)) = \left(\frac{3}{10}, \frac{1}{2}\right)$, $\varphi((1,0)) = (1,0)$, and $\varphi(x) = \left(\frac{1}{5}, \frac{1}{3}\right)$ for all other $x \in \Delta$.

In the second (5b) $\varphi((0,1)) = \left(\frac{3}{10}, \frac{7}{10}\right)$, $\varphi(x) = \left(\frac{1}{2}, \frac{1}{2}\right)$ for all x with $0 < x_1 < \frac{1}{2}$ and $x_2 = 1 - x_1$, $\varphi(x) = x$ for all x with $\frac{1}{2} \leq x_1 \leq 1$ and $x_2 = 1 - x_1$ or $0 \leq x_1 \leq 1$ and $x_2 = 0$, and $\varphi(x) = \left(\frac{1}{2}, \frac{1}{2}\right)$ for all other $x \in \Delta$.

We say that an updating rule continuously distorts beliefs if φ is continuous on Δ. We say that an updating rule smoothly distorts beliefs if φ is continuously differentiable on $\operatorname{int}\Delta$. When there are at least three states, as any point on the boundary $\partial\Delta$ is the limit of a sequence of beliefs on the interior of the sequence, we have the following corollary of Theorem 5.1.

Corollary 5.4. *If there are three or more states ($n \geq 3$), a non-trivial updating that continuously distorts beliefs respects the Blackwell order if and only if it is Bayes' law.*

When there are two states, imposing continuity of φ only refines the occasionally coarse updating rules only slightly. The only difference is that now $\varphi(0) = a$ and $\varphi(1) = b$.

Instead, imposing smoothness is needed to produce Bayes' law, as $\varphi(x) = \max\{a, \min\{x, b\}\}$ is continuous on $[0,1]$ but not differentiable at a (or b) for a (b) $\in (0,1)$:

Corollary 5.5. *A non-trivial updating that smoothly distorts beliefs respects the Blackwell order if and only if it is Bayes' law.*

6 Some Final Remarks

One could generalize the definition of an updating rule to where it is now a map $U\colon \Delta \times \Pi \times \mathcal{U}(A, \Theta) \to \mathbb{R}$, where $\mathcal{U}(A, \Theta)$ is a finite set of compact-action decision problems. That is, the updating rule could adapt to the decision problem itself. In this case, our approach would not work; in fact, there are other such updating rules (that systematically distort beliefs) that respect the Blackwell order. For example, any updating rule that maps beliefs within an element of the power diagram induced by the decision problem within the element renders the value of information positive.

6.1 Non-Expected-Utility Preferences

Our results extend beyond the expected-utility realm. Specifically, Karni and Safra (2022) introduce axioms that ultimately yield a value function

$$\tilde{V}(x) := \max_{a \in A} w\left(\mathbb{E}_{x(\theta)} u(a, \theta)\right),$$

where $w\colon \mathbb{R} \to \mathbb{R}$ is strictly increasing and convex.[1] Importantly, note that when A is finite, the projection of $\tilde{V}$ onto Δ produces the same **Power Diagram** as that corresponding to a DM with EU preferences, regardless of the shape of the (strictly increasing) w. It is clear, then, that all of our proofs go through, replacing $\cdot$ with $w(\cdot)$ in the value function.

[1] As Karni and Safra (2022) note, this formulation is similar to the smooth ambiguity representation of Klibanoff et al. (2005), with the modification that w, here, is *convex*.

6.2 Updating Rules that Do Not Systematically Distort Beliefs

Although many updating rules in the literature systematically distort beliefs, not all do, including some realistic ones. As noted by de Clippel and Zhang (2022), updating rules that correspond to information aggregation failures or correlation neglect for multiple signals may not systematically distort beliefs. What can we say about these rules?

Not much. Obviously, our sufficiency result continues to hold, but our necessity result does not. For instance, an updating rule that is Bayes' law for any experiment except a completely uninformative experiment, in which case it produces some posterior other than μ with probability one, respects the Blackwell order. Some insights do carry over; however, like the fact that updating rules with errors that produce more "extreme" beliefs than Bayes' law must do so for all less informative experiments. Mirroring Claim A.1,

Remark 6.1. Let experiments π and π' correspond to ρ_B and ρ'_B, respectively; and let $\pi \succeq \pi'$. If U respects the Blackwell order and is such that $\operatorname{conv}\operatorname{supp}\Phi(\rho_B) \not\subseteq \operatorname{conv}\operatorname{supp}\rho_B$, then $\operatorname{conv}\operatorname{supp}\Phi\left(\rho'_B\right) \not\subseteq \operatorname{conv}\operatorname{supp}\rho_B$.

References

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

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.

Edi Karni and Zvi Safra. Hybrid decision model and the ranking of experiments. *Journal of Mathematical Economics*, 101:102700, 2022.

Peter Klibanoff, Massimo Marinacci, and Sujoy Mukerji. A smooth model of decision making under ambiguity. *Econometrica*, 73(6):1849–1892, 2005.

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.

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.

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.

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

Wenhao Wu. A geometric blackwell's order. *Mimeo*, 2022.

A Omitted Proofs

A.1 Proposition 3.3 Proof

Proof. Step 1 ("Bringing the Error Point In"): Fix μ and π that yield ρ_B with k affinely independent points of support $\{x_1,\ldots,x_k\}$ ($k \in \mathbb{N}$, $2 \leq k \leq n$), and let U have an expansive error for ρ_B. That is, for $x_1 \in \operatorname{supp}\rho_B$, $\varphi(x_1) = \hat{x}_1 \notin \operatorname{conv}\operatorname{supp}\rho_B$. Let $p \equiv p_1 \in (0,1)$ denote $\mathbb{P}_{\rho_B}(x_1)$; and let $p_j \in (0,1)$ denote $\mathbb{P}_{\rho_B}\left(x_j\right)$ and $\hat{x}_j := \varphi\left(x_j\right)$ for all $j \neq 1$.

Consider first two additional Bayesian distributions over posteriors. The first, ρ_B', corresponding to experiment π', has support on a subset of $\left\{x_1', x_2, \ldots x_k\right\}$; that is, all of the support points except for the 1st one are also support points of ρ_B. Moreover, let $x_1' \in \operatorname{conv}\operatorname{supp}\rho_B \setminus \left\{x_1'\right\}$, so that ρ_B' is a strict MPC of ρ_B. Let $q \in (0,1)$ denote $\mathbb{P}_{\rho_B'}\left(x_1'\right)$. Let $q_j \in (0,1)$ denote $\mathbb{P}_{\rho_B'}\left(x_j\right)$ for all $j \neq 1$. Note that $q > p$ and $q_j \leq p_j$ for all $j \neq 1$, with at least one inequality strict.

The second, ρ_B'', corresponding to experiment π'', has support on a subset of $\left\{x_1'', x_2, \ldots x_k\right\}$; i.e., all but the 1st support point are also in support of ρ_B'. Moreover, let $x_1' = \gamma x_1 + (1-\gamma) x_1''$ for some $\gamma \in (0,1)$, so that this distribution is a strict MPC of ρ_B' (and therefore also of ρ_B) and so that x_1, x_1' and x_1'' are collinear. Let $r \in (0,1)$ denote $\mathbb{P}_{\rho_B''}\left(x_1''\right)$. Let $r_j \in (0,1)$ denote $\mathbb{P}_{\rho_B''}\left(x_j\right)$ for all $j \neq 1$. Note that $r > q$ and $r_j \leq q_j$ for all $j \neq 1$, with at least one inequality strict. Let $\hat{x}_1' := \varphi\left(x_1'\right)$ and $\hat{x}_1'' := \varphi\left(x_1''\right)$.

Claim A.1. *Step 2 ("Banishing the New Points"): $\hat{x}_1' \notin \operatorname{conv}\operatorname{supp}\rho_B$ and $\hat{x}_1'' \notin \operatorname{conv}\operatorname{supp}\rho_B$.*

Proof. Suppose for the sake of contraposition that $\hat{x}_1' \in \operatorname{conv}\operatorname{supp}\rho_B$. As $\hat{x}_1 \notin \operatorname{conv}\operatorname{supp}\rho_B$, $\left\{\hat{x}_1'\right\} \cup \operatorname{conv}\operatorname{supp}\rho_B$, (which equals $\operatorname{conv}\operatorname{supp}\rho_B$) and $\hat{x}_1$ can be strictly separated by some hyperplane

$$H_{\alpha,\beta} := \left\{x \in \mathbb{R}^{n-1} \middle| \alpha \cdot x = \beta\right\}.$$

WLOG we may assume $\operatorname{conv}\operatorname{supp}\rho_B$ is a strict subset of the closed half-space $\left\{x \in \mathbb{R}^{n-1} \middle| \alpha \cdot x \leq \beta\right\}$. Consider the value function $V(x) = \max\{0, \alpha \cdot x - \beta\}$. Define the sets

$$A := \left\{x_j \in \{x_2,\ldots,x_k\} \middle| \alpha \cdot \hat{x}_j > \beta\right\},$$

and

$$B := \left\{x_j \in \{x_2, \dots, x_k\} | \alpha \cdot \hat{x}_j \le \beta\right\}.$$

We may relabel the points so that the first l xs lie in A, the last $(k-1-l)$ xs lie in B. Thus, the agent's payoff under experiment π is

$$\sum_{j=2}^{l} p_j \left(\alpha \cdot x_j - \beta\right) + p\left(\alpha \cdot x_1 - \beta\right),$$

whereas the agent's payoff under experiment π' is

$$\sum_{j=2}^{l} q_j \left(\alpha \cdot x_j - \beta\right).$$

Taking the difference of these two expressions, we obtain

$$\sum_{j=2}^{l} \underbrace{\left(p_j - q_j\right)}_{\ge 0} \underbrace{\left(\alpha \cdot x_j - \beta\right)}_{<0} + \underbrace{p\left(\alpha \cdot x_1 - \beta\right)}_{<0} < 0.$$

We have obtained a violation of Blackwell's theorem. ■

Claim A.2. *Step 3 ("If Distinct, then More Contracted $\Rightarrow$ More Extreme Errors"): If $\hat{x}'_1 \ne \hat{x}''_1$, there exists a hyperplane that strictly separates $\hat{x}''_1$ and* $\operatorname{conv}\left(\left\{\hat{x}'_1\right\} \cup \operatorname{supp} \rho'_B\right)$.

Proof. Suppose for the sake of contraposition that there does not exist a hyperplane that strictly separates $\hat{x}''_1$ and $\operatorname{conv}\left(\left\{\hat{x}'_1\right\} \cup \operatorname{supp} \rho'_B\right)$. This implies, via the strict separating hyperplane theorem, that $\hat{x}''_1 \in \operatorname{conv}\left(\left\{\hat{x}'_1\right\} \cup \operatorname{supp} \rho'_B\right)$, since the latter set is by construction compact and convex. By Claim A.1, neither $\hat{x}'_1$ nor $\hat{x}''_1$ lie in $\operatorname{conv}\operatorname{supp} \rho'_B$ (as $\operatorname{supp} \rho'_B \subseteq \operatorname{supp} \rho_B$). Accordingly, $\hat{x}'_1 \notin \operatorname{conv}\left(\left\{\hat{x}''_1\right\} \cup \operatorname{supp} \rho'_B\right)$. Thus, by the strict separating hyperplane theorem, there exists an ($n-2$ dimensional) hyperplane that strictly separates $\hat{x}'_1$ and $\operatorname{conv}\left(\left\{\hat{x}''_1\right\} \cup \operatorname{supp} \rho'_B\right)$. From here, the proof replicates that of Claim A.1. ■

Claim A.3. *Step 4 ("All Meet the Same End"): $\hat{x}'_i = \hat{x}''_i =: \hat{x}^*_i$*

Proof. Suppose for the sake of contraposition not. Consider a third experiment $\pi^\dagger$ that produces a correct distribution over posteriors $\rho^\dagger_B$ with support on a subset of

$$\left\{x_1, \dots, x_{i-1}, x_i, x''_i, x_{i+1}, \dots x_k\right\},$$

where $\mathbb{P}\left(x_j\right) = q_j$ for all $j \neq i$ and the probabilities of x_i'' and x_i are precisely such that ρ_B' is a strict mean-preserving contraction of this distribution. By Claim A.2, there exists a hyperplane that strictly separates $\hat{x}_i''$ and $\operatorname{conv}\left(\left\{\hat{x}_i'\right\} \cup \operatorname{supp} \rho_B'\right)$, $H_{\gamma,\delta} := \left\{x \in \mathbb{R}^{n-1} \middle| \gamma \cdot x = \delta\right\}$, where $\operatorname{conv}\left(\left\{\hat{x}_i'\right\} \cup \operatorname{supp} \rho_B'\right)$ is a strict subset of the closed half-space $\left\{x \in \mathbb{R}^{n-1} \middle| \gamma \cdot x \leq \delta\right\}$. For value function $V(x) = \max\{0, \gamma \cdot x - \delta\}$, the difference in payoffs for the agent from experiments π' and $\pi^{\dagger}$ is strictly negative, a violation. ■

Step 5 ("Repeating Steps 1-4"): Consider second two additional Bayesian distributions over posteriors $\rho_B^{\triangle}$ and $\rho_B^{\triangle\triangle}$, corresponding to experiments $\pi^{\triangle}$ and $\pi^{\triangle\triangle}$, respectively. The first has support on a subset of $\left\{x_1, \ldots, x_{t-1}, x_t^{\triangle}, x_{t+1}, \ldots x_k\right\}$, where $t \in \{2, \ldots, k\}$. Again, all but one of the support points are also in support ρ_B, but now the support point we are changing is not the 1st. Moreover, let

$$x_t^{\triangle} \in \operatorname{conv} \operatorname{supp} \rho_B \setminus \left(\operatorname{conv}\left(\operatorname{supp} \rho_B \setminus \{x_1\}\right) \cup \{x_t\}\right),$$

i.e., $x_t^{\triangle}$ is neither x_t nor a convex combination of exclusively points in the support of ρ_B other than x_1. $\rho_B^{\triangle}$ is a strict MPC of ρ_B. Let $s_j \in (0,1)$ denote $\mathbb{P}_{\rho_B^{\triangle}}\left(x_j\right)$ for all j. Note that $s_t > p_t$, $s \equiv s_1 < p_1 \equiv p$, and $s_j \leq p_j$ for all $j \neq t, 1$.

The second has support on a subset of $\left\{x_1, \ldots, x_{t-1}, x_t^{\triangle\triangle}, x_{t+1}, \ldots x_k\right\}$, where $t \in \{2, \ldots, k\}$. Yet again, all but one of the support points also support ρ_B. Moreover, let

$$x_t^{\triangle\triangle} \in \operatorname{conv} \operatorname{supp} \rho_B^{\triangle} \setminus \left(\operatorname{conv}\left(\operatorname{supp} \rho_B^{\triangle} \setminus \{x_1\}\right) \cup \left\{x_t^{\triangle}\right\}\right),$$

i.e., $x_t^{\triangle\triangle}$ is neither $x_t^{\triangle}$ nor a convex combination of exclusively points in the support of $\rho_B^{\triangle}$ other than x_1. Let $x_t^{\triangle} = \gamma x_t + (1-\gamma) x_t^{\triangle\triangle}$ for some $\gamma \in (0,1)$ so that $x_t^{\triangle}$, $x^{\triangle\triangle}$, and x_t are collinear. This distribution is a strict MPC of $\rho_B^{\triangle}$ (and therefore also of ρ_B). Let $u_j \in (0,1)$ denote $\mathbb{P}_{\rho_B^{\triangle\triangle}}\left(x_j\right)$ for all j. Note that $u_t > s_t$, $u \equiv u_1 < s_i \equiv s$, and $u_j \leq s_j$ for all $j \neq t, 1$. Let $\hat{x}_t^{\triangle} := \varphi\left(x_t^{\triangle}\right)$ and $\hat{x}_t^{\triangle\triangle} := \varphi\left(x_t^{\triangle\triangle}\right)$.

We have the following three claims, the proofs for which mirror (*mutatis mutandis*) those for Claims A.1, A.2, and A.3, respectively, and which we, therefore, omit.

Claim A.4. *$\hat{x}_t^{\triangle} \notin \operatorname{conv} \operatorname{supp} \rho_B$ and $\hat{x}_t^{\triangle\triangle} \notin \operatorname{conv} \operatorname{supp} \rho_B$.*

Claim A.5. *If $\hat{x}_t^{\triangle} \neq \hat{x}_t^{\triangle\triangle}$, there exists a hyperplane that strictly separates $\hat{x}_t^{\triangle\triangle}$ and $\operatorname{conv}\left(\left\{\hat{x}_t^{\triangle}\right\} \cup \operatorname{supp} \rho_B^{\triangle}\right)$.*

Claim A.6. $\hat{x}_t^{\triangle} = \hat{x}_t^{\triangle\triangle} =: \hat{x}_t^*$.

The final step is to fill in a small gap.

Claim A.7. *Step 6 ("Filling in a Small Gap"): Let U have expansive errors for distinct posteriors $x, y \in \Delta$. If*

$$\varphi(\lambda x + (1-\lambda) y) = \hat{y} \ \forall\ \lambda \in [0, \lambda^*) \ (\lambda^* \in (0,1)) \quad \textit{and} \quad \varphi(\lambda x + (1-\lambda) y) = \hat{x} \ \forall\ \lambda \in (\lambda^*, 1],$$

$\hat{x} = \hat{y} = \varphi(\lambda x + (1-\lambda) y)$ *for all* $\lambda \in [0,1]$.

Proof. WLOG x, y, and μ are collinear. First, we show that $\hat{x} = \hat{y}$. Suppose not. Evidently, $\operatorname{conv}\{\hat{x}, x, y\}$ and $\{\hat{y}\}$ can be strictly separated by a hyperplane or $\operatorname{conv}\{\hat{y}, x, y\}$ and $\{\hat{x}\}$ can be strictly separated by a hyperplane. WLOG we assume the former. Let such a strictly separating hyperplane be

$$H_{\alpha,\beta} := \left\{ x \in \mathbb{R}^{n-1} \middle| \alpha \cdot x = \beta \right\}.$$

WLOG we specify that $\operatorname{conv}\{\hat{x}, x, y\}$ is a strict subset of the closed half-space $\left\{ x \in \mathbb{R}^{n-1} \middle| \alpha \cdot x \leq \beta \right\}$. Consider the value function $V(x) = \max\{0, \alpha \cdot x - \beta\}$ and two distributions ρ_B and ρ_B' with support on $\{x, y\}$ and $\{x', y\}$ where $x' = \lambda x + (1-\lambda) y$ for some $\lambda \in [1-\varepsilon, 1)$, where $\varepsilon > 0$ is small. Let $p := \mathbb{P}_{\rho_B}(y)$ and $p' := \mathbb{P}_{\rho_B'}(y)$ and observe that $p' < p$. By construction, ρ_B' is a strict MPC of ρ_B. However, the DM's payoff under the former is $p'(\alpha \cdot y - \beta) > p(\alpha \cdot y - \beta)$, her payoff under the latter, so U does not respect the Blackwell order. By contraposition, we must have $\hat{x} = \hat{y}$.

Finally, define $x^\circ := \lambda^* x + (1-\lambda^*) y$. We want to show that $\hat{x} = \varphi(x^\circ) =: \hat{x}^\circ$. Otherwise, we could construct three distributions ρ_B, with support on $\{x, x^\circ, y\}$; ρ_B', with support on $\{x', y'\}$, where $x' \in (x, x^\circ)$ and $y' \in (x^\circ, y)$; and ρ_B'', which either has support on $\{x^\circ, y''\}$ where $y'' \in (x^\circ, y')$ or is δ_{x°. By construction ρ_B'' is a strict MPC of ρ_B', which is a strict MPC of ρ_B. Following the previous paragraph's approach, it is apparent that if $\hat{x}^\circ \neq \hat{x}$ either ρ' is strictly preferred to ρ or ρ'' is strictly preferred to ρ', which establishes the result by contraposition. ∎

This concludes the proof of the proposition. ∎

A.2 Lemma 3.4 Proof

Proof. Let $x_1, x_2 \in \Delta$ be such that $\mu = \lambda x_1 + (1-\lambda) x_2$ for some $\lambda \in (0,1)$ and such that $\varphi(\mu) \notin \operatorname{conv}\{x_1, x_2\}$. Let ρ_B have support $\{x_1, x_2\}$. If U has an expansive error for ρ_B, we are done. Otherwise, consider instead $\rho_B' = \gamma \rho_B + (1-\gamma)\delta_\mu$ for some $\gamma \in (0,1)$. By construction, U has an expansive error for ρ_B'. But then we can generate a fusion of ρ_B', ρ_B'', with support on $\{x_3, x_4\}$ where $x_3 = \tau x_1 + (1-\tau)\mu$ and $x_4 = \iota x_2 + (1-\iota)\mu$, for some appropriately chosen $\tau, \iota \in (0,1)$. If U is not expansive for ρ_B'', the value of information is not positive, which proves the result by contraposition. ∎

A.3 Lemma 3.5 Proof

Proof. Consider ρ_B, corresponding to π with $\operatorname{supp} \rho_B = \{0, z\}$ ($z \in (\mu, 1]$), where $z > \varphi(z) = \hat{z} \geq \mu$. Also consider ρ_B', corresponding to π', with support $\{0, z'\}$ with $z' \in (\hat{z}, z)$. Let $p := \mathbb{P}_{\rho_B'}(z')$. Evidently, we cannot have $\hat{z}' := \varphi(z') > \hat{z}$ or else value function $V(x) := \max\left\{0, x - \frac{\hat{z}'+\hat{z}}{2}\right\}$ illustrates that U does not respect the Blackwell order.

Following the same logic, for ρ_B'' with support $\{0, z''\}$ with $z'' \in (\hat{z}', z')$, we must have $\hat{z}'' := \varphi(z'') \leq \hat{z}'$. Suppose for the sake of contraposition that $\hat{z}'' < \hat{z}'$. Consider the ternary distribution with support $\{0, z'', z\}$, ρ_B^m, corresponding to experiment π^m, with $\mathbb{P}_{\rho_B^m}(0) = 1 - p$. Observe that ρ_B' is a strict MPC of ρ_B^m. Consider value function $V(x) := \max\left\{0, x - \frac{\hat{z}'+\hat{z}''}{2}\right\}$. After some algebra, we see that π' is strictly superior to π^m under U if and only if

$$p \frac{z' - z''}{z - z''}\left(z - \frac{\hat{z}' + \hat{z}''}{2}\right) < p\left(z' - \frac{\hat{z}' + \hat{z}''}{2}\right) \Leftrightarrow z'' - \frac{\hat{z}' + \hat{z}''}{2} > 0,$$

which holds as $z'' > \hat{z}' > \hat{z}''$.

Accordingly, U does not respect the Blackwell order and so we must have $\hat{z}' = \hat{z}''$. Thus, we must have $\varphi(x) = x^* \geq \mu$ for all $x \in (x^*, z)$.

If $z = 1$, we are done. Suppose now that $z < 1$. Suppose there exists some $y \in (z, 1)$ with $\hat{y} := \varphi(y) > x^*$. Evidently, for all $y' \in (y, 1)$, $\hat{y} := \varphi(y') \geq \hat{y}$ or else we could construct a value function for which information has strictly negative value under U. Consider ρ_B, corresponding to π with support $\{0, z - \varepsilon, y'\}$, with $\varepsilon \in (0, z - x^*)$; and ρ_B', corresponding to π' with support $\{0, y\}$. Let $1 - p := \mathbb{P}_{\rho_B}(0) = \mathbb{P}_{\rho_B'}(0)$, $p := \mathbb{P}_{\rho_B'}(y)$, $q := \mathbb{P}_{\rho_B}(y')$, and $p - q := \mathbb{P}_{\rho_B}(z - \varepsilon)$, with $(p - q)(z - \varepsilon) + q y' = p y$.

For value function $V(x) = \max\left\{0, x - \frac{\min\{z-\varepsilon,\hat{y}\}+x^*}{2}\right\}$ the DM strictly prefers π' to π under U if and only if

$$p\left(y - \frac{\min\{z-\varepsilon,\hat{y}\}+x^*}{2}\right) > q\left(y' - \frac{\min\{z-\varepsilon,\hat{y}\}+x^*}{2}\right),$$

which holds if and only if

$$z - \varepsilon > \frac{\min\{z-\varepsilon,\hat{y}\}+x^*}{2},$$

which holds by assumption, yielding a strictly negative value for information. ■

A.4 Proposition 3.6 Proof

Proof. WLOG $x' \in (\mu, 1)$. By assumption $\varphi(x') =: \hat{x}' > x'$ By Proposition 3.3, for all $x \in (0, x')$, $\varphi(x) = x^*$, where $x^* \geq \hat{x}' > x'$.

Claim A.8. $\varphi(x^*) =: \hat{x}^* = x^*$.

Proof. If $\hat{x}^* > x^*$, then by Proposition 3.3, if U respects the Blackwell order, $\varphi(x') \geq \hat{x}^*$, a contradiction.

Next, suppose for the sake of contraposition that $\hat{x}^* < x^*$. Figure 6 illustrates this proof. Observe that i) $\hat{x}^* \geq x'$ and ii) for all $z \in (\hat{x}^*, x^*)$, $\varphi(z) =: \hat{z} \leq \hat{x}^*$ (or else we could generate a strictly negative value of information). However, then consider two experiments, ρ_B, with support on $\left\{0, \frac{x'+\mu}{2}, x^*\right\}$ with $p := \mathbb{P}_{\rho_B}(x^*)$ and $q := \mathbb{P}_{\rho_B}\left(\frac{x'+\mu}{2}\right)$; and ρ_B' with support on $\left\{0, \frac{\hat{x}^*+x^*}{2}\right\}$, where

$$(p+q)\frac{\hat{x}^*+x^*}{2} = px^* + q\frac{x'+\mu}{2}.$$

Then, consider value function

$$V(x) = \max\left\{0, x - \frac{\frac{\hat{x}^*+x^*}{2}+x^*}{2}\right\},$$

which reveals that ρ_B' (which yields a payoff of 0 to the DM under U, ignoring the payoff from 0 as it will cancel out) is strictly preferred by the DM to ρ_B (which yields a strictly negative payoff to the DM). ■

Claim A.9. $\varphi(x) = x^*$ *for all* $x \in [x', x^*]$.

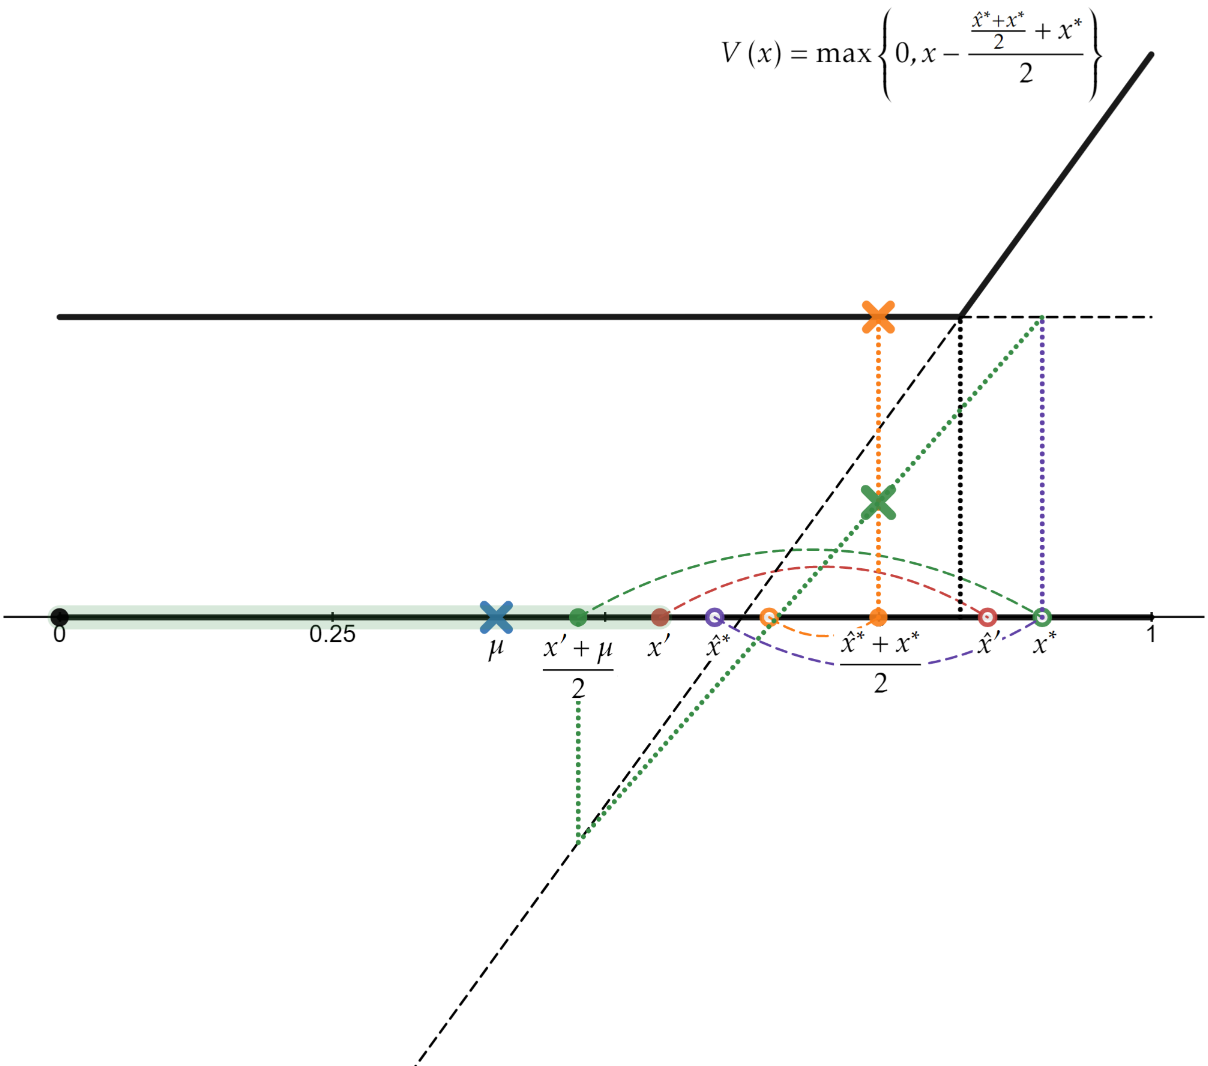

Figure 6: Claim A.8 proof

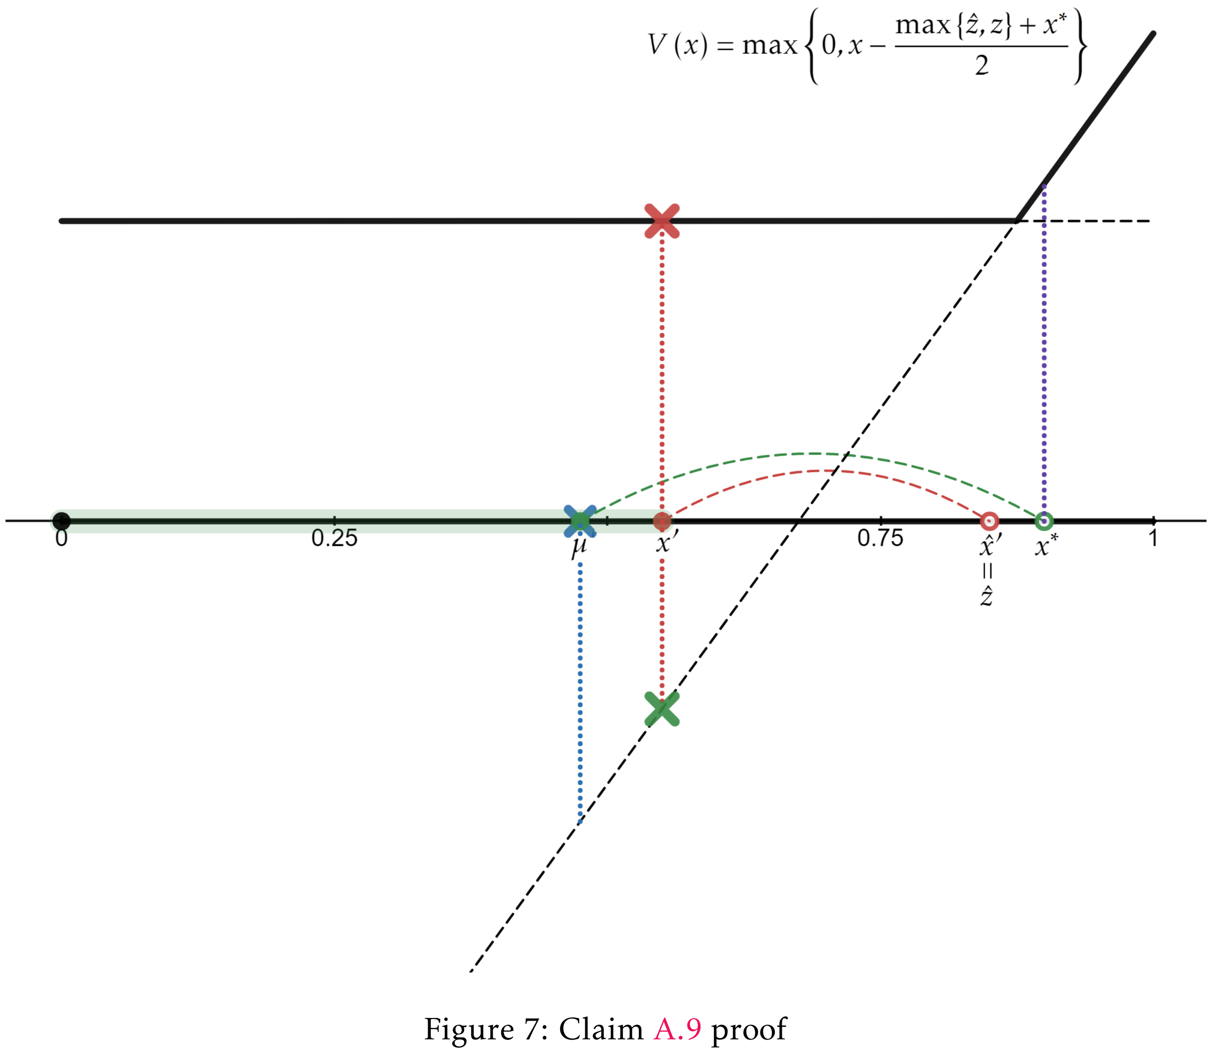

Figure 7: Claim A.9 proof

Proof. Figure 7 illustrates this proof. Evidently, we must have $\varphi(x) \leq x^*$ for all $x \in [x', x^*]$. Suppose for some $z \in [x', x^*)$ $\varphi(z) =: \hat{z} < x^*$. Consider ρ_B, corresponding to π, with support on $\{0, \mu, x^*\}$ (with respective probabilities $1-p$, $p-q$ and q) and ρ'_B, corresponding to π', with support on $\{0, z\}$ (with respective probabilities $1-p$ and p) and where we must have $pz = (p-q)\mu + qx^*$. Then for value function

$$V(x) = \max\left\{0, x - \frac{\max\{\hat{z}, z\} + x^*}{2}\right\},$$

the DM's payoff from π' is strictly higher than that from π under U if and only if

$$0 > q\left(x^* - \frac{\max\{\hat{z}, z\} + x^*}{2}\right) + (p-q)\left(\mu - \frac{\max\{\hat{z}, z\} + x^*}{2}\right),$$

which holds if and only if

$$\frac{\max\{\hat{z}, z\} + x^*}{2} > z,$$

which is true by assumption. By contraposition we obtain the result. ■

Evidently, $\varphi(y) =: \hat{y} \geq x^*$ for all $y \in (x^*, 1]$ or else we could get a strictly negative value for information. Moreover, if $\hat{y} > y$ for some $y \in (x^*, 1)$, that would imply $\hat{x}^* \geq \hat{y} > x^*$, a contradiction. Thus, $\hat{y} \leq y$ for all $y \in [x^*, 1]$. Finally, if $\hat{y} < y$ for some $y \in (x^*, 1)$, then Lemma 3.5 implies there exists some $\bar{x} \in [x^*, y)$ such that $\varphi(x) = \bar{x}$ for all $x \in [\bar{x}, 1)$. ■

A.5 Proof of Theorem 3.1

Proof. Our first step is to show that we may, without loss of generality, focus on errors produced by the updating rule for non-vertex beliefs.

Claim A.10. *If U respects the Blackwell order and produces an expansive error for some vertex e_i, it produces an expansive error for some x on the relative interior of every face S_m for which e_i is also a vertex.*

Proof. Fix some e_i for which U produces an expansive error and pick an arbitrary face of Δ, S_m, that has e_i as a vertex. By the definition of an expansive error, $\hat{e}_i := \varphi(e_i) \notin \ell_{e_i}$, the line segment between e_i and μ. If the face $S_m = \Delta$, we can construct the binary distribution, ρ_B, with support $\{e_i, x_2\}$, where x_2 is such that $\mu = \lambda x_2 + (1-\lambda)e_i$ for some

$\lambda \in (0,1)$. For any λ that is sufficiently close to 1, $\hat{e}_i \notin \operatorname{conv}\operatorname{supp}\rho_B$. By Proposition 3.3, U produces an expansive error for all $x \in \operatorname{int}\operatorname{conv}\operatorname{supp}\rho_B$, which includes some $x \in \operatorname{int}\Delta$.

If the face $S_m \neq \Delta$ then we construct the distribution ρ_B with support $\{e_i, x_1, x_2\}$, where $x_1 \in \operatorname{int} S_m$ and $x_2 \in \operatorname{int}\Delta$. For all x_1 sufficiently close to e_i and all x_2 sufficiently close to μ (with $\mu \in \operatorname{int}\operatorname{conv}\operatorname{supp}\rho_B$), $\hat{e}_i \notin \operatorname{conv}\operatorname{supp}\rho_B$. Thus, Proposition 3.3 implies U produces an expansive error for all $x \in \operatorname{int}\operatorname{conv}\operatorname{supp}\rho_B$, which includes some $x \in \operatorname{int} S_m$. ■

Thus, let $x_1 \in \operatorname{int} S_m$ for some $S_m \in \hat{\mathcal{S}}^1$ and U have an expansive error for x_1, with $\hat{x}_1 := \varphi(x_1)$. Pick an arbitrary $S_t \in \mathcal{S}_t(S_m)$. There are four cases to consider: 1. $\hat{x}_1 \notin \operatorname{int} S_t$ and $S_t \neq \Delta$, 2. $\hat{x}_1 \notin \operatorname{int} S_t$ and $S_t = \Delta$, 3. $\hat{x}_1 \in \operatorname{int} S_1$ (the S_t under scrutiny is an edge), and 4. $\hat{x}_1 \in \operatorname{int} S_t$ with $t \geq 2$ (the S_t under scrutiny is not an edge).

Case 1: $\hat{x}_1 \notin \operatorname{int} S_t$ and $S_t \neq \Delta$. Let γ_B have support on t affinely independent points $\{x_1, \ldots, x_t\}$ with $x_i \in S_t$ for all i, $x_i = e_i$ for all $i \neq 1$ where e_i are distinct vertices of S_t, and $p_i := \mathbb{P}(x_i)$.[2]

Claim A.11. *$\varphi(x) = x^* \notin S_t$ for all $x \in \operatorname{conv}\operatorname{supp}\gamma_B$ with $x = \lambda \cdot \operatorname{supp}\gamma_B$ for vector λ with $\sum_{j=1}^{t} \lambda_j = 1$, $\lambda_j \in [0,1]$ for all j and $\lambda_1 > 0$.*

Proof. Omitted, as the proof follows the proof of Proposition 3.3 nearly exactly. ■

Next, construct a γ_B' with support on t affinely independent points, of which $t-1$ are vertices of S_t and the last support point is $\lambda \cdot \operatorname{supp}\gamma_B$, where $\sum_{j=1}^{t} \lambda_j = 1$, $\lambda_j \in [0,1]$ for all j, $\lambda_1 > 0$), and λ_u is close to 1 for some $u \neq 1$ with $e_u \in \operatorname{supp}\gamma_B$. Evidently, for all $x \in \operatorname{int} S_t$, there exists a λ of this form such that $x \in \operatorname{int}\operatorname{conv}\operatorname{supp}\gamma_B'$. Accordingly, Claim A.11 implies $\varphi(x) = \hat{x}_1 = x^*$ for all $x \in \operatorname{int} S_t$.

Case 2: $\hat{x}_1 \notin \operatorname{int} S_t$ and $S_t = \Delta$. Let ρ_B have support on 2 points, x_1 and x_2, with $\mu \in \ell^\circ(x_1, x_2)$. By Proposition 3.3, $\varphi(x) = x^*$ for all $x \in \ell^\circ(x_1, x_2)$. Next construct ρ_B' with affinely independent support on $\left\{x_1', \ldots, x_n'\right\}$, where $\operatorname{conv}\operatorname{supp}\rho_B' \subset \operatorname{int}\Delta$, $x_1' \in \ell^\circ(x_1, x_2)$; and such that $x^* \notin \operatorname{int}\operatorname{conv}\operatorname{supp}\rho_B'$. By Proposition 3.3 $\varphi(x) = x^*$, for all $x \in \operatorname{int}\operatorname{conv}\operatorname{supp}\rho_B'$.

[2]Note that γ_B is not a Bayes-plausible distribution over posteriors, as $\mu \notin S_t$; however, $\operatorname{supp}\gamma_B \cup x_{t+1}$ with $x_{t+1} \notin S_t$ and μ in their convex hull is the support of some ρ_B with affinely independent support, so we may WLOG work with just γ_B. We carry this approach throughout this proof.

If $x^* \in \operatorname{int}\Delta$, we are in Case 4, below. If $x^* \notin \operatorname{int}\Delta$, then observe that for all $x \in \operatorname{int}\Delta$, we can find two points $x_1'' \in \operatorname{int}\operatorname{conv}\operatorname{supp}\rho_B'$ and $x_2'' \in \operatorname{int}\Delta$ such that $x, \mu \in \ell^\circ\left(x_1'', x_2''\right)$. Thus, Proposition 3.3 implies $\varphi(x) = x^*$.

Case 3: $\hat{x}_1 \in \operatorname{int} S_1$**.** Consider γ_B with binary support on $\{x_1, x_2\}$, where $x_1, x_2 \in S_1$ (recall x_1–which must also be in $\operatorname{int} S_1$–is the specified point for which U has an expansive error). Define

$$e_i^* := \left\{e_i \in E \mid x_1 \in \ell^\circ\left(\hat{x}_1, e_i^*\right)\right\}.$$

By construction, this is well-defined. Then,

Claim A.12. *$\varphi(x) = x^* = \hat{x}_1 \in \operatorname{int} S_1$ for all $x \in \ell^\circ\left(x^*, e_i^*\right)$. Moreover, either $\varphi(x) = x$ for all $x \in \operatorname{int} S_1 \setminus \ell^\circ\left(x^*, e_i^*\right)$ or $\varphi(x) = x^*$ for all $x \in \operatorname{int} S_1$.*

Proof. Following the proofs of Claims A.8 and A.9, $\varphi(x) = x^* = \hat{x}_1 \in \operatorname{int} S_1$ for all $x \in \ell^\circ\left(x^*, e_i^*\right)$. Moreover, following the remainder of the proof of Proposition 3.6, either i. $\varphi(x) = x$ for all $x \in \operatorname{int} S_1 \setminus \ell^\circ\left(x^*, e_i^*\right)$; or ii. there is some $\tilde{x} \in \operatorname{int} S_1 \setminus \ell^\circ\left(x^*, e_i^*\right)$ such that $\varphi(x) = \tilde{x}$ for all $x \in \operatorname{int} S_1 \setminus \ell^\circ\left(\tilde{x}, e_i^*\right)$ and $\varphi(x) = x$ for all $x = \lambda x^* + (1-\lambda)\tilde{x}$ for some $\lambda \in [0,1]$. Suppose for the sake of contradiction that $\tilde{x} \neq x^*$. However, then by the first case of this theorem's proof, $\varphi(x) = x^* = \tilde{x}$ for all $x \in \operatorname{int}\Delta$, which is false. Thus, $\tilde{x} = x^*$. ■

Case 4: $\hat{x}_1 \in \operatorname{int} S_t$**, with** $t \geq 2$. For $x, y \in S_t$, define

$$\wp(x, y) := \{x' \in S_t \mid \exists\ \lambda \in [0,1] : \lambda x' + (1-\lambda) y = x\},$$

i.e., these are the points on the line between x and y on the "opposite" side of x from y. If $S_t \neq \Delta$, observe that for any $x \in \operatorname{int} S_t \setminus (\{x_1\} \cup \wp(\hat{x}_1, x_1))$, we can find a γ_B with binary support $\{x_1, x_2\}$ such that $x \in \operatorname{int}\operatorname{conv}\operatorname{supp}\gamma_B$ (x is a strict convex combination of x_1 and x_2) and $\hat{x}_1 \notin \operatorname{conv}\operatorname{supp}\gamma_B$; and so, by Proposition 3.3, $\varphi(x) = x^*$ for all such x. Moreover by Claim A.7, we must also have $\varphi(x) = x^* = \hat{x}_1$ for all $x \in \operatorname{int} S_t$.

If $S_t = \Delta$, we construct the following two distributions: ρ_B^1, with support on $\{x_1, x_2\}$, where μ is a strict convex combination of x_1 and x_2, and x_2 is close to μ; and ρ_B^2, which has support on n affinely independent points, one of which is in the interior of $\operatorname{supp}\rho_B^1$, and all of which are close to μ. Proposition 3.3 implies that $\varphi(x) = x^*$ for all $x \in \operatorname{int}\operatorname{conv}\operatorname{supp}\rho_B^1$

and for all $x \in \operatorname{int}\operatorname{conv}\operatorname{supp}\rho_B^2$. Provided the support points of ρ_B^2 are sufficiently close to μ, which we assume, $x^*, \hat{x}_1 \notin \operatorname{conv}\operatorname{supp}\rho_B^2$.

Furthermore, for any $x \in \operatorname{int}\Delta \setminus (\wp(x^*,\mu) \cup \wp(\hat{x}_1,\mu) \cup \wp(x_1,\mu))$, we can find a ρ_B^3 with binary support $\{x_3, x_4\}$ such that $x_4 \in \operatorname{int}\operatorname{conv}\operatorname{supp}\rho_B^2$, $x, \mu \in \operatorname{int}\operatorname{conv}\operatorname{supp}\rho_B^3$ (x is a strict convex combination of x_3 and x_4), and $x^*, \hat{x}_1 \notin \operatorname{conv}\operatorname{supp}\rho_B^3$. Consequently, by Proposition 3.3, $\varphi(x) = x^*$ for all such x. Moreover by Claim A.7, we must also have $\varphi(x) = x^* = \hat{x}_1$ for all $x \in \operatorname{int} S_t$.

There is one gap left to fill: what else does respecting the Blackwell order necessitate, when the updating rule makes mistakes on a vertex (or multiple vertices)? Writing $\hat{e}_i := \varphi(e_i)$, we have

Claim A.13. *If an updating rule produces an error for some vertex e_i ($i \in \{1,\ldots,n\}$) $\hat{e}_i \in \ell(x^*, e_i)$.*

Proof. Suppose for the sake of contraposition that $\hat{e}_i$ and $\operatorname{conv}\{e_i, x^*\}$ can be strictly separated by a hyperplane

$$H_{\alpha,\beta} := \left\{x \in \mathbb{R}^{n-1} \,\middle|\, \alpha \cdot x = \beta\right\}.$$

WLOG, we may assume $\operatorname{conv}\{e_i, x^*\} \subset H_{\alpha,\beta}^{\leq} := \left\{x \in \mathbb{R}^{n-1} \,\middle|\, \alpha \cdot x \leq \beta\right\}$, which implies $\hat{e}_i \in \Delta \setminus H_{\alpha,\beta}^{\leq}$.

Consider the value function $V(x) = \max\{0, \alpha \cdot x - \beta\}$, and let ρ_B have support $\{e_i, y\}$ with $p := \mathbb{P}_{\rho_B}(e_i)$. The payoff to the DM from ρ_B is

$$p(\alpha e_i - \beta),$$

which is strictly decreasing in p, yielding a strictly negative value of information. ■

Note that we do not assume that the error is expansive in this claim. It holds regardless of the type of error. ■

A.6 Lemma 4.3 Proof

Proof. Let $\varphi(x) \neq x$ for some $x \in \operatorname{supp}\rho_B$. WLOG we may suppose that ρ_B is binary, since we can always just collapse the points other than the specified x to their barycenter, $\underset{\sim}{x}$. If

ρ_B, does not produce an expansive error and $\varphi(x) \notin \ell_x$ then $\varphi(x) \in \ell_{\underline{x}}$. However, as $\underline{x} \neq \mu$, ρ_B must produce an expansive error, a contradiction. ■

A.7 Proposition 4.5 Proof

Proof. Let ρ_B, corresponding to π, have n affinely independent points of support $\{x_1, \ldots, x_n\}$ and U have a contractive error for one of them, WLOG x_1. Let $p \equiv p_1 \in (0,1)$ denote $\mathbb{P}_{\rho_B}(x_1)$; and let $p_j \in (0,1)$ denote $\mathbb{P}_{\rho_B}\left(x_j\right)$ and $\hat{x}_j := \varphi\left(x_j\right)$ for all j.

Step 1 ("Edge Points Mapped to the Prior"): Consider another Bayesian distribution over posteriors, ρ_B', corresponding to π', with support on $\left\{x_1', x_2, \ldots, x_n\right\}$; that is, all of the support points except for the first one are also support points of ρ_B. Moreover, let $x_1' \in \ell^\circ(x_1, x_s)$ for some $s \neq 1$, so that ρ_B' is a strict MPC of ρ_B and x_1' lies on the edge between x_1 and x_s. Let $p' \in (0,1)$ denote $\mathbb{P}_{\rho_B'}\left(x_1'\right)$. Let $p_j' \in (0,1)$ denote $\mathbb{P}_{\rho_B'}\left(x_j\right)$ for all $j \neq 1$. Note that $p' > p$, $p_s' < p_s$ and $p_j' = p_j$ for all $j \neq 1, s$.

Claim A.14. $\hat{x}_1' := \varphi\left(x_1'\right) = \mu$.

Proof. As U does not produce an expansive error, $\hat{x}_1' \in \ell_{x_1'}$, where possibly $\hat{x}_1' = x_1'$. Suppose for the sake of contraposition that $\hat{x}_1' \neq \mu$. In that case the sets $\operatorname{conv}\left\{\hat{x}_1', x_1', x_1\right\}$ and $\operatorname{conv}\{\hat{x}_1, \mu\}$ can be strictly separated by a hyperplane

$$H_{\alpha,\beta} := \left\{x \in \mathbb{R}^{n-1} \middle| \alpha \cdot x = \beta\right\}.$$

WLOG we may assume that $\operatorname{conv}\left\{\hat{x}_1^1, x_1^1, x_1\right\}$ is a strict subset of the closed half-space $\left\{x \in \mathbb{R}^{n-1} \middle| \alpha \cdot x \geq \beta\right\}$. Consider the value function $V(x) = \max\{0, \alpha \cdot x - \beta\}$. We may ignore points x_j with $j \neq s, 1$. Since we are assuming that there are no expansive errors, we may WLOG assume $\alpha \cdot \hat{x}_s > \beta$ and $\alpha \cdot x_s > \beta$. Thus, the agent's payoff under experiment π is $p_s(\alpha \cdot x_s - \beta)$, and her payoff under experiment π' is $p_s'(\alpha \cdot x_s - \beta) + p_1'\left(\alpha \cdot x_1' - \beta\right)$. Taking the difference of these two expressions, we obtain

$$(p_s - p_s')(\alpha \cdot x_s - \beta) - p_1'(\alpha \cdot x_1' - \beta) = -p_1(\alpha \cdot x_1 - \beta) < 0,$$

as $p_s + p_1 = p_s' + p_1'$ and $p_s x_s + p_1 x_1 = p_s' x_s + p_1' x_1'$. ■

Now another Bayesian distribution over posteriors, $\rho_B^\dagger$, corresponding to $\pi^\dagger$, with support on

$$\left\{x_1, x_2, \ldots, x_s - 1, x_s^\dagger, x_{s+1}, \ldots x_n\right\},$$

where $s \neq 1$. By construction, all of the support points except for $x_s^\dagger$ are also support points of $\rho_B^\dagger$. Moreover, let $x_s^\dagger \in \ell^\circ(x_1, x_s)$, so that $\rho_B^\dagger$ is a strict MPC of ρ_B and $x_s^\dagger$ lies on the edge between x_1 and x_s. Let $p^\dagger \in (0,1)$ denote $\mathbb{P}_{\rho_B^\dagger}\left(x_1^\dagger\right)$. Let $p_j^\dagger \in (0,1)$ denote $\mathbb{P}_{\rho_B^\dagger}\left(x_j\right)$ for all $j \neq 1$. Note that $p^\dagger < p$, $p_s^\dagger > p_s$ and $p_j^\dagger = p_j$ for all $j \neq 1, s$.

Claim A.15. $\hat{x}_s^\dagger := \varphi\left(x_s^\dagger\right) = \mu$.

Proof. Omitted as it is virtually identical to the proof of Claim A.14. ■

Step 2 ("Face Points Mapped to the Prior"): The final step is to show that the points in $\operatorname{int}\operatorname{conv}\operatorname{supp}\rho_B$ must all be mapped to μ by φ.

i. Consider any 2-dimensional face of the simplex $\Delta_{\rho_B} := \operatorname{conv}\operatorname{supp}\rho_B$ for which two (of the three) edges, $S^1_{\rho_B,i}$ and $S^1_{\rho_B,l}$, share vertex x_1. Evidently, any point $x \in \operatorname{int}\Delta_{\rho_B}$ can be obtained as the strict convex combination of points $x_i \in \operatorname{int} S^1_{\rho_B,i}$ and $x_l \in S^1_{\rho_B,l}$. It is easy to see that for all such x, U respecting the Blackwell order (and not producing a contractive error) implies $\varphi(x) = \mu$. If $n = 2$, we are done.

ii. If $n > 2$, consider any 3-dimensional face of Δ_{ρ_B} for which three (of the four) 2-d faces are those specified in i. Following the same logic, any point in the relative interior of this collection of 3-d faces must be mapped to μ by φ. If $n = 3$, we are done.

iii. If $n > 3$, consider any 4-dimensional face...and so on.

This process continues until we arrive at a single face is of maximal dimension (Δ_{ρ_B}), when it terminates, allowing us to conclude the result. ■

A.8 Theorem 5.1 Proof

Proof. First, let $n = 2$. Necessity follows from Propositions 3.6 and 4.6 and Claim A.13. As for sufficiency, we define function $W(x) := \mathbb{E}_{x(\theta)} u(\hat{a}^*, \theta)$, where, understanding $\hat{x} := \varphi(x)$,

$$\hat{a}^* \in \underset{\hat{a} \in \arg\max_{a \in A} \mathbb{E}_{\hat{x}(\theta)} u(a,\theta)}{\arg\max} \mathbb{E}_{x(\theta)} u(\hat{a}, \theta) \ .$$

A positive value for information is implied by W's convexity in x.

By construction, (denoting $\varphi(0) = u$ and $\varphi(1) = v$)

$$W(x) = \begin{cases} \alpha x + \beta, & \text{if} \quad x = 0 \\ \sigma x + \eta, & \text{if} \quad 0 < x \leq a \\ V(x), & \text{if} \quad a < x < b \\ \gamma x + \delta, & \text{if} \quad b \leq x < 1 \\ \tau x + \rho, & \text{if} \quad x = 1, \end{cases}$$

where $\alpha y + \beta \geq \sigma y + \eta$ for all $y \leq u$, $\sigma a + \eta = V(a)$, $\sigma = V'(a^-)$, $V(b) = \gamma b + \delta$, $\gamma = V'(b^+)$, $\tau y + \rho \geq \gamma y + \delta$ for all $y \geq v$, and $V(x)$ is convex.

Second, let $n \geq 3$. Necessity is a consequence of Theorems 3.1 and 4.1. For sufficiency, consider again $W(x) \equiv \max_{\hat{a} \in \arg\max_{a \in A} \mathbb{E}_{\hat{x}(\theta)} u(a,\theta)} \mathbb{E}_{x(\theta)} u(\hat{a}, \theta)$. Observe that, letting s denote the vector parallel to line-segment $\ell^\circ(x^*, e_i^*)$ in the direction of x^* from e_i^*, $D_s(f(x))$ the directional derivative of function f at x along s and $D_s(f(x^-))$ the left-sided directional derivative of function f at x along s,

(i) $W(x) = \alpha \cdot x + \beta$ for all $x \in \operatorname{int} S_m$, for all $S_m \in \hat{\mathcal{S}}^2$, where $V(x^*) = \alpha \cdot x^* + \beta$.

(ii) If $x^* \notin \operatorname{int} S_1$ for some $S_1 \in \hat{\mathcal{S}}^1$, $W(x) = \alpha \cdot x + \beta$ for all $x \in \operatorname{int} S_1$, for all $S_1 \in \hat{\mathcal{S}}^1$.

(iii) If $x^* \in \operatorname{int} S_1'$ for some $S_1' \in \hat{\mathcal{S}}_1$, either

   (a) $W(x) = \alpha \cdot x + \beta$ for all $x \in \ell^\circ(x^*, e_i^*)$, for some vertex e_i^* of S_1', $W(x) = V(x) \geq \alpha \cdot x + \beta$ for all $x \in \operatorname{int} S_1' \setminus \ell^\circ(x^*, e_i^*)$, and $D_s(V(x^{*-})) = D_s(\alpha \cdot x^* + \beta)$; or

   (b) $W(x) = \alpha \cdot x + \beta$ for all $x \in \operatorname{int} S_1'$.

(iv) For all $x \in S_m$, for all $S_m \notin \hat{\mathcal{S}}$, $W(x) = V(x) \geq \alpha \cdot x + \beta$.

(v) For all $e_i \in \hat{\mathcal{E}}$, $W(e_i) \geq \alpha \cdot e_i + \beta$, as for a Bayesian, the regions of beliefs on which actions are optimal are convex.

Thus, W is convex. ■