

Inertial Updating*

Adam Dominiak

Matthew Kovach

Gerelt Tserenjigmid

February 26, 2023

Abstract: We introduce and characterize *inertial updating* of beliefs. Under inertial updating, a decision maker (DM) chooses a belief that minimizes the subjective distance between their prior belief and the set of beliefs consistent with the observed event. Importantly, by varying the subjective notion of distance, inertial updating provides a unifying framework that nests three different types of belief updating: (i) Bayesian updating, (ii) non-Bayesian updating rules such as the $\alpha - \beta$ rule (Grether, 1980), and (iii) updating rules for events with zero probability, like the conditional probability system (CPS) of Myerson (1986a,b). We demonstrate that our model is behaviorally equivalent to the Hypothesis Testing model (HT) of Ortoleva (2012), clarifying the connection between HT and CPS. We apply our model to a persuasion game.

Keywords: Inertial updating, Bayesian updating, non-Bayesian updating, zero-probability events, Bayesian divergence, conditional probability system, hypothesis testing.

JEL: D01, D81.

*Dominiak: Virginia Tech (dominiak@vt.edu); Kovach: Virginia Tech (mkovach@vt.edu); Tserenjigmid: UC Santa Cruz (gtserenj@ucsc.edu). This paper subsumes “Ordered Surprises and Conditional Probability Systems” by the same authors. We are very grateful to David Freeman, Paolo Ghirardato, Faruk Gul, Edi Karni, Shaowei Ke, Yusufcan Masatlioglu, Pietro Ortoleva, Burkhard Schipper, Dong Wei, and Chen Zhao for valuable comments and discussions, as well as the seminar participants at UC Riverside, Texas A&M, ISI Delhi and University of Michigan.

1 Introduction

How decision makers revise their beliefs after receiving information is a foundational problem in economics and game theory. While the benchmark model of Bayesian updating is broadly appealing for a variety of reasons, it has two major issues. First, it is descriptively limited; there is robust experimental evidence that people’s beliefs systematically deviate from what Bayesian updating prescribes.¹ Second, it is *incomplete*; a well-known limitation of Bayesian updating is that it is not defined for zero-probability events.² We resolve these limitations of Bayesian updating by introducing the Inertial Updating (**IU**) representation: a complete theory of belief updating that unifies Bayesian and non-Bayesian updating rules.

The **IU** addresses these two issues by recasting the problem of belief updating as an *optimization problem*; i.e., a problem of belief selection subject to two simple constraints. For each event E , our DM selects a new belief μ_E that is (i) consistent with E and (ii) closest to her initial belief. That is, her new belief μ_E is the element of $\Delta(E)$ that is “closest” to μ among all of the probability measures over E . Importantly, our DM utilizes a subjective notion of distance to determine her conditional belief. Since our DM minimizes the change in her beliefs relative to her prior, we refer to this behavior as Inertial Updating. We provide a complete behavioral analysis of **IU** and demonstrate that it provides a unifying framework for understanding many different forms of belief updating in the literature.

The **IU** representation is characterized by three axioms (see [section 3](#)). The first two postulates are standard: **SEU Postulates** imposes a subjective expected utility representation for each conditional preference, and **Consequentialism** ensures that for any event E , the DM only considers states within E possible. The third axiom, **Dynamic Coherence**, was introduced by Ortoleva (2012) to characterize the Hypothesis Testing model (HT).³ To interpret this axiom, say that an event A is *revealed implied by* event B if every state that the DM believes is possible after learning B is also an element of A . That is, once the DM learns that the “true state is contained in B ,” she is also convinced that that “true state is contained in A ,” and therefore A^c is believed to be null after B . **Dynamic Coherence** requires that this revealed implication over events is acyclic.

Our main theorem shows that the preceding three axioms are necessary and sufficient for the **IU** representation. Our proof is based on an extension of Afriat’s theorem (Afriat (1967)),

¹For experimental evidence, see Kahneman and Tversky (1972), Kahneman and Tversky (1983), Camerer (1987), Eil and Rao (2011), along with surveys by Camerer (1995) and Benjamin (2019).

²This is an especially important issue in sequential games, as particular off-path beliefs are used to support certain equilibria. Accordingly, complete theories of belief updating, such as the Conditional Probability System introduced by Myerson (1986a,b), have been proposed.

³In the HT, an agent’s behavior is in accord with SEU, yet she also has a second-order belief and thus has multiple beliefs in mind. She updates her prior according to Bayes’ rule if she receives “expected” information. When information is “unexpected,” she rejects her prior and uses her second-order belief to select a new belief according to a maximum likelihood rule. This suggests an interpretation of an essentially Bayesian agent who is nevertheless open to fundamentally shifting her worldview.

Varian (1982)) for general budget sets due to Matzkin (1991). We are able to apply this theorem by showing that **Dynamic Coherence** implies that the data set of “belief choices” satisfies the Strong Axiom of Revealed Preferences (SARP). As in the Afriat’s theorem, we get continuity and strict convexity of the distance for free. A corollary of our theorem is that **IU** and HT are behaviorally equivalent, despite their stark difference in appearance and the significantly different proof techniques.

The first key feature of **IU** is that it is descriptively rich; **IU** accommodates Bayesian and non-Bayesian updating. While it is well known that **Dynamic Consistency** ensures Bayesian updating, we provide a complimentary result showing that distance functions that are generalizations of the celebrated Kullback-Leibler (KL) divergence deliver posteriors that are consistent with Bayesian updating. We then build upon this insight and define a family of non-Bayesian updating rules that we call **Distorted Bayesian**. We provide an axiomatic characterization of this form of updating via two axioms. The first axiom, **Consistency**, ensures that the distortion only depends on the probability of the state, while the second axiom, **Independence of Irrelevant Information**, ensures that the relative distortions are event-independent. Consequently, this means that the distortion employed is independent of the realized event.

We then show that a variety of updating biases fall under **Distorted Bayesian**. In particular, **Distorted Bayesian** updating has a non-trivial connection to the well-known $\alpha - \beta$ rule from Grether (1980), capturing forms of under- or over-reaction. The **Distorted Bayesian** can also allow for asymmetric reactions, along with features of confirmation bias. Further, this rule allows for history-dependent updating, and therefore it can capture a wide array of context effects. The characterization of **Distorted Bayesian** and a discussion of the preceding examples can be found in section 3.2.

The second key feature of **IU** is that it is a complete theory of updating: conditional beliefs are well-defined for all events. This follows because the DM’s notion of distance is well-defined for all distributions. Of course, we are not the first to propose a complete theory of updating. The most prominent complete theory is Myerson’s Conditional Probability System (CPS) (Myerson, 1986a,b), which was motivated by the Sequential Equilibria of Kreps and Wilson (1982).

We provide a simple behavioral foundation for CPS in section 4.1. Our characterization relies upon a novel axiom, **Conditional Consistency**, that implies **Dynamic Consistency** among the non-null events and extends this consistency to “conditionally non-null” events. We then show that CPS is a special case of **IU** by providing an explicit distance function that generates any CPS. Because **IU** and HT are behaviorally equivalent, this also establishes that the CPS is a special case of HT.

We extend the **IU** to settings with signal structures and provide an explicit distance function that generates the $\alpha - \beta$ rule from Grether (1980) and its generalizations in section 5.

Using this extension, we illustrate the usefulness of our model by applying it to the Bayesian persuasion games of Kamenica and Gentzkow (2011) (section 6). In particular, we analyze the effect of non-Bayesian belief updating on the optimal signal structure in the Bayesian persuasion games.

2 Model

2.1 Basic Setup

We study choice under uncertainty in the framework of Anscombe and Aumann (1963). A DM faces uncertainty described by a nonempty and finite set of states of nature $S = \{s_1, \dots, s_n\}$.⁴ Let Σ be an arbitrary collection of nonempty subsets of S such that $S \in \Sigma$. Any element E of Σ is called an event. Let X be a nonempty, finite set of outcomes and $\Delta(X)$ be the set of all lotteries over X , i.e., $\Delta(X) := \{p : X \rightarrow [0, 1] \mid \sum_{x \in X} p(x) = 1\}$.

We are interested in a DM's preference over acts. An act is a mapping $f : S \rightarrow \Delta(X)$ that assigns a lottery to each state. Any act f that assigns the same lottery to all states ($f(s) = p$ for all $s \in S$) is called a constant act. Using a standard abuse of notation, we denote by $p \in F$ the corresponding constant act. Hence, we can identify the set of lotteries $\Delta(X)$ with the constant acts. The set of all acts is $\mathcal{F} := \{f : S \rightarrow \Delta(X)\}$. We define mixed lotteries and acts in the usual way: for any $\lambda \in [0, 1]$, $\lambda p + (1 - \lambda)q$ is the lottery providing x with probability $\lambda p(x) + (1 - \lambda)q(x)$, and $\lambda f + (1 - \lambda)g$ is the act that yields $\lambda f(s) + (1 - \lambda)g(s)$ in state s . A preference relation over \mathcal{F} is denoted by \succsim . As usual, \succ and \sim are the asymmetric and symmetric parts of \succsim , respectively. Moreover, fEg denotes that conditional act that returns $f(s)$ for $s \in E$ and $g(s)$ otherwise.

The DM's behavior is depicted by a family of preference relations $\{\succsim_E\}_{E \in \Sigma}$, each defined over \mathcal{F} . We write \succsim in place of \succsim_S , and we call \succsim the initial preference. Say that E is \succsim -null (or simply null) if $fEg \sim g$ for any $f, g \in \mathcal{F}$. Otherwise, E is non-null. Similarly, we say E is \succsim_A -null if $fEg \sim_A g$ for any $f, g \in \mathcal{F}$. If E is not \succsim_A -null, then it is \succsim_A -non-null.

We denote by $\Delta(S)$ the set of all probability distributions on S . For notational convenience, for each $\mu \in \Delta(S)$ and each $s_i \in S$, we will sometimes write μ_i in place of $\mu(s_i)$: the probability of state s_i according to μ . For any $\pi \in \Delta(S)$, let $sp(\pi)$ denote the support of π . For any μ and event E such that $\mu(E) > 0$, let $\text{BU}(\mu, E)$ denote the Bayesian update of μ conditional on E .

Finally, let $\|\cdot\|$ denote the Euclidean norm. For any set A and a function d on A , we write $\arg \min d(A) = \{x \in A \mid d(y) \geq d(x) \text{ for any } y \in A\}$ (whenever this is well-defined).

⁴We focus on a finite state space as it is more standard for decision theoretic analyses and general enough for most economic applications, but we can easily extend our model to an infinite state space.

2.2 Inertial Updating

When the DM observes an event $E \in \Sigma$, she revises her initial preference \succsim conditional on E . The new preference is denoted by \succsim_E and governs the DM's behavior in light of E . This setting is quite general and incorporates the standard signal structure as a special case.⁵ We provide additional analysis of this special case in [section 5](#).

Rather than specify the exact belief our DM utilizes, **IU** only requires that her new belief is (i) consistent with the information and (ii) of minimal distance to her prior. We now formally define our notion of distance.

Definition 1 (Distance Function). A function $d : \Delta(S) \rightarrow \mathbb{R}$ is a **distance function** with respect to $\mu \in \Delta(S)$, denoted by d_μ , if $d_\mu(\mu) < d_\mu(\pi)$ for any $\pi \in \Delta \setminus \{\mu\}$.

We impose only a very weak property; the prior is the global minimizer among all beliefs. This is a simple coherence property, because otherwise a DM should immediately adopt the other belief. Equipped with this notion of distance, we now introduce the **IU** representation.

Definition 2 (IU). A family of preference relations $\{\succsim_E\}_{E \in \Sigma}$ admits an **Inertial Updating** representation if there are a Bernoulli utility function $u : X \rightarrow \mathbb{R}$, a prior $\mu \in \Delta(S)$, a distance function $d_\mu : \Delta(S) \rightarrow \mathbb{R}$ such that for each $E \in \Sigma$, the preference relation \succsim_E admits a SEU representation with (u, μ_E) , meaning that for any $f, g \in \mathcal{F}$,

$$(1) \quad f \succsim_E g \quad \text{if and only if} \quad \sum_{s \in E} \mu_E(s) u(f(s)) \geq \sum_{s \in E} \mu_E(s) u(g(s)),$$

where

$$(2) \quad \mu_E \equiv \arg \min_{\pi \in \Delta(E)} d_\mu(\pi).$$

Since the prior is the global minimizer of d_μ , $\mu = \arg \min_{\pi \in \Delta(S)} d_\mu(\pi)$. For any $E \in \Sigma$, $\Delta(E)$ is convex, $\arg \min_{\pi \in \Delta(E)} d_\mu(\pi)$ will be unique when d_μ is strictly quasi-convex. In fact, the following much weaker condition will suffice: for any $\pi, \pi' \in \Delta(S)$ with $\pi \neq \pi'$, if $d_\mu(\pi) = d_\mu(\pi')$, then there is $\alpha \in (0, 1)$ such that $d_\mu(\alpha\pi + (1 - \alpha)\pi') < d_\mu(\pi)$.

2.3 Notions of Distance

By allowing for a subjective notion of distance, the **IU** generalizes Bayesian updating while also providing a unifying approach to non-Bayesian updating rules. In this section, we discuss a few examples of distance functions and the beliefs they generate. We begin by introducing a Bayesian distance, which will also be useful in defining non-Bayesian distances later.

⁵In particular, when $S = \Omega \times M$, for a set of payoff relevant states Ω and signals M , the signal m corresponds to the event $\{(\omega, m) \in S \mid \omega \in \Omega\}$.

Definition 3 (Bayesian Divergence). For a strictly increasing and strictly concave function $\sigma : \mathbb{R}_+ \rightarrow \mathbb{R}$, let d_μ be given by

$$(3) \quad d_\mu(\pi) = - \sum_{i=1}^n \mu_i \sigma \left(\frac{\pi_i}{\mu_i} \right).$$

Proposition 1. For any non-null $E \in \Sigma$,

$$\mu_E = \arg \min_{\pi \in \Delta(E)} - \sum_{i=1}^n \mu_i \sigma \left(\frac{\pi_i}{\mu_i} \right) = BU(\mu, E)$$

Notably, Equation 3 “includes” the KL divergence as a special case ($\sigma(x) = \ln(x)$). However, since $\ln(0) = -\infty$, the KL divergence is not well-defined when $\text{sp}(\mu) \subseteq \text{sp}(\pi)$. Therefore, we focus our attention on σ that are well defined on \mathbb{R}_+ . For example, $\sigma(x) = \ln(\alpha x + \beta)$ where $\alpha, \beta > 0$ is a well-defined, strictly increasing, and strictly concave function.

We now introduce the following notation to simplify our exposition.

Notation. The **Bayesian function** for a given σ is denoted by $\beta^\sigma : \mathbb{R}_+^n \times \mathbb{R}_+^n \rightarrow \mathbb{R}$; i.e.,

$$\beta^\sigma(\mathbf{x}, \mathbf{y}) = - \sum_{i=1}^n x_i \sigma \left(\frac{y_i}{x_i} \right) \text{ for any } \mathbf{x}, \mathbf{y} \in \mathbb{R}_+^n.$$

Bayesian updating is denoted by $BU : \mathbb{R}_+^n \times \Sigma \rightarrow \mathbb{R}_+^n$; i.e.,

$$BU(\mathbf{x}, E) = \left(\frac{x_i \mathbf{1}\{i \in E\}}{\sum_{j \in E} x_j} \right)_{i \in S} \text{ for any } \mathbf{x} \in \mathbb{R}_+^n.$$

Note that \mathbf{x} and \mathbf{y} are not necessarily probability distributions.

Following the intuition from **Bayesian Divergence**, we can introduce a distorted version of this distance notion to capture non-Bayesian beliefs.

Definition 4 (Distorted Bayesian). An **IU** DM admits a Distorted Bayesian distance if

$$d_\mu(\pi) = \beta^\sigma(\delta(\mu), \pi)$$

where $\delta : [0, 1] \rightarrow \mathbb{R}_+$ and σ is strictly increasing and strictly concave. Then by Proposition 1,

$$(4) \quad \mu_E = BU(\delta(\mu), E)$$

for any non-null $E \in \Sigma$. Further, we say that this distance is **Monotonic** if δ is strictly increasing.

If $\delta > 0$, then we also have $\mu_E = BU(\delta(\mu), E)$ for any E . Moreover, suppose δ is defined as

follows: $\delta(t) = t + \epsilon \mathbb{1}\{t = 0\}$ where ϵ is small enough. Then μ_E is approximately equivalent to $\text{BU}(\mu, E)$ when E is non-null and when E is a null-event, μ_E is equivalent to $\text{BU}(\mu^*, E)$ where μ^* is the uniform distribution over S . This example approximates a special case of Myerson’s CPS introduced in [Definition 6](#).

The **Distorted Bayesian** distance notion captures non-Bayesian updating through the distortion function δ .⁶ Intuitively, such an agent behaves as if they apply Bayes’ rule to a distorted prior. When $\delta(x) = x^\alpha$, this corresponds to a special case of Grether’s $\alpha - \beta$ rule (Grether, 1980) where $\alpha = \beta$. For $\alpha < 1$, this captures under-reaction to information and base-rate neglect, while $\alpha > 1$ captures over-reaction to information. In [section 5](#), we show that our model nests the general version of Grether’s $\alpha - \beta$ rule. It is also straightforward to generalize δ to capture a variety of belief distortions, including asymmetric reactions based on prior beliefs like confirmation bias (à la Rabin and Schrag (1999)) or over(under) reaction to small(large) probabilities (Kahneman and Tversky (1979)).

In [section 3.2](#) we characterize **Distorted Bayesian** and Monotonic **Distorted Bayesian**. Although δ is independent of the realized event, the **Distorted Bayesian** distance can capture features of history or reference dependence.

Definition 5 (Mixed Bayesian). Let d_μ be given by

$$(5) \quad d_\mu(\pi) = \beta^\sigma(\mu + \rho, \pi),$$

where σ is strictly increasing and strictly concave and $\text{sp}(\mu) \cup \text{sp}(\rho) = S$. Then for any $E \in \Sigma$, by [Proposition 1](#),

$$\mu_E = \text{BU}(\mu + \rho, E) = \alpha(E) \text{BU}(\mu, E) + (1 - \alpha(E)) \text{BU}(\rho, E),$$

where $\alpha(E) = \frac{\mu(E)}{\mu(E) + \rho(E)}$. Note that when E is a null-event, $\mu_E = \text{BU}(\rho, E)$. Through ρ , the Mixed Bayesian distance can capture motivated reasoning Kunda (1990) or wishful thinking (Mayraz (2011); Caplin and Leahy (2019); Kovach (2020b)).

IU is a complete updating rule; it is defined for all events, even those which are null. While To illustrate updating rules for zero-probability events, we can define a support-dependent Bayesian divergence.

Definition 6 (Support-Dependent Bayesian Divergence). Let

$$d_\mu(\pi) = \begin{cases} \beta^\sigma(\mu, \pi) & \text{if } \mu(\text{sp}(\pi)) > 0, \\ \beta^\sigma(\mu^*, \pi) + \sigma(1) + |\sigma(0)| & \text{otherwise.} \end{cases}$$

⁶For example, δ captures the DM’s imperfect memory or recall of her previously updated belief – prior (e.g., see Mullainathan (2002), Wilson (2014), Gennaioli and Shleifer (2010), and Bordalo et al. (2016)).

for some μ^* with full-support.

Proposition 2. For any $E \in \Sigma$,

$$\mu_E = \begin{cases} BU(\mu, E) & \text{if } \mu(E) > 0, \\ BU(\mu^*, E) & \text{otherwise.} \end{cases}$$

This distance yields Bayesian updating whenever possible. After a null event, the DM switches to μ^* and then utilizes Bayes' rule. This belief updating rule was used in Galperti (2019), and is a special case of both Myerson (1986a,b) and Ortoleva (2012).

A final example that we wish to mention is the Euclidian distance.

Definition 7 (Euclidean distance). Let $d_\mu(\pi) = \|\mu - \pi\|$. Then

$$\mu_E(s) = \mu(s) + \frac{1 - \mu(E)}{|E|} \text{ for any } E \in \Sigma \text{ and } s \in E.$$

This distance has several nice features. First, it is well-defined for every event, even those that are null. Second, the Euclidean distance is a metric, unlike KL divergence. On the other hand, it is always non-Bayesian and “under utilizes” prior odds when updating beliefs: probability is allocated to the remaining states (i.e., those in E) uniformly. These features echo two consistent findings from experiments: DM's exhibit base-rate neglect (Benjamin, 2019) and are biased toward uniform distributions or the “ignorance prior” (Fox and Clemen, 2005).

3 Axiomatic Characterization

In this section, we present three behavioral postulates that characterize **IU**. Our first axiom imposes the standard SEU conditions of Anscombe and Aumann (1963) on each (conditional) preference relation \succsim_E . Because these conditions are well-understood, we will not provide a formal discussion of the conditions.

AXIOM 1 (SEU Postulates). For each $E \in \Sigma$, the following conditions hold.

- (i) **Weak Order:** \succsim_E is complete and transitive.
- (ii) **Archimedean:** For any $f, g, h \in \mathcal{F}$, if $f \succ_E g$ and $g \succ_E h$, then there are $\alpha, \beta \in (0, 1)$ such that $\alpha f + (1 - \alpha)h \succ_E g$ and $g \succ_E \beta f + (1 - \beta)h$.
- (iii) **Monotonicity:** For any $f, g \in \mathcal{F}$, if $f(s) \succsim_E g(s)$ for each $s \in S$, then $f \succsim_E g$.
- (iv) **Nontriviality:** There are $f, g \in \mathcal{F}$ such that $f \succ_E g$.

(v) **Independence:** For any $f, g, h \in \mathcal{F}$ and $\alpha \in (0, 1]$, $f \succsim_E g$ if and only if $\alpha f + (1 - \alpha)h \succsim_E \alpha g + (1 - \alpha)h$.

(vi) **Invariant Risk Preference:** For all lotteries $p, q \in \Delta(X)$, $p \succsim q$ if and only if $p \succsim_E q$.

The next axiom is standard and ensures that the DM forms a new belief that is consistent with the available information.

AXIOM 2 (**Consequentialism**). For any $E \in \Sigma$ and all $f, g \in F$,

$$f(s) = g(s) \text{ for all } s \in E \implies f \sim_E g.$$

The next axiom, **Dynamic Coherence**, was introduced in Ortleva (2012), and a careful discussion may be found there. In our setting, we say that an event A is *revealed implied* by event B if every state that the DM believes is possible after learning B is also an element of A . **Dynamic Coherence** requires that this “revealed preference” over events is acyclic.

AXIOM 3 (**Dynamic Coherence**). For any $A_1, \dots, A_n \subseteq S$, if $S \setminus A_i$ is $\succsim_{A_{i+1}}$ -null for each $i \leq n - 1$ and $S \setminus A_n$ is \succsim_{A_1} -null, then $\succsim_{A_1} = \succsim_{A_n}$.

If $S \setminus A_i$ is $\succsim_{A_{i+1}}$ -null, then A_i is revealed implied by A_{i+1} . Since **Dynamic Coherence** implies this relation is acyclic, the revealed preference satisfies SARP. Using the result of Matzkin (1991), an extension of Afriat (1967) to general budget sets, SARP is a necessary and sufficient condition for having a subjective distance function for belief selection.

Theorem 1. *The following are equivalent.*

- (i) *A family of preference relations $\{\succsim_E\}_{E \in \Sigma}$ admits an **IU** representation.*
- (ii) *It satisfies **SEU Postulates**, **Consequentialism**, and **Dynamic Coherence**.*
- (iii) *It admits an **IU** representation with respect to a continuous, strictly convex distance function.*

For a simple intuition behind our result, note that **SEU Postulates** and **Consequentialism** imply that our DM has a conditional belief μ_E with support contained in E . Consequently, we may view each event E as generating a “budget set,” $\Delta(E)$, from which the DM must choose her conditional belief. The conditional belief, μ_E , is therefore “revealed preferred” to any other belief in the budget set. **Dynamic Coherence** ensures that this revealed preference satisfies SARP, allowing for the construction of a “utility function” (i.e., a distance function) that generates these beliefs.

Similar to Afriat’s theorem, (i) we obtain a continuous, strictly convex distance function without additional restrictions on preferences; (ii) the above result holds for arbitrary collection

Σ of events. One advantage of our proof is that it is easy to extend to more general models. In Section 7, we consider a generalization of **IU** that satisfies a weakening of **Consequentialism** and the corresponding characterization theorem uses the same generalization of Afriat's theorem.

3.1 Bayesian Updating

Our main theorem does not require **Dynamic Consistency**, and in fact our axioms are independent of this classic postulate. Similar to results from Ghirardato (2002) and Epstein and Breton (1993), imposing **Dynamic Consistency** in our setting ensures that conditional beliefs are consistent with Bayesian updating whenever possible. Recall that fEh denotes that conditional act that returns $f(s)$ for $s \in E$ and $h(s)$ otherwise.

AXIOM 4 (**Dynamic Consistency**). For all non-null events $E \in \Sigma$ and $f, g, h \in \mathcal{F}$,

$$fEh \succsim gEh \text{ if and only if } f \succsim_E g.$$

Proposition 3. *A family of preference relations $\{\succsim_E\}_{E \in \Sigma}$ satisfies **SEU Postulates**, **Consequentialism**, **Dynamic Coherence**, and **Dynamic Consistency** if and only if it admits an **IU** representation and $\mu_E = BU(\mu, E)$ for each non-null E .*

Since **Dynamic Consistency** has been discussed extensively, (both Ghirardato (2002) and Epstein and Breton (1993) include excellent discussions), we will not discuss this result further. Instead, we simply wish to remark that **Dynamic Consistency** places no restrictions on conditional beliefs after null events, which is a major drawback of the standard model.

A strength of **IU** is that it provides a coherent framework for belief revision after null events, which we discuss in section 4. Notably, in section 4.1 we introduce a strengthening of **Dynamic Consistency**, which we call **Conditional Consistency**, that extends the logic of **Dynamic Consistency** to all conditional events and show that this condition characterizes the CPS of Myerson (1986a,b).

3.2 Distorted Bayesian Updating

One of the key insights provided by **IU** is that distance minimization can be viewed as a unifying framework that accommodates many updating behaviors. In this section, we expand upon this insight by characterizing several special cases of **IU** that allow for non-Bayesian updating.

AXIOM 5 (**Consistency**). For any $E \in \Sigma$, $s, s' \in E$, and $x, y \in X$,

$$x\{s\}y \sim x\{s'\}y \text{ if and only if } x\{s\}y \sim_E x\{s'\}y.$$

Consistency requires that if the DM initially thinks that two states are equally likely, then she still thinks they are equally likely after learning an event that contains both of them.

We characterize **Distorted Bayesian** with one additional condition that we call **Independence of Irrelevant Information**. This axiom ensures that updating behavior only depends on the probability of a state and not on the name of the state. Further, this condition also ensures that the relative distortions are independent of the realized event.

AXIOM 6 (Independence of Irrelevant Information). For any $E_1, E_2 \in \Sigma \setminus S$, $s, s' \in E_1 \cap E_2$, and $p, q, r \in \Delta(X)$,

$$p\{s\}r \sim_{E_1} q\{s'\}r \text{ implies } p\{s\}r \sim_{E_2} q\{s'\}r.$$

Proposition 4. *Consider a family of preference relations $\{\succsim_E\}_{E \in \Sigma}$ with an **IU** representation. The **IU** representation admits a **Distorted Bayesian** distance if and only if **Consistency** and **Independence of Irrelevant Information** hold.*

We can now characterize Monotonic **Distorted Bayesian** distance by introducing a condition ensuring that the DM preserves the “more likely than” judgments implied by her prior.

AXIOM 7 (Monotonicity). For any $E \in \Sigma$, $s, s' \in E$, and $x, y \in X$,

$$x\{s\}y \succsim x\{s'\}y \text{ if and only if } x\{s\}y \succsim_E x\{s'\}y.$$

To understand **Monotonicity**, consider $S = \{s_1, s_2, s_3\}$, $\mu = (12/20, 7/20, 1/20)$, and $E = \{s_2, s_3\}$. Under **Dynamic Consistency**, relative likelihoods are exactly preserved and so a Bayesian DM continues to believe that s_2 is seven times as likely as s_3 upon learning E . Without **Dynamic Consistency**, the **IU** would place no restrictions on the conditional relative likelihoods of s_2 and s_3 . Since our DM believed that E was relatively unlikely, it is plausible that she is now less confident in her judgment about the relative odds of s_2 and s_3 . Consequently, she may desire to further modify her belief so that s_3 becomes more likely. For example, she may now think that s_2 is only twice as likely as s_3 , resulting in the posterior $\mu_E = (2/3, 1/3)$. Notice that s_2 is still more likely than s_3 ; she does not entirely disregard her previous judgments. This restriction is precisely the content of **Monotonicity**.

Proposition 5. *Consider a family of preference relations $\{\succsim_E\}_{E \in \Sigma}$ with an **IU** representation. The **IU** representation admits a Monotonic Distorted Bayesian distance if and only if **Monotonicity** and **Independence of Irrelevant Information** hold.*

For each of the following examples, we let $S = \{s_1, s_2, s_3\}$, and suppose $\mu = (12/20, 7/20, 1/20)$. In each of the tables, blue (light) shading indicates that the state is under-weighted relative to Bayes’ rule, while red (dark) shading indicates the state is over-weighted.

Example 1 (Bayesian). Our Distorted Bayesian model includes Bayesian updating as the special case $\delta(x) = x$. These posteriors are given in the table below and will serve as the benchmark to describe our other examples.

s/A	$\{s_1, s_2\}$	$\{s_2, s_3\}$	$\{s_1, s_3\}$
s_1	0.63	—	0.92
s_2	0.37	0.875	—
s_3	—	0.125	0.08

Table 1: Bayes’ Posteriors for various events

Example 2 (Under/Over-Reaction). Suppose for some $\alpha > 0$,

$$\delta(x) = x^\alpha.$$

Note that for $\alpha = 1$ this reduces to Bayes’ rule (see Table 1). For $\alpha < 1$, the relative probabilities are “compressed,” capturing under-reaction to the higher probability state. On the other hand, when $\alpha > 1$, relative probabilities are “exaggerated,” capturing over-reaction to the higher probability state.

s/A	$\{s_1, s_2\}$	$\{s_2, s_3\}$	$\{s_1, s_3\}$
s_1	0.606	—	0.88
s_2	0.394	0.826	—
s_3	—	0.174	0.12

s/A	$\{s_1, s_2\}$	$\{s_2, s_3\}$	$\{s_1, s_3\}$
s_1	0.656	—	0.95
s_2	0.344	0.912	—
s_3	—	0.088	0.05

Table 2: Posteriors for $\alpha = 0.8$ and $\alpha = 1.2$.

Comparing to the Bayesian posteriors in Table 1, it is simple to see that when $\alpha < 1$ the DM always under-weights the more likely state, and when $\alpha > 1$ the DM always over-weights the more likely state.

Example 3 (S -reaction). When δ has a sigmoid shape, it simultaneously captures under-reaction to “expected states” and over-reaction to “unexpected states.” For some $x_0 \in \mathbb{R}$ and $a > 0$,

$$\delta(x) = \frac{1}{1 + e^{a(x_0 - x)}}.$$

s/A	$\{s_1, s_2\}$	$\{s_2, s_3\}$	$\{s_1, s_3\}$
s_1	0.69	—	0.91
s_2	0.31	0.821	—
s_3	—	0.179	0.09

Table 3: Posteriors for S -reaction with $a = 6, x_0 = 0.5$.

Compared to the Bayesian posteriors, the DM over-weights s_1 after $\{s_1, s_2\}$, exhibiting features of over-reaction, while the DM under-weights s_1 after $\{s_1, s_3\}$ and s_2 in after $\{s_2, s_3\}$.

This is because the shape of δ induces over-reaction to rare events, thereby increasing the probability of s_3 .

Example 4 (Confirmation Bias). Confirmation bias refers to the tendency to give extra credence to “believed hypothesis.” For some $b > 0$, let

$$\delta(x) = x + b \mathbb{1} \left\{ x > \frac{1}{2} \right\}.$$

Under this rule, states which are believed to be more likely are biased by b .

s/A	$\{s_1, s_2\}$	$\{s_2, s_3\}$	$\{s_1, s_3\}$
s_1	$\frac{12+20b}{19+20b}$	—	$\frac{12+20b}{13+20b}$
s_2	$\frac{7}{19+20b}$.875	—
s_3	—	.125	$\frac{1}{13+20b}$

Table 4: Posteriors under Confirmation Bias.

The DM always over-reacts to s_1 , her favored state, whenever information allows. When the information precludes s_1 she behaves in accordance with Bayes’ rule.

4 Updating After Zero-probability Events

The most well-known weakness of Bayesian updating is that it is incomplete; it is not defined for zero-probability events. This is particularly problematic in game theoretic settings, where beliefs are induced by the equilibrium strategies and any action off the equilibrium path leads to a zero-probability event. In contrast, our notion of belief updating is well-defined for zero-probability events. Thus, **IU** provides a way to extend (non-)Bayesian updating to all events.

4.1 Conditional Probability System

Perhaps the most well-known method for handling choice conditional on null-events is the conditional probability system (CPS) introduced by Myerson (1986a,b). The development of CPS is closely related the developments of Perfect Bayesian Equilibrium and its refinements. PBE requires that agents’s beliefs are Bayes-consistent with the prior whenever possible. However, PBE does not make any restrictions when Bayes’ rule is not applicable. Hence, PBE may allow for some unreasonable beliefs for actions off the equilibrium path. The Sequential Equilibria of Kreps and Wilson (1982) refines the PBE by requiring that any belief in sequential equilibria should be a limit of full-support beliefs after applying Bayes rule accordingly. Checking whether conditional beliefs can be supported by full-support beliefs is not easy task and Myerson (1986a,b) shows that this limit requirement of sequential equilibria is equivalent to the following simple condition.

Definition 8. A **conditional probability system** is a collection $\{\mu_E\}_{E \in \Sigma}$ of conditional probability distributions such that for all $s \in F \subseteq E$,

$$(6) \quad \mu_E(s) = \mu_F(s) \mu_E(F).$$

When $\mu_E(F) \neq 0$, [Equation 6](#) reduces to Bayes' rule. However, when $\mu_E(F) = 0$, it implies that $\mu_E(s) = 0$ as well, and so it places no restriction directly on $\mu_F(s)$.

As we will show below, CPS is a special case of our model. A major distinction between CPS and **IU** is that CPS requires Bayesian updating whenever possible, while **IU** provides a unifying framework that allows for Bayesian and non-Bayesian updating. To characterize CPS, we introduce the following strengthening of **Dynamic Consistency**.

AXIOM 8 (Conditional Consistency). For all $E \in \Sigma$, \succsim_E -non-null $A \subset E$, and $f, g, h \in \mathcal{F}$,

$$f A h \succsim_E g A h \text{ if and only if } f \succsim_A g.$$

Conditional Consistency implies **Dynamic Consistency** but also has bite on events that are \succsim -null. In essence, **Conditional Consistency** extends the logic of **Dynamic Consistency** to all conditional preferences E and nested events that are \succsim_E -non-null.

To illustrate **Conditional Consistency**, imagine a coin flip. The states h and t are the usual outcomes of heads or tails, e and e' denote edges where e' has been worn thin, while l_1 and l_2 denote landing on a marked location, which yields the state space $S = \{h, t, e, e', l_1, l_2\}$. Initially, the DM believes that $\mu(h) = \mu(t) = \frac{1}{2}$, and treats the other states as null.

Suppose the DM is informed that, astonishingly, the coin did not land on a face; $A = \{e, e', l_1, l_2\}$ was realized. Further, suppose that our DM believes that the coin landing on either of the marked locations is *more impossible* than its landing on an edge. Accordingly, her conditional beliefs are $\mu_A(\{e, e'\}) = 1$ and $\mu_A(\{l_1, l_2\}) = 0$. If this information is further refined so that $\{e\}$ is ruled out and our DM continues to utilize Bayes' rule, then we expect $\mu_B(\{e'\}) = 1$ (where $B = \{e', l_1, l_2\}$). **Conditional Consistency** imposes Dynamic Consistency between μ_A and μ_B because B becomes \succsim_A -non-null and $B \subset A$.

Our next theorem states that **Conditional Consistency** is the precise strengthening of **Dynamic Consistency** required to characterize CPS.

Theorem 2. *A family of preference relations $\{\succsim_E\}_{E \in \Sigma}$ satisfies **SEU Postulates**, **Consequentialism**, and **Conditional Consistency** if and only if it admits a **CPS** representation.*

While our theorem ensures that the collection of beliefs satisfies the requirement of a CPS ([Definition 8](#)), it does not directly shed light on the structure of the CPS. Our next proposition shows that any CPS can be described by a collection of beliefs whose supports partition S . Further, the DM moves between these beliefs in an “ordered” fashion and this CPS representation is generated by a support-dependent bayesian distance.

Proposition 6. *Suppose a family of preferences $\{\succsim_E\}_{E \in \Sigma}$ admits a CPS representation. Then there are $\mu^0, \dots, \mu^K \in \Delta(S)$ such that $sp(\mu^0), \dots, sp(\mu^K)$ is a partition of S and for any $E \in \Sigma$,*

$$\mu_E = BU(\mu^{k^*}, E) \text{ where } k^* = \min\{k \mid \mu^k(E) > 0\}.$$

Moreover, $\{\succsim_E\}_{E \in \Sigma}$ has an **IU** representation with respect to the following distance function:

$$d_\mu(\pi) = \beta^\sigma(\mu_i^{k^*}, \pi) + k^* (\sigma(1) + |\sigma(0)|)$$

where $k^* = \min\{k \mid \mu^k(sp(\pi)) > 0\}$.⁷

Note that Proposition 2 is a special case of the above result when $K = 1$.

Example 5 (Coin Flip). Recall the coin flip example from before, where the states are $S = \{h, t, e, e', l_1, l_2\}$, where h and t correspond to heads or tails, e and e' correspond to the coin landing on an edge, where one edge is thinner than the other, while l_1 and l_2 correspond to the coin landing on precisely marked locations. These possibilities are described by the probability distributions

$$\mu^0(s) = \begin{cases} \frac{1}{2} & s \in \{h, t\} \\ 0 & \text{otherwise} \end{cases}; \mu^1(s) = \begin{cases} \frac{7}{8} & s = e \\ \frac{1}{8} & s = e' \\ 0 & \text{otherwise} \end{cases}; \text{ and } \mu^2(s) = \begin{cases} \frac{1}{2} & s \in \{l_1, l_2\} \\ 0 & \text{otherwise} \end{cases}.$$

Our DM has the initial prior μ^0 (i.e., \succsim has an SEU representation with (u, μ^0)). Suppose she observes $A = \{e, e', l_1, l_2\}$. Since $\mu^0(A) = 0$, Bayesian updating is not defined. After A , the DM selects μ^1 (i.e., $\mu_A = \mu^1$) because it is of “lower order” than μ^2 and therefore it takes precedence.

4.2 Hypothesis Testing

A recent and elegant addition to the literature on updating after zero-probability events is the Hypothesis Testing model (HT) of Ortoleva (2012). Such an agent will update using Bayes’ rule for expected events: events with probability above some threshold ϵ . When an event E is unexpected (i.e., under the agent’s prior $\mu(E) \leq \epsilon$), the agent rejects her prior, updates a second-order prior over beliefs, and selects a new belief according to a maximum likelihood procedure. Formally, a HT representation is given by a triple, (μ, ρ, ϵ) , consisting of a prior $\mu \in \Delta(S)$, a second order prior $\rho \in \Delta(\Delta(S))$, and a threshold $\epsilon \in [0, 1)$ with the requirement

⁷The first part this proposition is not entirely new. Kreps and Wilson (1982) already points out that a connection between sequential equilibria beliefs and a collection of linearly ordered priors μ^0, \dots, μ^K .

that $\mu = \arg \max_{\pi \in \Delta(S)} \rho(\pi)$. Then, for any $E \in \Sigma$,

$$\mu_E = \begin{cases} \text{BU}(\mu, E) & \text{if } \mu(E) > \epsilon, \\ \text{BU}(\pi^E, E) & \text{otherwise.} \end{cases}$$

where $\pi^E = \arg \max_{\pi \in \Delta(S)} \rho(\pi)\pi(E)$. It turns out that HT is behaviorally equivalent to **IU**.

Corollary 1. A family of preference relations $\{\succsim_E\}_{E \in \Sigma}$ admits an HT representation if and only if it admits an **IU** representation.

This corollary immediately follows from our [Theorem 1](#) and Theorem 1 of Ortoleva (2012). However, it is important to note that our proofs are quite different.

4.3 Relating HT and CPS

The formal relationship between HT and CPS has not previously been established. Our results, [Corollary 1](#) and [Proposition 6](#), indirectly show that CPS is a special case of HT. Further, since every CPS satisfies Bayes' rule, it is a special case of HT with $\epsilon = 0$.

Corollary 2. For each family of $\{\succsim_E\}_{E \in \Sigma}$ CPS preferences, there exists an equivalent family of HT preferences with $\epsilon = 0$.

However, these models are not equivalent; even when $\epsilon = 0$, HT preferences may be inconsistent with CPS preferences. The reason for this is due to the way in which the selection of new beliefs occurs in HT. The following example illustrates this distinction by constructing an HT that is not consistent with any CPS.

Example 6. Recall our coin example ([Example 5](#)), but suppose instead that the DM has a HT representation with $\epsilon = 0$, $\rho(\mu^0) = \frac{1}{2}$, $\rho(\mu^1) = \frac{1}{3}$, and $\rho(\mu^2) = \frac{1}{6}$. When $A = \{e', l_1, l_2\}$,

$$\mu^1(A)\rho(\mu^1) = \frac{1}{8} \cdot \frac{1}{3} < \frac{1}{6} = \mu^2(A)\rho(\mu^2),$$

since $\mu^1(A) = \frac{1}{8}$ and $\mu^2(A) = 1$. Thus, $\mu_A^{HT} = \text{BU}(\mu^2, A)$. However, when $E = \{e, e', l_1, l_2\}$,

$$\mu^1(E)\rho(\mu^1) = \frac{1}{3} > \frac{1}{6} = \mu^2(E)\rho(\mu^2)$$

since $\mu^1(E) = \mu^2(E) = 1$. Thus, $\mu_E^{HT} = \text{BU}(\mu^1, E)$. Hence, the described HT representation does not admit an OS representation because

$$\mu_E^{HT}(e') = \frac{1}{8} \neq 0 \cdot \frac{1}{8} = \mu_A^{HT}(e') \cdot \mu_E^{HT}(A),$$

violates (6) in [Definition 8](#).

The relationship between CPS and HT is illustrated in [Figure 1](#).

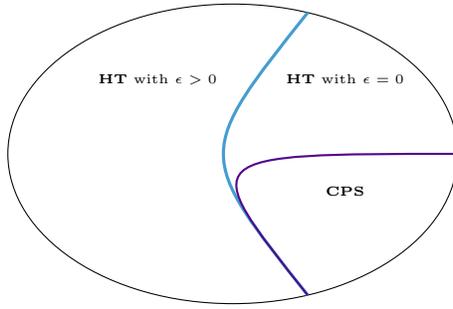


Figure 1: Relationship between CPS and HT.

4.4 A Non-Bayesian CPS

A natural way to generalize the CPS is to retain the sequential selection of new beliefs while incorporating the idea of “non-Bayesian reaction to unexpected events” from the HT model. To do so, we introduce ϵ -CPS, a one-parameter, non-Bayesian extension of the CPS. This extension may lead to an interesting, non-Bayesian generalization of sequential equilibria.

Definition 9. A family of preferences $\{\succsim_E\}_{E \in \Sigma}$ admits an ϵ -CPS representation if there are probability distributions $\mu_0, \dots, \mu_K \in \Delta(S)$ and $\epsilon \in [0, 1)$ such that

$$\mu_E = \text{BU}(\mu_{k^*}, E) \text{ where } k^* = \min\{k \leq K \mid \mu_k(E) > \epsilon\},$$

for every $E \in \Sigma$.

The ϵ -CPS representation incorporates the key idea of HT by allowing for non-Bayesian reactions to unexpected events: $\mu_j(E) \leq \epsilon$. However, it provides additional structure to the posterior selection process. The ϵ -CPS remains a special case of HT.

Theorem 3. *Any ϵ -CPS representation also has a HT representation. Moreover, if $\epsilon = 0$, then the threshold for the HT representation is also zero.*

5 Incorporating a Signal Structure

While our setting is quite general, it is often useful to make explicit reference to a signal structure. We therefore illustrate that our framework can incorporate standard signal structures utilized in experimental settings and game theory by introducing more structure to the state space (e.g., S has a product structure).

Let Ω be the payoff relevant state space and M be the set of all signals. For each $\omega \in \Omega$ and $m \in M$, let $P(\omega)$ be the (unconditional) probability that the payoff relevant state ω occurs and $P(m|\omega)$ be the (conditional) probability that the DM receives the signal m when the state is ω . Indeed, learning a signal will be equivalent to learning an event in an expanded state

space, $S = \Omega \times M$. Specifically, receiving the signal m is equivalent to learning the event $\{(\omega, m)\}_{\omega \in \Omega}$ in S .

Let μ be the prior on S , so that $\mu_{\omega m} = P(m|\omega)P(\omega)$ for each (ω, m) . In the case of Bayesian updating, the connection between our framework and the signal structure is straightforward. Note that the KL divergence generates Bayesian updating in the signal structure framework:

$$P(\omega|m) = \frac{\mu_{\omega m}}{\sum_{\omega' \in \Omega} \mu_{\omega' m}} = \frac{P(m|\omega)P(\omega)}{\sum_{\omega' \in \Omega} P(m|\omega')P(\omega')}.$$

A similar connection is possible for non-Bayesian updating rules. For example, consider now the following distance function. For $\alpha, \beta \geq 0$,

$$d_\mu(\pi) = \sum_{(\omega, m) \in \Omega \times M} \left(\sum_{m' \in M} \mu_{\omega m'} \right)^{\alpha - \beta} \mu_{\omega m}^\beta \log \left(\frac{\pi_{\omega m}}{\mu_{\omega m}} \right).$$

This distance generates the posterior

$$P(\omega|m) = \frac{\left(\sum_{m' \in M} \mu_{\omega m'} \right)^{\alpha - \beta} \mu_{\omega m}^\beta}{\sum_{\omega' \in \Omega} \left(\sum_{m' \in M} \mu_{\omega' m'} \right)^{\alpha - \beta} \mu_{\omega' m}^\beta} = \frac{(P(m|\omega))^\beta (P(\omega))^\alpha}{\sum_{\omega' \in \Omega} (P(m|\omega'))^\beta (P(\omega'))^\alpha},$$

which is precisely the non-Bayesian updating rule proposed by Grether (1980). This is a simple generalization of Bayes' rule, where α captures the influence of the prior and β captures the influence use of the signals.

In general, the following distance function

$$d_\mu(\pi) = \sum_{(\omega, m) \in \Omega \times M} g \left(\sum_{m' \in M} \mu_{\omega m'} \right) f \left(\frac{\mu_{\omega m}}{\sum_{m' \in M} \mu_{\omega m'}} \right) \log \left(\frac{\pi_{\omega m}}{\mu_{\omega m}} \right)$$

generates Distorted Bayesian updating in the signal structure framework:

$$P(\omega|m) = \frac{f \left(\sum_{m' \in M} \mu_{\omega m'} \right) g \left(\frac{\mu_{\omega m}}{\sum_{m' \in M} \mu_{\omega m'}} \right)}{\sum_{\omega' \in \Omega} f \left(\sum_{m' \in M} \mu_{\omega' m'} \right) g \left(\frac{\mu_{\omega' m}}{\sum_{m' \in M} \mu_{\omega' m'}} \right)} = \frac{f(P(m|\omega))g(P(\omega))}{\sum_{\omega' \in \Omega} f(P(m|\omega'))g(P(\omega'))}.$$

The above updating rule reduces to Grether's rule when $f(x) = x^\beta$ and $g(x) = x^\alpha$. We apply this updating rule to the Bayesian persuasion games in Section 6.

Example 7. Consider the following example, with $\Omega = \{\omega_H, \omega_L\}$ and $M = \{h, l\}$. We suppose $P(\omega_H) = \frac{5}{8}$ and $P(h|\omega_H) = P(l|\omega_L) = \frac{3}{5}$.

Table 5 illustrates the prior over S and the resulting posterior beliefs under Bayesian updating. Applying the Distorted Bayesian distance yields the following conditional probabilities for ω_H after signals h and l :

	h	l		$\mu(\cdot h)$	$\mu(\cdot l)$
ω_H	0.375	0.25	ω_H	0.7143	0.5263
ω_L	0.15	0.225	ω_L	0.2857	0.4737

Table 5: Induced prior μ over $S = \Omega \times M$ and corresponding Bayes' posteriors.

$$\mu(\omega_H|h) = \frac{f(0.6)g(0.625)}{f(0.6)g(0.625) + f(0.4)g(0.375)},$$

$$\mu(\omega_H|l) = \frac{f(0.4)g(0.625)}{f(0.4)g(0.625) + f(0.6)g(0.375)}.$$

To further illustrate, we consider several specifications for $f(x) = x^\beta$ and $g(x) = x^\alpha$ in the table below.

	$\mu(\cdot h)$	$\mu(\cdot l)$		$\mu(\cdot h)$	$\mu(\cdot l)$		$\mu(\cdot h)$	$\mu(\cdot l)$
ω_H	0.7264	0.4603	ω_H	0.6755	0.5211	ω_H	0.7347	0.5517
ω_L	0.2736	0.5397	ω_L	0.3245	0.4789	ω_L	0.2653	0.4483

Table 6: The left table reports posteriors with Base-rate neglect ($\alpha = 0.8$) and over-reaction to signals ($\beta = 1.4$), the middle table reports posteriors with Base-rate neglect ($\alpha = 0.8$) and under-reaction to signals ($\beta = 0.8$), and the right table reports posteriors with Base-rate bias ($\alpha = 1.2$) and an accurate reaction to signals ($\beta = 1$).

6 Application to Bayesian Persuasion

In this section, we demonstrate the usefulness of our model by applying it to the Bayesian persuasion games of Kamenica and Gentzkow (2011). In particular, we analyze the effects of non-Bayesian updating rules on the optimal information structure. We first describe the general Bayesian persuasion environment. Let Ω be the set of payoff-relevant states and ρ be a prior over $\Omega = \{\omega_1, \dots, \omega_n\}$. Let A and M be the sets of actions and messages, respectively. A signal structure is a function $\pi : \Omega \rightarrow \Delta(M)$. Given action $a \in A$ and state $\omega \in \Omega$, the receiver's payoff is $u(a, \omega)$ and the sender's payoff is $v(a, \omega)$. Given message realization m and signal structure π , the receiver's optimal action is determined by

$$a_\pi^*(m) = \arg \max_{a \in A} \mathbb{E}_{\omega \sim \mu_\pi(\cdot|m)} u(a, \omega)$$

where $\mu_\pi(\cdot|m)$ is a conditional probability distribution over Ω . The sender's goal is to persuade the receiver to take certain actions by choosing a signal structure π . The optimal signal structure for the sender must solve

$$\max_{\pi \in \Pi} \mathbb{E}_{\omega \sim \rho} \mathbb{E}_{m \sim \pi(\omega)} v(a_\pi^*(m), \omega).$$

To illustrate the implications of our model, we now consider a simpler environment with two actions; $A = \{a, b\}$. Then

$$a_\pi^*(m) = a \text{ if } \sum_{i=1}^n \mu_\pi(\omega_i|m) u_i \geq 0,$$

where $u_i = u(a, \omega_i) - u(b, \omega_i)$. We assume that the sender always prefers action a , which is captured by $v(a, \omega) = 1$ and $v(b, \omega) = 0$. Hence, the sender maximizes

$$\sum_{m \in M} \sum_{\omega \in \Omega} \mu(\omega, m) \mathbb{1}\{a^*(m) = a\}.$$

This simple environment is rich enough to nest the judge-prosecutor example of Kamenica and Gentzkow (2011) and the police-driver example of Kamenica (2019). Since $\min\{|\Omega|, |A|\} = 2$, we will assume that $|M| = 2$.⁸

To apply our model, let $S = \Omega \times M$ and let μ be a prior over S determined by ρ and π : $\mu(\omega_i, m) = \rho_i \pi_m(\omega_i)$.⁹ Our model determines the conditional probability $\mu_\pi(\cdot|m)$ and the rest is standard. We assume biased Bayesian updating defined in section 5, where the conditional probability is given by

$$\mu_\pi(\omega_i|m) = \frac{g(\rho_i) f(\pi_m(\omega_i))}{\sum_{j=1}^n g(\rho_j) f(\pi_m(\omega_j))}.$$

We assume that f and g are positive valued and f is strictly increasing.¹⁰ Below we demonstrate how the curvature of f determines the form of the optimal signal structure for the sender.

Given this updating rule, the optimal action of the receiver is

$$a_\pi^*(m) = a \text{ iff } \sum_{i=1}^n g(\rho_i) f(\pi_m(\omega_i)) u_i \geq 0.$$

The sender's optimization problem is

$$\max_{\pi \in \Pi} V(\pi) = \sum_{m \in M} \left(\sum_{i=1}^n \rho_i \pi_m(\omega_i) \right) \mathbb{1}\left\{ \sum_{i=1}^n g(\rho_i) f(\pi_m(\omega_i)) u_i \geq 0 \right\}.$$

To simplify the exposition, we first rule out some uninteresting scenarios in which persuasion does not matter. Note that the maximum value for V is 2. To focus on the interesting cases, suppose now that we have u_1, \dots, u_n and ρ such that $V = 2$ cannot be achieved. We also assume that the fully revealing signal structure is not optimal; i.e., $\sum_{i=1}^n g(\rho_i) u_i < 0$. Since messages m_1 and m_2 are symmetric, we will focus on signal structures such that $a_\pi^*(m_1) = a$.

⁸We are restricting attention to “straightforward” signal structures. See de Clippel and Zhang (2022).

⁹Bayesian plausibility is already satisfied with this Cartesian structure.

¹⁰Although it is not essential, for simplicity, we assume f is differentiable.

Then the sender's problem is simply to maximize

$$\max_{\pi_{m_1} \in [0,1]^{|\Omega|}} \sum_{i=1}^n \rho_i \pi_{m_1}(\omega_i) \text{ subject to } \sum_{i=1}^n g(\rho_i) f(\pi_{m_1}(\omega_i)) u_i \geq 0.$$

Since the objective function is linear, the curvature of f essentially dictates the form of the optimal signal structure.

Proposition 7. *Suppose either $f(x) = x$ and $\frac{g(\rho_i)u_i}{\rho_i} \neq \frac{g(\rho_j)u_j}{\rho_j}$ for any i, j with $u_i, u_j < 0$ or f is strictly concave. For any optimal signal structure π^* , there is $\bar{\omega} \in \Omega$ and non-empty $\Omega_1 \subseteq \Omega \setminus \{\bar{\omega}\}$ such that $\pi_{m_1}^*(\bar{\omega}) \in (0, 1)$ and*

$$\pi_{m_1}^*(\omega_1) = 1 \text{ for any } \omega_1 \in \Omega_1 \text{ and } \pi_{m_1}^*(\omega_2) = 0 \text{ for any } \omega_2 \in \Omega \setminus (\Omega_1 \cup \{\bar{\omega}\}).$$

The proposition essentially shows that the sender randomizes at a single state when f is concave. However, when f is not concave, in particular when $f'(0) = 0$, the sender randomizes at as many as states as possible.

Proposition 8. *Suppose $f'(0) = 0$. For any optimal signal structure π^* , there is a partition Ω_1, Ω_2 of Ω such that*

$$\pi_{m_1}(\omega_1) = 1 \text{ for any } \omega_1 \in \Omega_1 \text{ and } \pi_{m_1}(\omega_2) \in (0, 1) \text{ for any } \omega_2 \in \Omega_2.$$

The above propositions show that the optimal signal structure in the case of $\beta \leq 1$ (including Bayesian updating) is qualitatively different from the case of $\beta > 1$. The following example illustrates the difference between the cases $\beta \leq 1$ and $\beta > 1$.

Example. Suppose $|\Omega| = 3$. Let $\rho = (\frac{1}{7}, \frac{3}{7}, \frac{3}{7})$ and $(u_1, u_2, u_3) = (1, -\frac{1}{2}, -1)$. Then when $\beta \leq 1$,

$$\pi_{m_1}^H = 1, \pi_{m_1}^M = \left(\frac{2}{3}\right)^{\frac{1}{\beta}} \in (0, 1), \text{ and } \pi_{m_1}^L = 0.$$

However, when $\beta > 1$,

$$\pi_{m_1}^H = 1, \pi_{m_1}^M = \frac{2^{\frac{1}{\beta-1}}}{\left(3(1+2^{\beta-1})\right)^{\frac{1}{\beta}}} \in (0, 1), \text{ and } \pi_{m_1}^L = \frac{1}{\left(3(1+2^{\beta-1})\right)^{\frac{1}{\beta}}} \in (0, 1).$$

7 Partial Consequentialism and Weighted IU

In this section, we generalize our main result by relaxing **Consequentialism**. Following our analogy to revealed preference theory, **Consequentialism** ensures that E is equivalent to the budget set $\Delta(E)$. By dropping **Consequentialism**, we allow for the DM to perceive a subjective budget set from which she may choose. For instance, this may be because the DM

perceives the information as less reliable than the analyst, or the DM may have an imperfect memory and her uncertainty about which event transpired is reflected in her beliefs. We do however impose two natural conditions on her behavior.

Definition 10 (wIU). A family of preference relations $\{\succsim_E\}_{E \in \Sigma}$ admits a **Weighted Inertial Updating** representation if there are a Bernoulli utility function $u : X \rightarrow \mathbb{R}$, a prior $\mu \in \Delta(S)$, a distance function $d_\mu : \Delta(S) \rightarrow \mathbb{R}$, and a weight $\gamma \in [0, 1)$ such that for each $E \in \Sigma$, the preference relation \succsim_E admits a SEU representation with (u, μ_E) , where

$$(7) \quad \mu_E \equiv \gamma \mu + (1 - \gamma) \arg \min_{\pi \in \Delta(E)} d_\mu(\pi).$$

We first demonstrate that **IU** representations can be generated from **wIU** representations by imposing **Consequentialism**.

Proposition 9. *If a family of preference relations $\{\succsim_E\}_{E \in \Sigma}$ admits a **wIU** representation and satisfies **Consequentialism**, then it also admits an **IU** representation.*

To characterize **wIU**, we need to weaken **Dynamic Coherence** and **Consequentialism** to accommodate the DM's partial reaction to information. While our DM does not fully incorporate the informational content of the event A , her belief in A increases and, consequently, she necessarily gives lower credence to $S \setminus A$ and any $E \subseteq S \setminus A$.

We introduce the following definition to capture the DM's subjective perception of events that become relatively less likely after A .

Definition 11 (Unfavored Event). We say E is \succsim_A -unfavored if for any $E' \subseteq S$ and $p, q \in \Delta(X)$,

$$p E' w \sim q E w \text{ implies } p E' w \succsim_A q E w,$$

with at least one strict inequality for some E' . We then say E is a \succsim_A -favored event if E^c is \succsim_A -unfavored.

Similar to how **Dynamic Coherence** ensures a consistent reaction to null events, **Partial Dynamic Coherence** ensures a consistent reaction to favored events.

AXIOM 9 (Partial Dynamic Coherence). For any $A_1, \dots, A_n \subseteq S$, if A_i is $\succsim_{A_{i+1}}$ -favored for each $i \leq n - 1$ and A_n is \succsim_{A_1} -favored, then $\succsim_{A_1} = \succsim_{A_n}$.

Next, we require that her subjective belief in E weakly increases after she is told that E has occurred. While **Consequentialism** demands that the DM is convinced of E , our novel axiom **Partial Consequentialism** only demands that she puts more stock in E .

AXIOM 10 (Partial Consequentialism). For any $E \subseteq S$, E is \succsim_E -favored.

Finally, we require a condition to ensure a consistent reaction to all events. That is, the following condition guarantees that γ is event independent.

AXIOM 11 (**Relative Tradeoff Consistency**). For any $A, B \in \Sigma$, $p, q, r \in \Delta$, and $\alpha \in (0, 1)$,

If $w A q \sim p$ and $w A q \sim_A \alpha p + (1 - \alpha)w$, then

$w B q \sim r$ implies $w B q \sim_B \alpha r + (1 - \alpha)w$.

Theorem 4. *Suppose \succsim has a full-support. The following are equivalent.*

- (i) *A family of preference relations $\{\succsim_E\}_{E \in \Sigma}$ admits a **wIU** representation.*
- (ii) *It satisfies **SEU Postulates**, **Partial Dynamic Coherence**, **Partial Consequentialism**, and **Relative Tradeoff Consistency**.*
- (iii) *It admits a **wIU** representation with respect to a continuous, strictly convex distance function.*

8 Related Literature

A few papers have studied the idea of distance minimization and how it relates to belief updating. Perea (2009) axiomatized *imaging* rules, which are minimum distance rules utilizing Euclidean distance. Under imaging, for each $E \subseteq S$ a posterior π is selected that minimizes $d_\mu(\pi) = \|\phi(\mu) - \phi(\pi)\|$, where $\pi \in \Delta(E)$ and ϕ is an affine function. This is a special case of the **IU**. More recently, Basu (2019) studies AGM (Alchourrón et al., 1985) belief revision. Within this setting, he establishes an equivalence between updating rules that are AGM-consistent, Bayesian, and weak path independent and lexicographic updating rules. He then turns to minimum distance updating rules and shows that every support-dependent lexicographic updating rule admits a minimum distance representation. In contrast, we allow for non-Bayesian updating. Zhao (2022) and Dominiak et al. (2022) both study distance minimization with a more general notion of information; information is a subset I of $\Delta(S)$ rather than an event. This more general notion of information requires significantly different axioms. Moreover, Zhao (2022) focuses on Bayes' rule.

There is a large literature in experimental economics and psychology documenting various belief biases, and excellent surveys can be found in Camerer (1995) and Benjamin (2019). There is also growing number of papers taking axiomatic approaches to studying forms of non-Bayesian updating.¹¹ Of course, Ortoleva (2012) is the most closely related among these,

¹¹For behavioral models of non-Bayesian updating, see, for example Barberis et al. (1998); Rabin and Schrag (1999); Mullainathan (2002); Rabin (2002); Mullainathan et al. (2008); Gennaioli and Shleifer (2010); and Bordalo et al. (2016).

but it has already been discussed in detail. Other papers include Suleymanov (2021), which studies deviations from Bayesian updating caused by ambiguity; Epstein (2006) and Kovach (2020a), both of which study a prior-biased updating rule in which posterior beliefs are a convex combination of the prior and the Bayesian posterior; and Epstein et al. (2008), which extends Epstein (2006) to an infinite horizon setting. The updating rule in Epstein (2006), Kovach (2020a), and Epstein et al. (2008) is a special case of Weighted **IU** characterized in [section 7](#).

Our paper also contributes to a growing literature applying models of non-standard belief updating rules to games of strategic information transmission. Recent contributions in this area include Galperti (2019), de Clippel and Zhang (2022), and Lee et al. (2023).

As we carefully discussed in [section 4](#), updating under zero-probability events is studied in Myerson (1986a,b) and Ortoleva (2012). Another approach to dealing with null events is the Lexicographic Probability System (LPS) of Blume et al. (1991). While LPS also involves a collection of probability distributions, LPS utilizes the entire collection of distributions in the evaluation process via a lexicographic ordering. Consequently, a DM described by LPS will violate Archimedean Continuity, (see [AXIOM 1\(ii\)](#)) of the initial preference. Further, LPS replaces (Savage) null-events with “infinitely more likely than,” so that null-events are effectively precluded. While LPS necessarily deviates from SEU, there is a mathematical equivalence between conditional probabilities generated by LPS and CPS (e.g., see Brandenburger et al. (2006)). Hence, our results further clarify the connections between HT, CPS, and LPS.

A Proofs

A.1 Proof of [Proposition 1](#)

Take any non-null $E \in \Sigma$. Let $\text{sp}(\mu) = A \cup C$ and $E = B \cup C$ where $\text{sp}(\mu) \cap E = C$. We then solve the following optimization problem:

$$\max_{\pi \in \Delta(E)} \sum_{i=1}^n \mu_i \sigma \left(\frac{\pi_i}{\mu_i} \right) = \sum_{i \in A \cup C} \mu_i \sigma \left(\frac{\pi_i}{\mu_i} \right) = \sum_{i \in C} \mu_i \sigma \left(\frac{\pi_i}{\mu_i} \right) + \mu(A) \sigma(0).$$

Hence we want to maximize $f((\pi_i)_{i \in C}) = \sum_{i \in C} \mu_i \sigma \left(\frac{\pi_i}{\mu_i} \right)$ subject to the constraint $\sum_{i \in C} \pi_i = 1 - \pi(A)$. Let us first fix $\alpha = 1 - \pi(A)$ and $C' = \{i \in C | \pi_i > 0\}$. Then we need to maximize

$$\sum_{i \in C'} \mu_i \sigma \left(\frac{\pi_i}{\mu_i} \right) - \lambda \left(\sum_{i \in C'} \pi_i - \alpha \right).$$

The first order condition gives $\sigma' \left(\frac{\pi_i}{\mu_i} \right) = \lambda$ for each $i \in C'$ (Since σ is strictly concave, the FOC is sufficient). Hence, $\pi_i = \mu_i c'^{-1}(\lambda)$. After finding λ from the constraint $\sum_{i \in C'} \pi_i = \alpha$, we

have

$$\pi_i = \alpha \frac{\mu_i}{\mu(C')}.$$

If we calculate the objective function at the above values:

$$f((\pi_i)_{i \in C}) = \mu(C') \sigma\left(\frac{\alpha}{\mu(C')}\right) + \mu(C \setminus C') \sigma(0)$$

We need to find the optimal α and C' . Note that, $\mu(C) > \mu(C')$,

$$\mu(C) \sigma\left(\frac{\alpha}{\mu(C)}\right) > \mu(C') \sigma\left(\frac{\alpha}{\mu(C')}\right) + \mu(C \setminus C') \sigma(0);$$

equivalently,

$$\mu(C) \left(\sigma\left(\frac{\alpha}{\mu(C)}\right) - \sigma(0)\right) > \mu(C') \left(\sigma\left(\frac{\alpha}{\mu(C')}\right) - \sigma(0)\right).$$

To obtain the above inequality, it is sufficient to show that $x(\sigma(\frac{\alpha}{x}) - \sigma(0))$ is strictly increasing; i.e., $(x(\sigma(\frac{\alpha}{x}) - \sigma(0)))' = \sigma(\frac{\alpha}{x}) - \sigma(0) - \frac{\alpha}{x} \sigma'(\frac{\alpha}{x}) > 0$. This inequality holds since σ is strictly concave.

Since σ is strictly increasing, we also have $\mu(C) \sigma\left(\frac{1}{\mu(C)}\right) > \mu(C) \sigma\left(\frac{\alpha}{\mu(C)}\right)$ when $1 > \alpha$. Hence, f is maximized when $\alpha = 1$ and $C' = C$. In other words, $\pi_i = \frac{\mu_i}{\mu(C)}$; i.e., $\mu_E = \text{BU}(\mu, E)$.

A.2 Lemma 1

The following result will be useful.

Lemma 1. *For any $\mu, \pi \in \Delta(S)$,*

$$-\sigma(0) \geq \beta^\sigma(\mu, \pi) \geq -\sigma(1).$$

Proof of Lemma 1. Since σ is strictly increasing, it is immediate that $\beta^\sigma(\mu, \pi) \leq -\sigma(0)$. For any $C \in \Sigma$, let

$$f(C) = \mu(C) \sigma\left(\frac{1}{\mu(C)}\right) + (1 - \mu(C)) \sigma(0).$$

As we showed in the proof of Proposition 1, $x \sigma(\frac{1}{x}) + (1 - x) \sigma(0)$ is strictly increasing when $x \in (0, 1)$. Hence we have, $\sigma(1) \geq f(C)$. Let $A = \text{sp}(\mu) \cap \text{sp}(\pi)$. By Proposition 1, $f(A) \geq \beta^\sigma(\mu, \pi)$. Hence, $\beta^\sigma(\mu, \pi) \geq -\sigma(1)$. \square

A.3 Proof of Proposition 2

We first consider the scenario where E is a null-event. Then for any $\pi \in \Delta(E)$, we have

$$d_\mu(\pi) = \beta^\sigma(\mu^*, \pi).$$

Then by Proposition 1, we have $\mu_E = \text{BU}(\mu^*, E)$. Suppose now E is non-null. Let $\text{sp}(\mu) = A \cup C$ and $E = B \cup C$ where $\text{sp}(\mu) \cap E = C$.

$$d_\mu(\pi) = \begin{cases} \beta^\sigma(\mu, \pi) & \text{if } \pi \in \Delta(E) \setminus \Delta(B), \\ \beta^\sigma(\mu^*, \pi) + \sigma(1) + |\sigma(0)| & \text{if } \pi \in \Delta(B). \end{cases}$$

Let $\mu^1 = \text{BU}(\mu, E)$ and $\mu^2 = \text{BU}(\mu, B)$. By Proposition 1, μ^1 maximizes $\beta^\sigma(\mu, \pi)$ subject to the constraint $\pi \in \Delta(E)$. Again, by Proposition 1, μ^2 maximizes $\beta^\sigma(\mu^*, \pi)$ subject to the constraint $\pi \in \Delta(B)$. Hence, to show that $\mu_E = \mu^1$, it is sufficient to prove that $d_\mu(\mu^1) < d_\mu(\mu^2)$; equivalently,

$$d_\mu(\mu^1) = \beta^\sigma(\mu, \mu^1) < d_\mu(\mu^2) = \beta^\sigma(\mu^*, \mu^2) + \sigma(1) + |\sigma(0)|.$$

The above inequality is implied by Lemma 1.

A.4 Proof of Theorem 1

Note that (iii) trivially implies (i). Let us first show that (i) implies (ii). Suppose $\{\succsim_E\}$ admits an **IU** representation with respect to (μ, u, d_μ) . The **IU** representation indeed satisfies SEU postulates. We now prove the necessity of Consequentialism and Dynamic Coherence.

Consequentialism. Take any $E \in \Sigma$ and $f, g \in F$ such that $f(s) = g(s)$ for all $s \in E$. Since $\mu_E(E) = 1$ and $f(s) = g(s)$ for all $s \in E$, we have

$$\sum_{s \in S} \mu_E(s) f(s) = \sum_{s \in E} \mu_E(s) f(s) = \sum_{s \in S} \mu_E(s) g(s) = \sum_{s \in E} \mu_E(s) g(s);$$

i.e., $f \sim_E g$.

Dynamic Coherence. Take any $A_1, \dots, A_n \subseteq S$ such that $S \setminus A_i$ is $\succsim_{A_{i+1}}$ -null for each $i \leq n-1$ and $S \setminus A_n$ is \succsim_{A_1} -null. Equivalently, $\mu_{A_{i+1}}(A_i) = 1$ for each $i \leq n-1$ and $\mu_{A_1}(A_n) = 1$. Since $\mu_{A_{i+1}} \in \Delta(A_i)$ and $\mu_{A_i} = \arg \min_{\pi \in \Delta(A_i)} d_\mu(\pi)$, $d_\mu(\mu_{A_i}) \leq d_\mu(\mu_{A_{i+1}})$. Similarly, we have $d_\mu(\mu_{A_n}) \leq d_\mu(\mu_{A_1})$. Therefore, we have

$$d_\mu(\mu_{A_1}) \leq d_\mu(\mu_{A_2}) \leq \dots \leq d_\mu(\mu_{A_n}) \leq d_\mu(\mu_{A_1});$$

i.e., $d_\mu(\mu_{A_1}) = d_\mu(\mu_{A_n})$. Since μ_{A_n} is the unique minimizer of d_μ in $\Delta(A_n)$ and $\mu_{A_1} \in \Delta(A_n)$, $d_\mu(\mu_{A_1}) = d_\mu(\mu_{A_n})$ implies that $\mu_{A_1} = \mu_{A_n}$; i.e., $\succsim_{A_1} = \succsim_{A_n}$.

Let us now show that (ii) implies (iii). Suppose $\{\succsim_E\}_{E \in \Sigma}$ satisfies **SEU Postulates**, **Consequentialism**, and **Dynamic Coherence**. Since \succsim satisfies SEU postulates, there is (μ, u) such that \succsim has a SEU representation with (μ, u) . Since \succsim_E satisfies SEU postulates, there is (μ_E, u_E) such that \succsim_E has a SEU representation with (μ_E, u_E) . By Invariant Risk

Preference, $u_E(p) \geq u_E(q)$ and $u(p) \geq u(q)$ for any $p, q \in \Delta(X)$. Without loss of generality, let us assume that $u_E = u$. Hence, \succsim_E has a SEU representation with (μ_E, u) .

Let us now discuss implications of **Consequentialism**. Take any $E \in \Sigma$ and any $f, g \in F$ and $p, q \in \Delta(X)$ such that $p \succ q$ and $f(s) = g(s) = p$ for all $s \in E$ and $f(s) = p$ and $g(s) = q$ for any $s \in E^c$. By **Consequentialism**, we have $f \sim_E g$; equivalently,

$$\sum_{s \in S} \mu_E(s) f(s) = u(p) = \sum_{s \in E} \mu_E(s) g(s) = \mu_E(E) u(p) + (1 - \mu_E(E)) u(q).$$

In other words, we have $\mu_E(E) = 1$; i.e., $\mu_E \in \Delta(E)$.

Afriat's theorem for general budget sets. To obtain the IU representation, we use an extension of Afriat's theorem (Afriat (1967)) for general budget sets due to Matzkin (1991). To state Afriat's theorem for general budget sets, some notations are necessary. Let Z be a convex, bounded subset of \mathbb{R}_+^n . Let $\mathcal{D} = (\mathbf{x}^t, B^t)_{t \in T}$ be a data set where $\mathbf{x}^t \in B^t$ is the observed consumption bundle that is chosen from the budget set $B^t \subset Z$ at observation $t \in T$. We say that (\mathbf{x}^t, B^t) is a co-convex subset of Z if the following three conditions hold: (i) $Z \setminus B^t$ is open and convex; (ii) for any $\mathbf{e} \geq 0$ and $\mathbf{x} \in Z \setminus B^t$, $\mathbf{x} + \mathbf{e} \in Z$ implies $\mathbf{x} + \mathbf{e} \in Z \setminus B^t$; and (iii) for any $\mathbf{e} > 0$, $\mathbf{x}^t + \mathbf{e} \in Z$ implies $\mathbf{x}^t + \mathbf{e} \in Z \setminus B^t$.

Let us now define the following revealed preference relation on $\{\mathbf{x}^t\}_{t \in T}$. We say \mathbf{x}^t is revealed preferred to \mathbf{x}^s , denoted by $\mathbf{x}^t \succsim_R \mathbf{x}^s$ if $\mathbf{x}^s \in B^t$. We say \mathbf{x}^t is strictly revealed preferred to \mathbf{x}^s , denoted by $\mathbf{x}^t \succ_R \mathbf{x}^s$ if $\mathbf{x}^s \in B^t$ and $\mathbf{x}^t \neq \mathbf{x}^s$. Finally, we say the data set $\mathcal{D} = (\mathbf{x}^t, B^t)_{t \in T}$ satisfies the Strong Axiom of Revealed Preferences (SARP) if \succsim_R is acyclic; i.e., there is no sequence $\mathbf{x}^{t_1}, \mathbf{x}^{t_2}, \dots, \mathbf{x}^{t_L}$ such that $\mathbf{x}^{t_l} \succsim_R \mathbf{x}^{t_{l+1}}$ for each $l \leq L - 1$ and $\mathbf{x}^{t_L} \succ_R \mathbf{x}^{t_1}$.

Theorem 1 of Matzkin (1991). Suppose for each $t \in T$, (\mathbf{x}^t, B^t) is a co-convex subset of Z . Then the data set $\mathcal{D} = (\mathbf{x}^t, B^t)_{t \in T}$ satisfies SARP if and only if there is a strictly increasing, continuous, strictly concave utility function $u : Z \rightarrow \mathbb{R}$ such that for any $t \in T$,

$$u(\mathbf{x}^t) > u(\mathbf{x}) \text{ for any } \mathbf{x} \in B^t \setminus \{\mathbf{x}^t\}.$$

To apply the above theorem, let us arbitrarily label the set of all events: $\Sigma = \{E_t\}_{t \in T}$. Then let $Z = \Delta(S)$ and $\mathbf{x}^t = \mu_{E_t}$ and $B^t = \Delta(E_t)$ for each $t \in T$. Let $\mathcal{D} = (\mathbf{x}^t, B^t)_{t \in T}$.

Note that Z is a convex, bounded subset of \mathbb{R}_+^n . Let us show that (\mathbf{x}^t, B^t) is a co-convex subset of Z . First, $Z \setminus B^t$ is open and convex in Z . Second, for any $\mathbf{x} \in Z$ and $\mathbf{e} \geq 0$, $\mathbf{x} + \mathbf{e} \in Z$ implies $\mathbf{e} = 0$. Hence, (ii) and (iii) of co-convexity are trivially satisfied.

Let us now show that **Dynamic Coherence** implies that $\mathcal{D} = (\mathbf{x}^t, B^t)_{t \in T}$ satisfies SARP. Take any sequence $\mathbf{x}^{t_1}, \mathbf{x}^{t_2}, \dots, \mathbf{x}^{t_L}$ such that $\mathbf{x}^{t_l} \succsim_R \mathbf{x}^{t_{l+1}}$ for each $l \leq L - 1$ and $\mathbf{x}^{t_L} \succ_R \mathbf{x}^{t_1}$. To prove SARP, we shall show that $\mathbf{x}^{t_L} = \mathbf{x}^{t_1}$. By definition of the revealed preference relation

$\succsim_R, \mathbf{x}^{t_l} \succsim_R \mathbf{x}^{t_{l+1}}$ is equivalent to $\mathbf{x}^{t_{l+1}} \in \Delta(E^{t_l})$. In other words, $\mu_{E^{t_{l+1}}} \in \Delta(E^{t_l})$ for each $l \leq L-1$. Similarly, $\mu_{E^{t_1}} \in \Delta(E^{t_L})$.

Note that $\mu_{E^{t_{l+1}}} \in \Delta(E^{t_l})$ implies $\mu_{E^{t_{l+1}}}(E^{t_l}) = 1$; equivalently, $\mu_{E^{t_{l+1}}}(S \setminus E^{t_l}) = 0$. In other words, $S \setminus E^{t_l}$ is $\succsim_{E^{t_{l+1}}}$ -null for each $l \leq L-1$. Similarly, $S \setminus E^{t_L}$ is $\succsim_{E^{t_1}}$ -null. By **Dynamic Coherence**, $\succsim_{E^{t_1}} = \succsim_{E^{t_L}}$; equivalently, $\mu_{E^{t_1}} = \mu_{E^{t_L}}$. In other words, $\mathbf{x}^{t_1} = \mathbf{x}^{t_L}$.

Since $\mathcal{D} = (\mathbf{x}^t, B^t)_{t \in T}$ satisfies SARP, by Theorem 1 of Matzkin (1991), there is a strictly increasing, continuous, strictly concave utility function $u : Z \rightarrow \mathbb{R}$ such that for any $t \in T$,

$$u(\mathbf{x}^t) > u(\mathbf{x}) \text{ for any } \mathbf{x} \in B^t \setminus \{\mathbf{x}^t\}.$$

Let $d_\mu = -u$. Then since $B^t = \Delta(E^t)$ and $\mathbf{x}^t = \mu_{E^t}$,

$$\mu_{E^t} = \arg \min_{\pi \in \Delta(E^t)} d_\mu(\pi).$$

Finally, note that d_μ is continuous and strictly convex.

A.5 Proof of Proposition 3

This follows directly from existing results on Dynamic Consistency. For example, see Ghirardato (2002).

A.6 Proof of Proposition 4

We start by constructing a distortion δ_E for an arbitrary event E . Without loss, suppose $|E| \geq 2$. Fix $s^* \in E$ and let $\delta_E(\mu(s^*)) = \mu_{E^*}(s^*)$. For all $s \in E$, let $\delta_E(\mu(s)) = \frac{\mu_E(s)}{\mu_E(s^*)}$. Consider any $s_1, s_2 \in E$ such that $\mu(s_1) = \mu(s_2)$. Then by **Consistency** it follows that $\mu_E(s_1) = \mu_E(s_2)$, and by construction of δ_E , it follows that $\delta_E(\mu(s_1)) = \frac{\mu_E(s_1)}{\mu_E(s^*)} = \frac{\mu_E(s_2)}{\mu_E(s^*)} = \delta_E(\mu(s_2))$, hence δ_E is well-defined. Finally, note that for any $s, s' \in E$, $\frac{\delta_E(\mu(s'))}{\delta_E(\mu(s))} = \frac{\mu_E(s')}{\mu_E(s)}$. Summing over $s' \in E$ and using $\sum_{s' \in E} \mu_E(s') = 1$ yields $\mu_E(s) = \frac{\delta_E(\mu(s))}{\sum_{s' \in E} \delta_E(\mu(s'))}$, hence $\mu_E = BU(\delta_E(\mu), E)$.

Next, we use **Independence of Irrelevant Information** to show that δ_E is in fact independent of E . Fix any E_1, E_2 with $s, s' \in E_1 \cap E_2$. Consider some p, q such that $p\{s\}r \sim_{E_1} q\{s'\}r$. It is without loss to suppose $u(r) = 0$, and hence $u(p)\mu_{E_1}(s) = u(q)\mu_{E_1}(s')$. By the previous result, it follows that

$$u(p) \frac{\delta_{E_1}(\mu(s))}{\sum_{\tilde{s} \in E_1} \delta_{E_1}(\mu(\tilde{s}))} = u(q) \frac{\delta_{E_1}(\mu(s'))}{\sum_{\tilde{s} \in E_1} \delta_{E_1}(\mu(\tilde{s}))},$$

and so

$$\frac{\delta_{E_1}(\mu(s))}{\delta_{E_1}(\mu(s'))} = \frac{u(q)}{u(p)}.$$

By applying **Independence of Irrelevant Information**, it follows that $p\{s\}r \sim_{E_2} q\{s'\}r$,

and hence

$$\frac{\delta_{E_1}(\mu(s))}{\delta_{E_1}(\mu(s'))} = \frac{u(q)}{u(p)} = \frac{\delta_{E_2}(\mu(s))}{\delta_{E_2}(\mu(s'))}.$$

Hence there exists a $\delta : [0, 1] \rightarrow \mathbb{R}$ such that for any E , $\mu_E = BU(\delta(\mu), E)$. Finally, since δ is clearly only unique up to a scalar, it is without loss to suppose that $\delta : [0, 1] \rightarrow [0, 1]$.

A.7 Proof of Proposition 5

It is clear that **Monotonicity** implies **Consistency**, and by the previous result we have some $\delta : [0, 1] \rightarrow [0, 1]$ such that $\mu_E = BU(\delta(\mu), E)$ for any non-null E . Consider any E and $s, s' \in E$ and suppose $\mu(s) > \mu(s')$. Then from **Monotonicity**, if $x \succ y$ it follows that $x\{s\}y \succ x\{s'\}y$ and thus $x\{s\}y \succ_E x\{s'\}y$, which implies $\mu_E(s) > \mu_E(s')$. From here it is immediate that $\delta(\mu(s)) > \delta(\mu(s'))$. Since δ is arbitrary outside of $\{\mu(s)\}_{s \in S}$, it can be extended to $[0, 1]$ so that δ is strictly increasing.

A.8 Proof of Theorem 2

Necessity of the axioms is trivial, so we only prove sufficiency. By **SEU Postulates**, there are U and $\{\mu_E\}_{E \in \Sigma}$ such that for any $E \in \Sigma$, \succsim_E admits a SEU representation with (μ_E, U) ; for all $f, g \in \mathcal{F}$:

$$f \succsim_E g \quad \text{if and only if} \quad \sum_{s \in E} U(f(s))\mu_E(s) \geq \sum_{\omega \in S_E} U(g(s))\mu_E(s).$$

By **Consequentialism**, $\mu_E(E) = 1$. We now shall show Equation 6. Take any $E \in \Sigma$ and $A \subseteq E$. Let $A \in \Sigma$ be a \succsim_E -non-null event; i.e., $\mu_E(A) > 0$. Consider acts $f_A h$ and $g_A h$ such that $f_A h \succsim_E g_A h$; i.e.,

$$(8) \quad \sum_{s \in S} U(f_A h(s))\mu_E(s) \geq \sum_{s \in S} U(g_A h(s))\mu_E(s).$$

By **Consequentialism**,

$$(9) \quad \sum_{s \in E} U(f_A h(s))\mu_E(s) \geq \sum_{s \in E} U(g_A h(s))\mu_E(s),$$

or equivalently,

$$(10) \quad \sum_{s \in A} U(f(s))\mu_E(s) \geq \sum_{s \in A} U(g(s))\mu_E(s).$$

By **Conditional Consistency**, $f \succsim_A g$; i.e.,

$$(11) \quad \sum_{s \in A} U(f(s))\mu_A(s) \geq \sum_{s \in A} U(g(s))\mu_A(s).$$

Since Equations (10) hold for any f, g , we have Bayesian updating

$$(12) \quad \mu_A(s) = \frac{\mu_E(s)}{\mu_E(A)} \text{ for each } s \in A.$$

Finally, if $\mu_E(A) = 0$, then we have $\mu_E(s) = 0$ for any $s \in A$. Hence, Equation 6 holds.

A.9 Proof of Proposition 6

Suppose $\{\succsim_E\}_{E \in \Sigma}$ admits a CPS representation. To prove the first part of this proposition, we construct $\mu^0, \dots, \mu^K \in \Delta(S)$ inductively. Let $S_0 = S$, and consider \succsim_{S_0} . We let $\mu^0 = \mu_{S_0}$, and if μ^0 has full support, stop. Otherwise, let S_1 denote the set of all \succsim_{S_0} -null states, and let $\mu^1 = \mu_{S_1}$. By **Consequentialism**, $\mu^1(S_0 \setminus S_1) = 0$. If $\text{sp}(\mu^1) = S_1$, stop. Otherwise, let S_2 denote the set of all \succsim_{S_1} -null states, and $\mu^2 = \mu_{S_2}$. We proceed in this fashion until we reach a K such that $\text{sp}(\mu^K) = S_K$. Since S is finite, we must eventually stop. Note that we have constructed $\mu^0, \dots, \mu^K \in \Delta(S)$ such that $\text{sp}(\mu^0), \dots, \text{sp}(\mu^K)$ is a partition of S . We now shall prove that for any $E \in \Sigma$, $\mu_E = \text{BU}(\mu^{k^*}, E)$ where $k^* = \min\{k \mid \mu^k(E) > 0\}$.

Since $k^* = \min\{k \mid \mu^k(E) > 0\}$, $E \subseteq S_{k^*} = \bigcup_{k \geq k^*} \text{sp}(\mu^k)$. By the construction, $\mu^{k^*} = \mu_{S_{k^*}}$. Hence, $\mu_{S_{k^*}}(E) > 0$. Then by Equation 6, for any $s \in E$, $\mu_E(s) = \frac{\mu^{k^*}(s)}{\mu^{k^*}(E)}$; equivalently, $\mu_E = \text{BU}(\mu^{k^*}, E)$.

We now shall show that $\{\succsim_E\}_{E \in \Sigma}$ has an **IU** representation with respect to the following distance function:

$$d_\mu(\pi) = \beta^\sigma(\mu^{k^*}, \pi) + k^* (\sigma(1) + |\sigma(0)|),$$

where $k^* = \min\{k \mid \mu^k(\text{sp}(\pi)) > 0\}$. It is enough to show that for any $E \in \Sigma$,

$$\mu_E = \arg \min_{\pi \in \Delta(E)} d_\mu(\pi).$$

Take any E and let $k^* = \min\{k \mid \mu^k(E) > 0\}$. Note that for any $\pi \in \Delta(E)$, $\min\{k \mid \mu^k(\text{sp}(\pi)) > 0\} \geq k^*$. Hence, $\Delta_{k^*}, \dots, \Delta_K$ be the partition of $\Delta(E)$ such that for any $\pi \in \Delta(E)$, $\pi \in \Delta_l$ if and only if $l = \min\{k \mid \mu^k(\text{sp}(\pi)) > 0\}$. Let $\rho^l = \arg \min_{\pi \in \Delta_l} d_\mu(\pi)$. By Proposition 1, $\rho^{k^*} = \text{BU}(\mu^{k^*}, E)$. Take any $l > k^*$. We shall show $d_\mu(\rho^{k^*}) < d_\mu(\rho^l)$; equivalently,

$$\beta^\sigma(\mu^{k^*}, \rho^{k^*}) - \beta^\sigma(\mu^l, \rho^l) < (l - k^*)(\sigma(1) + |\sigma(0)|).$$

The above equality is implied by Lemma 1.

A.10 Proof of Corollary 2

See the proof of [Theorem 3](#) as this corollary is a special case of [Theorem 3](#) when $\epsilon = 0$.

A.11 Proof of Theorem 3

Let $\{\succsim_E\}$ be a family of preference relations with an ϵ -OS representation for some $\epsilon \in [0, 1)$. Then, there are probability distributions μ_0, \dots, μ_K such that

$$\mu_E = \text{BU}(\mu_{k^*}, E) \text{ where } k^* = \min\{k \leq K \mid \mu_k(E) > \epsilon\}$$

for every $E \in \Sigma$. Let $\Sigma_0, \dots, \Sigma_K$ be a partition of Σ such that for each k , Σ_k is the collection of events for which the prior μ_k is used for updating:

$$\Sigma_k = \{E \in \Sigma \mid k = \min\{\tilde{k} \leq K \mid \mu_{\tilde{k}}(E) > \epsilon\}\}.$$

Throughout this proof, we assume that for any $k \leq K$, E_k is an element of Σ_k . Take $\bar{\rho}_0, \underline{\rho}_0, \dots, \bar{\rho}_K, \underline{\rho}_K$ with

$$\bar{\rho}_0 > \underline{\rho}_0 > \bar{\rho}_1 > \underline{\rho}_1 > \dots > \bar{\rho}_K > \underline{\rho}_K > \delta \bar{\rho}_0 > 0$$

and $\underline{\rho}_k > \bar{\rho}_k \mu^{E'_k}(E_k)$ for any E_k, E'_k with $\mu^{E'_k}(E_k) < 1$.

Let $\mu_k^E = \text{BU}(\mu_k, E)$ for any $E \in \Sigma$. Let ρ be an element of $\Delta(\{\mu_k^{E_k}\}_{k \leq K, E_k \in \Sigma_k})$ such that (i) $\rho(\mu_k^{E_k}) \in (\underline{\rho}_k, \bar{\rho}_k)$ for any $k \leq K$ and (ii) $\rho(\mu_k^{E_k}) > \rho(\mu_k^{E'_k})$ if $\mu_k^{E_k} \neq \mu_k^{E'_k}$ and $\mu_k^{E'_k}(E_k) = 1$.

Let us first show that there is ρ that satisfies (ii). Let $\mu_k^{E_k} \succ^* \mu_k^{E'_k}$ if $\mu_k^{E_k} \neq \mu_k^{E'_k}$ and $\mu_k^{E'_k}(E_k) = 1$. It is enough to show that \succ^* is acyclic. To show acyclicity, suppose that there are E_k^1, \dots, E_k^T such that $\mu_k^{E_k^t}(E_k^{t+1}) = 1$ for each $t \leq T-1$ and $\mu_k^{E_k^T}(E_k^1) = 1$. Note that $\mu_k^{E'_k}(E_k) = 1$ is equivalent to $\text{supp}(\mu_k) \cap E'_k \subseteq E_k$. Hence, $\mu_k^{E'_k}(E_k) = 1$ implies $\text{supp}(\mu_k) \cap E'_k \subseteq \text{supp}(\mu_k) \cap E_k$. Then, $\mu_k^{E_k^t}(E_k^{t+1}) = 1$ implies $\text{supp}(\mu_k) \cap E_k^t \subseteq \text{supp}(\mu_k) \cap E_k^{t+1}$ and $\mu_k^{E_k^T}(E_k^1) = 1$ implies $\text{supp}(\mu_k) \cap E_k^T \subseteq \text{supp}(\mu_k) \cap E_k^1$. Hence, $\text{supp}(\mu_k) \cap E_k^t = \text{supp}(\mu_k) \cap E_k^{t'}$ for any t, t' ; i.e., $\mu_k^{E_k^t} = \mu_k^{E_k^{t'}}$.

We now show that $\{\succsim_E\}$ has a HT representation with (ρ, δ) when δ is large enough. Hence, we shall show that for any E_k ,

$$\rho(\mu_k^{E_k}) \mu_k^{E_k}(E_k) = \rho(\mu_k^{E_k}) > \rho(\mu_j^{E_j}) \mu_j^{E_j}(E_k) \text{ for any } j \neq k.$$

For any $j > k$, the above holds since $\rho(\mu_k^{E_k}) > \rho(\mu_j^{E_j})$. Suppose now $j < k$. In this case, $\mu_j(E_k) \leq \epsilon$ since k is the lowest index such that $\mu_k(E_k) > \epsilon$. Then, $\mu_j^{E_j}(E_k) = \text{BU}(\mu_j, E_j)(E_k) = \frac{\mu_j(E_k \cap E_j)}{\mu_j(E_j)}$. Since $\mu_j(E_k) \leq \epsilon$ and $\mu_j(E_j) > \epsilon$, there is a large enough

$\delta \in [0, 1)$ such that $\mu_j^{E_j}(E_k) \leq \delta$. Hence, by the construction of ρ ,

$$\rho(\mu_k^{E_k}) > \delta \rho(\mu_j^{E_j}) \geq \rho(\mu_j^{E_j}) \mu_j^{E_j}(E_k).$$

We finally show that the HT representation correctly chooses $\mu_k^{E_k}$ among $\{\mu_k^{E'_k}\}_{E'_k}$ for each E_k .

When $\mu_k^{E'_k}(E_k) < 1$, we have

$$\rho(\mu_k^{E_k}) \mu_k^{E_k}(E_k) = \rho(\mu_k^{E_k}) > \underline{\rho}_k > \bar{\rho}_k \mu_j^{E'_k}(E_k) > \rho(\mu_k^{E'_k}) \mu_j^{E'_k}(E_k).$$

When $\mu_k^{E'_k}(E_k) = 1$ and $\mu_k^{E_k} \neq \mu_k^{E'_k}$,

$$\rho(\mu_k^{E_k}) \mu_k^{E_k}(E_k) = \rho(\mu_k^{E_k}) > \rho(\mu_k^{E'_k}) = \rho(\mu_k^{E'_k}) \mu_j^{E'_k}(E_k).$$

It is immediate from the above construction of δ , $\delta = 0$ whenever $\epsilon = 0$.

A.12 Proof of Proposition 7

Let $A = \{i \leq n : g(\rho_i) u_i \geq 0\}$ and $B = \Omega \setminus A$. Let $x_i = \pi_{m_1}(\omega_i)$ and $\delta_i = |g(\rho_i) u_i|$. The sender's problem reduces to

$$\max_{\mathbf{x} \in [0,1]^n} \sum_{i=1}^n \rho_i x_i \text{ subject to } \sum_{i \in A} \delta_i f(x_i) \geq \sum_{i \in B} \delta_i f(x_i).$$

It is immediate that $x_i^* = 1$ whenever $i \in A$. Then

$$\max_{x_i \in [0,1]} \sum_{i \in B} \rho_i x_i \text{ subject to } M \geq \sum_{i \in B} \delta_i f(x_i),$$

where $M = \sum_{i \in A} \delta_i f(1)$.

Case 1. $f(x) = x$ and $\frac{g(\rho_i) u_i}{\rho_i} \neq \frac{g(\rho_j) u_j}{\rho_j}$ for any i, j with $u_i, u_j < 0$

Note that when $\frac{\rho_i}{\delta_i} > \frac{\rho_j}{\delta_j}$, we cannot have $1 > x_i^*$ and $x_j^* > 0$. The optimal signal structure takes a form

$$x_{i_1}^* = \dots = x_{i_k}^* = 1 > x_{i_{k+1}}^* = \frac{M - \sum_{s=1}^k \delta_{i_s}}{\delta_{i_{k+1}}} \geq x_{i_{k+2}}^* = x_{i_{|B|}}^* = 0,$$

where $\frac{\rho_{i_1}}{\delta_{i_1}} > \dots > \frac{\rho_{i_k}}{\delta_{i_k}} > \dots > \frac{\rho_{i_{|B|}}}{\delta_{i_{|B|}}}$.

Case 2. f is strictly concave.

Let us show that for any i, j , we cannot have $x_i^*, x_j^* \in (0, 1)$. Take any i, j and let $\delta_i f(x_i^*) +$

$\delta_j f(x_j^*) = m$. Then x_i^*, x_j^* must be the solution to the following maximization problem

$$\max_{x_i, x_j \in [0,1]} \rho_i x_i + \rho_j x_j \text{ subject to } \delta_i f(x_i) + \delta_j f(x_j) = m.$$

From the constraint, we have $x_j = f^{-1}\left(\frac{m - \delta_i f(x_i)}{\delta_j}\right)$. Hence, the above maximization problem reduces to

$$\max_{x_i \in [a_1, a_2]} \rho_i x_i + \rho_j f^{-1}\left(\frac{m - \delta_i f(x_i)}{\delta_j}\right),$$

where $a_1 = \max\{0, f^{-1}\left(\frac{m - \delta_j f(1)}{\delta_i}\right)\}$ and $a_2 = \min\{1, f^{-1}\left(\frac{m - \delta_j f(0)}{\delta_i}\right)\}$. The objective function is strictly convex since f is strictly concave and f is increasing. Hence, either $x_i^* = a_1$ or $x_i^* = a_2$. Note that $x_i^* = a_1$ means that either $x_i^* = 0$ or $x_j^* = 1$ and $x_i^* = a_2$ means that either $x_i^* = 1$ or $x_j^* = 0$.

Hence, the optimal signal structure takes a form

$$x_{i_1}^* = \dots = x_{i_k}^* = 1 > x_{i_{k+1}}^* = \frac{M - \sum_{s=1}^k \delta_{i_s} f(1) - \sum_{s=k+2}^{|B|} \delta_{i_s} f(0)}{\delta_{i_{k+1}}} \geq x_{i_{k+2}}^* = x_{i_{|B|}}^* = 0,$$

where $\{i_1, \dots, i_{|B|}\}$ is a permutation of B and k be the largest integer such that $M \geq k$.

A.13 Proof of Proposition 8

Similar to the argument in the proof of Proposition 7, we need to solve

$$\max_{x_i \in [0,1]} \sum_{i \in B} \rho_i x_i \text{ subject to } M \geq \sum_{i \in B} \delta_i f(x_i),$$

where $M = \sum_{i \in A} \delta_i f(1)$. As long as $M > 0$, there exists $x_j^* > 0$. Take any $i \neq j$. Let us show that $x_i^* = 0$. Let $\delta_i f(x_i^*) + \delta_j f(x_j^*) = m$. Then x_i^*, x_j^* must be the solution to the following maximization problem

$$\max_{x_i, x_j \in [0,1]} \rho_i x_i + \rho_j x_j \text{ subject to } \delta_i f(x_i) + \delta_j f(x_j) = m.$$

From the constraint, we have $x_j = f^{-1}\left(\frac{m - \delta_i f(x_i)}{\delta_j}\right)$. Hence, the above maximization problem reduces to

$$\max_{x_i \in [a_1, a_2]} \rho_i x_i + \rho_j f^{-1}\left(\frac{m - \delta_i f(x_i)}{\delta_j}\right),$$

where $a_1 = \max\{0, f^{-1}\left(\frac{m - \delta_j f(1)}{\delta_i}\right)\}$ and $a_2 = \min\{1, f^{-1}\left(\frac{m - \delta_j f(0)}{\delta_i}\right)\}$. Since $f'(0) = 0$, $(\rho_i x_i + \rho_j f^{-1}\left(\frac{m - \delta_i f(x_i)}{\delta_j}\right))'_{x_i} |_{x_i=0} = \rho_i > 0$. Hence $x_i^* = 0$ cannot be optimal solution. Hence, $x_i^* > 0$.

A.14 Proof of Theorem 4

Proof. (ii) \Rightarrow (iii). Take any $A \subset S$. Since \succsim and \succsim_A have SEU representations w.r.s.t (u, μ) and (u, μ_A) , E is \succsim_A -unfavored if for any $E' \subseteq S$ and $p, q \in \Delta(X)$, for any $u(p)\mu(E') = u(q)\mu(E)$ implies $u(p)\mu_A(E') \geq u(q)\mu_A(E)$, with at least one strict inequality for some E' . Note that when $\mu = \mu_A$, there is no \succsim_A -unfavored event since $u(p)\mu(E') = u(q)\mu(E)$ implies $u(p)\mu_A(E') = u(q)\mu_A(E)$ for every E' and p, q . However, by Partial Consequentialism, A^c is \succsim -unfavored. Hence, $\mu \neq \mu_A$.

If E is \succsim_A -unfavored, then $\delta(A) = \frac{\mu_A(E)}{\mu(E)}$ where $\delta(A) = \min_{E'} \frac{\mu_A(E')}{\mu(E')}$. Since $\mu \neq \mu_A$, $\delta(E) < 1$. Therefore,

$$E \text{ is } \succsim_A \text{-unfavored iff } \frac{\mu_A(E)}{\mu(E)} = \delta(A).$$

Consider the vector $\mu_A^* = \frac{\mu_A - \delta(A)\mu}{1 - \delta(A)}$. For each $s \in S$, since $\frac{\mu_A(s)}{\mu(s)} \geq \delta(A)$, $\mu_A^*(s) = \frac{\mu_A(s) - \delta(A)\mu(s)}{1 - \delta(A)} \geq 0$. Moreover, $\sum_{s \in S} \mu_A^*(s) = \sum_{s \in S} \frac{\mu_A(s) - \delta(A)\mu(s)}{1 - \delta(A)} = 1$. Hence, $\mu_A^* \in \Delta(S)$ and

$$\mu_A = \delta(A)\mu + (1 - \delta(A))\mu_A^*.$$

Note that E is \succsim_A -unfavored iff $\mu_A(E) = \delta(A)\mu(E)$ iff $\mu_A^*(E) = 0$. Then by Partial Consequentialism, A^c is \succsim_A -unfavored iff $\mu_A^*(A^c) = 0$. Hence, $\mu_A^* \in \Delta(A)$. We now shall show that there is a function d that $\mu_A^* = \arg \min_{\pi \in \Delta(A)} d_\mu(\pi)$.

We now essentially repeat the part of Theorem 1 for the data set $\mathcal{D}^* = \{(\mu_A^*, \Delta(A))\}_{A \in \Sigma}$ where $\mu_S^* = \mu$. To apply the aforementioned generalization of Afriat's theorem for general budget sets, we first define the following revealed preference relation. We say that μ_A^* is strictly revealed preferred to μ_B^* , denoted by $\mu_A^* R^* \mu_B^*$, if $\mu_B^* \in \Delta(A)$ and $\mu_A^* \neq \mu_B^*$. First, note that $\mu_S^* R^* \mu_A^*$ for any $A \in \Sigma \setminus \{S\}$. Second, $\neg \mu_A^* R^* \mu_S^*$ since $\mu \notin \Delta(A)$. Third, for any $A, B \in \Sigma \setminus \{S\}$, $\mu_A^* R^* \mu_B^*$ implies that A is \succsim_B -favored. Hence, Partial Dynamic Coherence is equivalent to the acyclicity of R^* .

By the arguments provided in the proof of Theorem 1, $(\mu_A^*, \Delta(A))$ is co-convex. Since \mathcal{D}^* satisfies GARP, by Theorem 1 of Matzkin (1991), there is a strictly increasing, continuous, strictly concave utility function $u : \Delta(S) \rightarrow \mathbb{R}$ such that for any $A \in \Sigma$,

$$\mu_A^* = \arg \max_{\pi \in \Delta(A)} u(\pi).$$

Let $d_\mu = -u$ and note that μ is the global minimizer of d_μ by the previous equation. Moreover,

$$\mu_A^* = \arg \min_{\pi \in \Delta(A)} d_\mu(\pi).$$

Finally, note that d_μ is continuous and strictly convex. To sum up, we have

$$\mu_A = \delta(A) \mu + (1 - \delta(A)) \arg \min_{\pi \in \Delta(A)} d_\mu(\pi)$$

for any $A \subseteq S$. We now shall show that $\delta(A) = \delta(B)$.

Take any $A, B \in \Sigma \setminus S$. There are p, q, r such that $\mu(A^c)u(q) = u(p)$ and $\mu(B^c)u(q) = u(r)$; equivalently, $w A q \sim p$ and $w B q \sim_B r$. Since $\mu_A(A^c) = \delta(A) \mu(A^c)$, we have $\delta(A) \mu(A^c)u(q) = \mu_A(A^c)u(q) = \delta(A) u(p)$; equivalently, $w A q \sim_A \delta(A) p + (1 - \delta(A))w$. By **Relative Tradeoff Consistency**, we have $w B q \sim \delta(A) r + (1 - \delta(A))w$; equivalently, $\mu_B(B^c)u(q) = \delta(B) \mu(B^c)u(q) = \delta(A)u(r) = \delta(A) \mu(B^c)u(q)$. Hence, $\delta(A) = \delta(B) = \delta$. Finally, we set $\delta(S) = \delta$ and obtain a Weighted **IU** representation.

(i) \Rightarrow (ii). SEU postulates are trivially satisfied. Since μ has full-support, $\mu_A \neq \mu$ for any $A \subset S$. We now shall prove the necessity of the other three axioms. By the argument above, E is \succsim_A -unfavored iff $\mu_A^*(E) = 0$ where $\mu_A^* = \arg \min_{\pi \in \Delta(A)} d_\mu(\pi)$. Equivalently, E is \succsim_A -favored iff $\mu_A^*(E) = 1$.

Partial Consequentialism is satisfied because A is \succsim_A -favored; i.e., $\mu_A^* \in \Delta(A)$.

To prove Partial Dynamic Coherence, take any $A_1, \dots, A_n \subseteq S$ such that A_i is $\succsim_{A_{i+1}}$ -favored for each $i \leq n - 1$ and A_n is \succsim_{A_1} -favored. In other words, $\mu_{A_{i+1}}^*(A_i) = 1$ for each $i \leq n - 1$ and $\mu_{A_1}^*(A_n) = 1$. Note that $\mu_{A_{i+1}}^*(A_i) = 1$ means that $\mu_{A_{i+1}}^* \in \Delta(A_i)$. Since $\mu_{A_i}^*$ is the unique minimizer of d_μ in $\Delta(A_i)$, we have $d_\mu(\mu_{A_i}^*) \leq d_\mu(\mu_{A_{i+1}}^*)$, the inequality is strict when $\mu_{A_i}^* \neq \mu_{A_{i+1}}^*$. We will obtain a contradiction if there is at least one strict inequality. Hence, $\mu_{A_1}^* = \dots = \mu_{A_n}^*$, which implies $\succsim_{A_1} = \succsim_{A_n}$.

To prove Relative Tradeoff Consistency, take any $A, B \in \Sigma$, $p, q \in \Delta$, and $\alpha \in (0, 1)$ such that

$$w A q \sim p \text{ and } w A q \sim_A \alpha p + (1 - \alpha)w;$$

equivalently, $\mu(A^c)u(q) = u(p)$ and $\mu_A(A^c)u(q) = \alpha u(p)$. Since $\mu_A(A^c) = \delta \mu(A)$, we have $\alpha = \delta$. Take any r such that $w B q \sim r$; equivalently, $\mu(B^c)u(q) = u(r)$. Since $\mu_B(B^c) = \alpha \mu(B)$, we have $\mu_B(B^c)u(q) = \alpha u(r)$; equivalently,

$$w B q \sim_B \alpha r + (1 - \alpha)w.$$

□

References

- AFRIAT, S. N. (1967): “The construction of utility functions from expenditure data,” *International economic review*, 8, 67–77. [2](#), [9](#), [27](#)
- ALCHOURRÓN, C. E., P. GÄRDENFORS, AND D. MAKINSON (1985): “On the logic of theory

- change: Partial meet contraction and revision functions.” *The Journal of Symbolic Logic*, 50. 23
- ANSCOMBE, F. AND R. AUMANN (1963): “A Definition of Subjective Probability,” *Annals of Mathematical Statistics*, 34, 199–205. 4, 8
- BARBERIS, N., A. SHLEIFER, AND R. VISHNY (1998): “A model of investor sentiment,” *Journal of financial economics*, 49, 307–343. 23
- BASU, P. (2019): “Bayesian updating rules and AGM belief revision,” *Journal of Economic Theory*, 179, 455 – 475. 23
- BENJAMIN, D. J. (2019): “Errors in probabilistic reasoning and judgment biases,” in *Handbook of Behavioral Economics: Applications and Foundations*, ed. by B. D. Bernheim, S. DellaVigna, and D. Laibson, North-Holland, vol. 2, chap. 2, 69–186. 2, 8, 23
- BLUME, L., A. BRANDENBURGER, AND E. DEKEL (1991): “Lexicographic Probabilities and Choice under Uncertainty,” *Econometrica*, 59, 61–79. 24
- BORDALO, P., K. COFFMAN, N. GENNAIOLI, AND A. SHLEIFER (2016): “Stereotypes,” *The Quarterly Journal of Economics*, 131, 1753–1794. 7, 23
- BRANDENBURGER, A., A. FRIEDENBERG, AND H. J. KEISLER (2006): “Notes on the relationship between strong belief and assumption,” *Unpublished Manuscript. Available at www.stern.nyu.edu/~abranden.[319]*. 24
- CAMERER, C. (1987): “Do Biases in Probability Judgement Matter in Markets? Experimental Evidence,” *American Economic Review*, 79. 2
- (1995): “Individual Decision Making,” in *The Handbook of Experimental Economics*, ed. by J. H. Kagel and A. E. Roth, Princeton University Press. 2, 23
- CAPLIN, A. AND J. LEAHY (2019): “Wishful Thinking,” NBER Working Paper 25707, National Bureau of Economic Research. 7
- DE CLIPPEL, G. AND X. ZHANG (2022): “Non-bayesian persuasion,” *Journal of Political Economy*, 130, 2594–2642. 20, 24
- DOMINIAK, A., M. KOVACH, AND G. TSERENJIGMID (2022): “Minimum Distance Belief Updating with General Information,” *working paper*. 23
- EIL, D. AND J. M. RAO (2011): “The Good News-Bad News Effect: Asymmetric Processing of Objective Information about Yourself,” *American Economic Journal: Microeconomics*, 3, 114–138. 2
- EPSTEIN, L. G. (2006): “An Axiomatic Model of Non-Bayesian Updating,” *Review of Economic Studies*. 24
- EPSTEIN, L. G. AND M. L. BRETON (1993): “Dynamically Consistent Beliefs Must Be Bayesian,” *Journal of economic theory*, 61, 1–22. 10
- EPSTEIN, L. G., J. NOOR, AND A. SANDRONI (2008): “Non-Bayesian Updating: A Theoretical Framework,” *Theoretical Economics*. 24

- FOX, C. R. AND R. T. CLEMEN (2005): “Subjective Probability Assessment in Decision Analysis: Partition Dependence and Bias Toward the Ignorance Prior,” *Management Science*, 51, 1417–1432. 8
- GALPERTI, S. (2019): “Persuasion: The art of changing worldviews,” *American Economic Review*, 109, 996–1031. 8, 24
- GENNAIOLI, N. AND A. SHLEIFER (2010): “What comes to mind,” *The Quarterly journal of economics*, 125, 1399–1433. 7, 23
- GHIRARDATO, P. (2002): “Revisiting Savage in a conditional world,” *Economic Theory*, 20, 83–92. 10, 28
- GREYER, D. M. (1980): “Bayes Rule as a Descriptive Model: The Representativeness Heuristic,” *Quarterly Journal of economics*. 1, 3, 7, 18
- KAHNEMAN, D. AND A. TVERSKY (1972): “The Concept of Probability in Psychological Experiments pp 25–48 Cite as Subjective Probability: A Judgment of Representativeness,” *Cognitive Psychology*, 3. 2
- (1979): “Prospect theory: An analysis of decision under risk,” *Econometrica: Journal of the Econometric Society*, 263–291. 7
- (1983): “Extensional vs. Intuitive Reasoning: The Conjunction Fallacy in Probability Judgement,” *Psychological review*, 90. 2
- KAMENICA, E. (2019): “Bayesian persuasion and information design,” *Annual Review of Economics*, 11, 249–272. 20
- KAMENICA, E. AND M. GENTZKOW (2011): “Bayesian persuasion,” *American Economic Review*, 101, 2590–2615. 4, 19, 20
- KOVACH, M. (2020a): “Conservative Updating,” *mimeo*. 24
- (2020b): “Twisting the Truth: Foundations of Wishful Thinking,” *Theoretical Economics*, 15, 989–1022. 7
- KREPS, D. M. AND R. WILSON (1982): “Reputation and imperfect information,” *Journal of Economic Theory*, 27, 253–279. 3, 13, 15
- KUNDA, Z. (1990): “The Case for Motivated Reasoning,” *Psychological Bulletin*. 7
- LEE, Y.-J., W. LIM, AND C. ZHAO (2023): “Cheap talk with prior-biased inferences,” *Games and Economic Behavior*. 24
- MATZKIN, R. L. (1991): “Axioms of revealed preference for nonlinear choice sets,” *Econometrica: Journal of the Econometric Society*, 1779–1786. 3, 9, 27
- MAYRAZ, G. (2011): “Wishful Thinking,” *working paper available at SSRN: <https://ssrn.com/abstract=1955644>*. 7
- MULLAINATHAN, S. (2002): “A memory-based model of bounded rationality,” *The Quarterly Journal of Economics*, 117, 735–774. 7, 23
- MULLAINATHAN, S., J. SCHWARTZSTEIN, AND A. SHLEIFER (2008): “Coarse thinking and

- persuasion,” *The Quarterly journal of economics*, 123, 577–619. 23
- MYERSON, R. B. (1986a): “Axiomatic Foundations of Bayesian Decision Theory,” Discussion Papers 671, Northwestern University, Center for Mathematical Studies in Economics and Management Science. 1, 2, 3, 8, 10, 13, 24
- (1986b): “Multistage Games with Communication,” *Econometrica*, 54, 323–358. 1, 2, 3, 8, 10, 13, 24
- ORTOLEVA, P. (2012): “Modeling the Change of Paradigm: Non-Bayesian Reactions to Unexpected News,” *American Economic Review*, 102, 2410–36. 1, 2, 8, 9, 15, 16, 23, 24
- PEREA, A. (2009): “A Model of Minimal Probabilistic Belief Revision,” *Theory and Decision*, 163–222. 23
- RABIN, M. (2002): “Inference by believers in the law of small numbers,” *The Quarterly Journal of Economics*, 117, 775–816. 23
- RABIN, M. AND J. L. SCHRAG (1999): “First Impressions Matter: A Model of Confirmatory Bias,” *Quarterly Journal of Economics*, 114, 37–82. 7, 23
- SULEYMANOV, E. (2021): “Robust Maximum Likelihood Updating,” *working paper*. 24
- VARIAN, H. R. (1982): “The nonparametric approach to demand analysis,” *Econometrica: Journal of the Econometric Society*, 945–973. 3
- WILSON, A. (2014): “Bounded memory and biases in information processing,” *Econometrica*, 82, 2257–2294. 7
- ZHAO, C. (2022): “Pseudo-Bayesian Updating,” *Theoretical Economics*, 17. 23