Coarse Bayesian Updating

Alexander M. Jakobsen*

February 28, 2021

Abstract

I introduce a model of belief updating—*Coarse Bayesian updating*—where, upon receipt of new information, an agent applies subjective criteria to select among competing theories of the world. The agent is characterized by a partition of the probability simplex and a representative distribution for each cell of the partition. When information arrives, the agent determines which cell contains the Bayesian posterior and adopts its representative as posterior belief. I characterize this procedure, analyze how it relates to existing models and evidence on non-Bayesian updating, and apply it to a standard setting of decision under risk.

1 Introduction

Bayesian updating plays a central role in economic theory. A wide body of evidence, however, suggests that actual behavior cannot be reconciled with Bayes' rule in a variety of settings. For example, individuals often display *conservatism bias*: they under-react to new evidence, possibly ignoring it altogether. Others *overreact* to information by falsely extrapolating or, more generally, engaging in pattern-seeking behavior. Combinations of these forces may lead individuals to under-weight some signals while over-weighting others. In this paper, I introduce and analyze a simple generalization of Bayesian updating—*Coarse Bayesian updating*—accommodating these, and other, behavioral tendencies.

*Department of Economics, University of Calgary; alexander.jakobsen@ucalgary.ca. I am grateful to Shaowei Ke, Dimitri Migrow, Rob Oxoby, and Benjamin Young for helpful conversations and feedback. I also thank seminar audiences at Michigan, UNSW, and UTS as well as participants at CETC (McGill, 2019), RUD (PSE, 2019), and SAET 2019 for their comments. Alaina Olson provided excellent research assistance. This research was supported by a SSHRC Insight Development Grant.

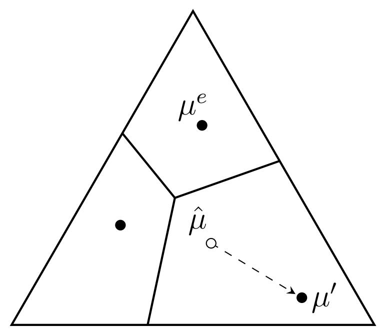

Figure 1: Coarse Bayesian Updating. In this example, the agent entertains three distributions (solid dots). The point μ^e is the prior. After observing a signal s, the agent determines which cell of the partition contains the Bayesian posterior $\hat{\mu}$, then adopts the representative of that cell (in this case, μ') as his new belief.

At its core, Coarse Bayesian updating is a model of bounded rationality stemming from a single key assumption: agents simplify the world by considering only a subset of the probability space. Members of this set represent competing theories or beliefs, and agents apply subjective criteria to select among them. More precisely, a Coarse Bayesian agent is characterized by a partition of the set of all probability distributions over a state space, together with a representative distribution for each cell of the partition. One of the distributions is the prior. After observing a signal, the agent determines which cell contains the Bayesian posterior and adopts the representative of that cell as his posterior belief (see Figure 1). Since realized posteriors typically differ from their Bayesian counterparts, Coarse Bayesians may exhibit overreaction, under-reaction, or other biases depending on the realized signal, the shape of the partition, and the positions of representative distributions within their cells.

The model can be interpreted in different ways. First, as indicated above, one might view the representative distributions as competing *theories of the world* and the partition as the agent's criteria, or *standard of proof*, for selecting among them. In this interpretation, agents correctly assess the informational content of signals and decide which, if any, competing theory ought to replace their prior. Given that the agent entertains a restricted set of theories, the updating procedure is a minimal deviation from Bayes' rule: the agent must have *some* criteria for choosing among feasible theories, and will adopt a given theory if it coincides with the Bayesian posterior.

Second, one can interpret Coarse Bayesian behavior as the result of *signal dis-*

tortion. Rather than actively weighing competing theories against the evidence, the agent mentally transforms the observed signal before applying Bayes' rule to the modified signal. Thus, in this interpretation, violations of Bayes' rule are the result of imperfect perception or attention.

Finally, Coarse Bayesian updating can be interpreted as a form of categorical thinking. Here, each cell of the partition represents a category of beliefs, and the representative of a cell an "archetype" of that category. Thus, the agent uses information to select a category, then adopts the archetype of that category. This interpretation can also be applied at the level of signals: the agent groups signals into categories, and updates beliefs based on the category of the realized signal.

In all cases, the parameters of the model are subjective characteristics of the individual: two Coarse Bayesians may differ in their sets of feasible beliefs, their partitions, or both. In contrast to the canonical framework of Savage (1972), then, Coarse Bayesians exhibit subjectivity not only in their prior beliefs, but also in their criteria for revising those beliefs. Some agents may tend to disregard evidence while others falsely extrapolate from it; some might be biased in favor of a particular theory, while others seek to discredit it. Even with a common prior, compelling evidence in the eyes of one agent may be completely unpersuasive for another, or result in radically different posterior beliefs. In general, Coarse Bayesians may disagree on the strength of evidence required to adopt a particular belief, or on the set of admissible beliefs to begin with, so that a given piece of evidence may yield (or magnify) disagreement among individuals.

Section 2 provides a simple characterization of Coarse Bayesian updating. I take as primitive a finite, exogenous state space and an updating rule specifying an individual's probabilistic beliefs at every possible signal. In my framework, signals represent messages that can be generated by stochastic information structures. Thus, a signal consists of a profile of numbers indicating the likelihood of the associated message being generated in different states. By employing such primitives, the model is readily adaptable to any standard economic or game-theoretic setting.

The characterization involves three testable assumptions on the updating rule. The first, *Homogeneity*, states that beliefs are invariant to scalar transformations of signals. Thus, like Bayes' rule, Coarse Bayesian updating rules only depend on the likelihood ratios of the observed signal. Second, *Cognizance* states that if two signals result in the same belief, then so does a "garbled" signal indicating that one of those signals was generated. A natural interpretation of this assumption is that the agent understands, or is cognizant of, his own updating procedure: if he is uncertain about which of two signals was generated, but recognizes that each would lead to the same posterior belief, then he adopts that belief. Finally, *Confirmation* states that if a signal exactly supports (or confirms) some feasible belief, then the updating rule associates that belief to the given signal. Theorem 1 establishes that an updating rule has a Coarse Bayesian representation if and only if it is Homogeneous, Cognizant, and Confirmatory. Moreover, the associated partition, representative elements, and prior are unique. Theorem 2 establishes that an updating rule has a Coarse Bayesian representation if and only if it has a *Signal Distortion* representation satisfying three analogous properties, thereby formalizing the second interpretation of the model described above.

Section 3 explores the main implications of Coarse Bayesian updating and examines connections to related models and evidence. In section 3.1, I discuss evidence on biased belief updating and demonstrate how Coarse Bayesian models can accommodate such behavior. In addition to under- and over-reaction, I show how Coarse Bayesians may exhibit "motivated" belief updating, limited perception, extreme-belief aversion, or susceptibility to logical fallacies. Section 3.2 examines the relationship to "paradigm shifts," including the Hypothesis-Testing model of Ortoleva (2012). I show that Coarse Bayesian updating can accommodate similar behavior and, motivated by this fact, examine whether Coarse Bayesians can be represented as Bayesians with second-order priors. I show that such models, dubbed *Maximum-Likelihood* updating rules, intersect the class of Coarse Bayesian rules but that neither class subsumes the other—unless there are exactly two states, in which case every Coarse Bayesian rule can be expressed as a Maximum-Likelihood rule. Finally, section 3.3 explores some basic properties of Coarse Bayesian updating in dynamic settings. I show that most Coarse Bayesians (other than standard Bayesians) are sensitive to the way histories of signals are pooled and ordered. However, this depends on how the updating rule is represented: if one employs a Signal Distortion rule, rather than its associated partitional representation in probability space, then signals can be reordered without affecting terminal beliefs. Thus, the equivalence between Coarse Bayesian (partitional) representations and their associated Signal Distortion representations only holds in static settings.

Section 4 applies the model to a standard setting of decision under risk. In particular, I analyze how Coarse Bayesians value information (Blackwell experiments) when faced with menus of actions with state-dependent payoffs. I show that a Coarse Bayesian's ex-ante value of information can be expressed in a familiar posterior-separable form, making the model amenable to the well-known Bayesian Persuasion framework (Kamenica and Gentzkow, 2011). Then, I establish that standard Bayesians have a higher value of information than Coarse Bayesians and that, unlike Bayesians, Coarse Bayesians typically exhibit violations of the Blackwell (1951) information ordering—they need not assign higher ex-ante value to more informative experiments. Finally, I examine how a Coarse Bayesian's value of information changes as he becomes more sophisticated, or "more Bayesian," in that the set of feasible beliefs expands and the associated partition becomes finer. I characterize the sophistication ordering by showing that more sophisticated agents exhibit more responsiveness to information, as captured by ex-ante value of information. Moreover, greater sophistication increases both the value of information and the degree of adherence to the Blackwell ordering in settings where it results in more "agreement" with Bayesian decision makers.[1]

1.1 Related Literature

Economists and psychologists have developed a large body of research documenting systematic violations of Bayesian updating; early contributions include Kahneman and Tversky (1972), Tversky and Kahneman (1974), and Grether (1980). As seen in the surveys of Camerer (1995), Rabin (1998), and Benjamin (2019), there is substantial variation in both the patterns of behavior displayed by subjects and the settings in which experiments are carried out. For example, studies differ in whether subjects observe individual pieces of information or larger samples (or sequences) of evidence; whether prior beliefs are objectively induced or subjectively formed by participants; whether choices are incentivized with monetary rewards; and how problems and information are framed.

Motivated by this evidence, several authors have developed models to better understand the mechanisms behind, and consequences of, non-Bayesian updating. Mod-

[1] In an experimental setting, Ambuehl and Li (2018) find that subjects tend to undervalue improvements to instrumentally valuable information, and argue that this is due to non-Bayesian belief updating. In addition, subjects differ in their responsiveness to information. Coarse Bayesian updating, which permits different agents to employ different updating rules, yields similar predictions.

els focusing on implications of biased updating are typically cast in simplified frameworks (eg, two states of the world; particular protocols or functional form assumptions), or involve non-standard elements like ambiguous signals or framing effects. See, among others, Barberis et al. (1998), Fryer et al. (2019), Gennaioli and Shleifer (2010), Rabin and Schrag (1999), and Mullainathan et al. (2008).

Decision theorists have developed axiomatic approaches to non-Bayesian updating. Kovach (2020), for example, develops a model of conservative updating. Epstein (2006) provides a model of non-Bayesian updating accommodating under-reaction, overreaction, and other biases; Epstein et al. (2008) extend this model to an infinite-horizon setting. Zhao (2016) axiomatizes a particular updating rule for signals indicating that one event is more likely than another.

Three studies are especially relevant to Coarse Bayesian updating. First, the *hypothesis testing* model introduced and axiomatized by Ortoleva (2012) posits that agents apply standard Bayesian updating except when news is sufficiently "surprising," in which case posterior beliefs are selected by applying a maximum-likelihood criterion to a second-order prior. In particular, an agent applies Bayes' rule if the prior probability of the observed signal weakly exceeds a threshold $\varepsilon \geq 0$; otherwise, the agent updates a second-order prior via Bayes' rule and selects a belief of maximal probability under the new second-order beliefs. In section 3.2, I show that Coarse Bayesian updating can accommodate similar behavior, and compare Coarse Bayesian updating rules to a general class of Maximum-Likelihood updating rules. Importantly, Maximum-Likelihood rules may violate the Confirmation property—perfect evidence for a candidate belief does not guarantee that that belief is adopted.

Second, Wilson (2014) studies optimal updating rules for a boundedly rational agent facing a binary decision problem and a stochastic sequence of signals. There are two states, and the agent has limited memory: only K memory states are available. In an optimal protocol, each memory state is associated with a convex set of posterior beliefs and a representative distribution for that set; if an interim Bayesian posterior belongs to some cell, then the representative of that cell is adopted as the agent's belief. Thus, the optimal protocol emerging from Wilson's model is a Coarse Bayesian updating procedure. Naturally, the parameters of this representation (cells and their representative points) depend on features of the environment such as the signal structure and the bound K. Coarse Bayesian updating procedures—like standard Bayesian updating—do not depend on any factors other than the informational content of realized signals. In particular, I do not require Coarse Bayesian representations to be optimal in any sense, nor do I impose cognitive bounds such as a restriction on the number of cells. One could endogenize Coarse Bayesian updating rules in various ways. For example, by introducing costs (or bounds) on the number of cells and fixing both an information structure and subsequent decision problem, one could focus on representations that maximize ex-ante expected utility subject to those constraints (see also footnote 6 in section 3.1 for a slightly different direction).

Third, a working paper, Mullainathan (2002), develops a model of categorical thinking. Agents in this model follow a procedure similar to Coarse Bayesian updating where feasible posteriors represent categories and the mapping from Bayesian posteriors to categories is determined by a partition of the simplex. A key difference is that the partition is derived from the set of feasible posteriors: given a set of feasible posteriors, an optimality condition similar in spirit to maximization of a likelihood function is used to select a posterior. The resulting partition has convex cells, as in a Coarse Bayesian representation, but cells need not contain their representative elements. In other words, behavior in this model need not satisfy Confirmation—see appendix B for an explicit example.

2 Model

I consider a single agent who updates beliefs after observing a noisy signal. Let Ω denote a finite set of $N \geq 2$ states and Δ the set of probability distributions over Ω. A distribution $\hat{\mu} \in \Delta$ assigns probability $\hat{\mu}_\omega$ to state $\omega \in \Omega$.

An **experiment** is a matrix σ with entries in $[0, 1]$ and N rows where each row is a probability distribution and each column has at least one nonzero entry. Each column represents a message that might be generated, and each row represents a state-contingent probability distribution over messages. Let $\mathcal{E}$ denote the set of all experiments.

As in Jakobsen (2020), a **signal** is a profile $s = (s_\omega)_{\omega\in\Omega} \in [0, 1]^\Omega$ such that $s_\omega \neq 0$ for at least one state ω. Let S denote the set of all signals. Intuitively, a signal represents a column (message) of some experiment, and its entries s_ω are the likelihoods of the message being generated in different states of the world. The notation $s \in \sigma$ indicates that s is a column of σ. I reserve e to denote the **uninformative signal**; that is, $e \in S$ and $e_\omega = 1$ for all $\omega \in \Omega$. Note that e qualifies as an experiment.

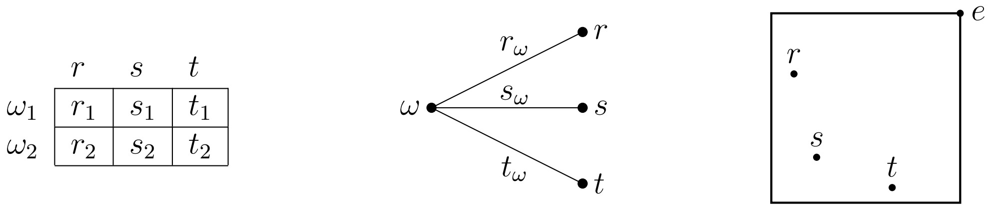

(a) Matrix Form (b) Tree Form (c) Signal Space

Figure 2: Three representations of an experiment $\sigma = [r, s, t]$.

Using the notation of signals, any experiment can be viewed as a collection (matrix) of signals, as a series of state-contingent distributions over signals, or as a collection of points in S; see Figure 2.

For two profiles $v = (v_\omega)_{\omega\in\Omega}$ and $w = (w_\omega)_{\omega\in\Omega}$ of real numbers, let $vw := (v_\omega w_\omega)_{\omega\in\Omega}$ denote the profile formed by multiplying v and w component-wise. Similarly, if $w_\omega > 0$ for all ω, let $v/w := (v_\omega/w_\omega)_{\omega\in\Omega}$. The dot product of v and w is given by $v \cdot w := \sum_{\omega\in\Omega} v_\omega w_\omega$. The notation $v \approx w$ indicates that $v = \lambda w$ for some $\lambda > 0$, where $\lambda w := (\lambda w_\omega)_{\omega\in\Omega}$ is the scalar product of λ with w. The standard Euclidean norm of v is denoted $\|v\|$. Clearly, if $\mu, \mu' \in \Delta$, then $\mu = \mu'$ if and only if $\mu \approx \mu'$.

If $\hat{\mu} \in \Delta$ and $s \in S$ such that $s\hat{\mu} \neq 0$, then $B(\hat{\mu}|s) := \frac{s\hat{\mu}}{s\cdot\hat{\mu}} \in \Delta$ denotes the **Bayesian posterior** of $\hat{\mu}$ at signal s. Note that $B(\hat{\mu}|s)$ is the unique $\mu' \in \Delta$ such that $\mu' \approx s\hat{\mu}$.

Finally, an **updating rule** is a function $\mu : S \to \Delta$ assigning probability distributions $\mu^s \in \Delta$ to signals $s \in S$. For each $s \in S$, μ^s is the agent's posterior belief conditional on observing signal s. I assume μ^e, the **prior**, has full support. Updating rules will often be written as profiles: $\mu = (\mu^s)_{s\in S}$.[2]

2.1 Coarse Bayesian Representations

Let $\mu : S \to \Delta$ be an updating rule such that μ^e has full support. I impose three testable assumptions on μ.

Assumption 1 (Homogeneity). If $s \approx t$, then $\mu^s = \mu^t$.

[2]Note that updating rules condition beliefs on signal realizations s but not on experiments σ. In practice, a signal must be generated by an experiment, in which case one may wish to denote posterior beliefs by $\mu^{(\sigma,s)}$ where s is a column of σ. Like standard Bayesian updating, however, Coarse Bayesian updating depends on s but not the other columns of σ. To minimize notation, I have omitted the underlying experiment(s) σ.

Homogeneity requires the agent's analysis of a signal to depend only on the likelihood ratios $s_\omega/s_{\omega'}$. This is a key feature of standard Bayesian updating: $B(\mu^e|s)$ coincides with $B(\mu^e|\lambda s)$, provided $\lambda > 0$ and $\lambda s \in S$. An implication of this assumption is that the agent is not susceptible to certain types of framing effects. For example, whether information is stated in terms of frequencies or likelihoods has no effect on the agent's cognitive process.

By Homogeneity, the notation μ^s can be extended to all non-zero profiles $\tilde{s}$ such that $\tilde{s}_\omega \geq 0$ for all ω because such profiles can be scaled by a factor $\lambda > 0$ to yield a signal $\lambda\tilde{s} \in S$. This will be convenient for expressing the remaining assumptions.

Assumption 2 (Cognizance). If $\mu^s = \mu^t$, then $\mu^{s+t} = \mu^s$.

Cognizance states that if signals s and t result in the same posterior belief, then the agent adopts that belief if he knows that either s or t was generated. This holds because $s + t$ represents a "garbled" signal indicating that either s or t has realized. Thus, an interpretation of Cognizance is that *the agent understands his own updating rule*: if he knows that one of two signals was generated and realizes that either one would lead him to the same posterior belief—that is, if he is cognizant of his own updating procedure—then he ought to adopt that belief. Although Cognizance is mainly motivated by normative considerations, it is also potentially important in applications. For example, section 4 studies how Coarse Bayesians value information. This involves ex-ante rankings of information structures that rely on correct forecasts about updating behavior. For such exercises to make sense, an assumption like Cognizance is required.

Assumption 3 (Confirmation). For all s, $\mu^{\mu^s/\mu^e} = \mu^s$.

To understand Confirmation, observe that for any s, μ^s is a feasible posterior because it is in the range of the updating rule. Moreover, any signal $t \approx \mu^s/\mu^e$ satisfies $B(\mu^e|t) = \mu^s$. Thus, Confirmation states that if a signal t exactly supports (or confirms) a feasible posterior, then the agent adopts that posterior after observing t. Although quite intuitive and normatively appealing, this property is not always satisfied by some closely-related models—see section 3.2 and appendix B.

Theorem 1. *An updating rule μ is Homogeneous, Cognizant, and Confirmatory if and only if there is a partition $\mathcal{P}$ of Δ and a profile $\mu^{\mathcal{P}} = (\mu^P)_{P\in\mathcal{P}}$ of distributions* *such that*

(i) each cell $P \in \mathcal{P}$ is convex,

(ii) for all $P \in \mathcal{P}$, $\mu^P \in P$, and

(iii) for all $s \in S$, $B(\mu^e|s) \in P$ implies $\mu^s = \mu^P$.

Such a pair $\langle \mathcal{P}, \mu^{\mathcal{P}} \rangle$ is a ***Coarse Bayesian Representation*** *of μ. If $\langle \mathcal{Q}, \mu^{\mathcal{Q}} \rangle$ is another Coarse Bayesian Representation of μ, then $\mathcal{P} = \mathcal{Q}$ and $(\mu^P)_{P \in \mathcal{P}} = (\mu^Q)_{Q \in \mathcal{Q}}$.*

Proof. First, observe that if $\alpha, \beta \geq 0$ and $s, t, \alpha s + \beta t \in S$, then

$$\begin{aligned} B(\mu^e|\alpha s + \beta t) &= \frac{(\alpha s + \beta t)\mu^e}{(\alpha s + \beta t) \cdot \mu^e} \\ &= \frac{\alpha s \cdot \mu^e}{(\alpha s + \beta t) \cdot \mu^e} \frac{s\mu^e}{s \cdot \mu^e} + \frac{\beta t \cdot \mu^e}{(\alpha s + \beta t) \cdot \mu^e} \frac{t\mu^e}{t \cdot \mu^e} \\ &= \frac{\alpha s \cdot \mu^e}{(\alpha s + \beta t) \cdot \mu^e} B(\mu^e|s) + \frac{\beta t \cdot \mu^e}{(\alpha s + \beta t) \cdot \mu^e} B(\mu^e|t). \end{aligned} \tag{1}$$

Thus, $B(\mu^e|\alpha s + \beta t)$ is a convex combination of $B(\mu^e|s)$ and $B(\mu^e|t)$; the weight attached to $B(\mu^e|s)$ is the prior probability of signal αs given that either αs or βt is generated. It is now straightforward to verify that if μ has a Coarse Bayesian Representation, then Assumptions 1–3 are satisfied (Assumption 2 follows from equation (1) and convexity of cells $P \in \mathcal{P}$).

For the converse, we construct a Coarse Bayesian Representation as follows. First, note that Homogeneity and Cognizance imply μ is **Convex**: if $\mu^s = \mu^t$ and $\alpha \in [0, 1]$, then $\mu^{\alpha s + (1-\alpha)t} = \mu^s$. It follows that μ is measurable with respect to a partition of S into convex cones. That is, there is a partition $\mathcal{C}$ of S such that (i) $\mu^s = \mu^t$ if and only if there exists $C \in \mathcal{C}$ such that $s, t \in C$, and (ii) every $C \in \mathcal{C}$ is a convex cone: if $s, t \in C$ and $\alpha, \beta \geq 0$ such that $\alpha s + \beta t \in S$, then $\alpha s + \beta t \in C$. Every $C \in \mathcal{C}$ can be identified with a subset of Δ by letting $P^C := \{B(\mu^e|s) : s \in C\}$. Each set P^C is convex by (1) and the fact that sets $C \in \mathcal{C}$ are convex cones. In addition, $\mathcal{P} := \{P^C : C \in \mathcal{C}\}$ is a partition of Δ because $B(\mu^e|s) = B(\mu^e|t)$ if and only if $s \approx t$, forcing s and t to belong to the same cone $C \in \mathcal{C}$. Cognizance implies $\mu^s \in P$ whenever $B(\mu^e|s) \in P \in \mathcal{P}$. □

Theorem 1 formalizes the concept of a Coarse Bayesian Representation and establishes that an updating rule has such a representation if and only if it is Homogeneous, Cognizant, and Confirmatory. Each of these testable assumptions expresses some element of "rational" information processing—indeed, each assumption is satisfied by a standard Bayesian. Nonetheless, the class of Coarse Bayesian Representations accommodates a large variety of behavioral biases and other violations of Bayes' rule. As we shall see, the three assumptions also help make meaningful comparisons to related models.

As described in the introduction, Coarse Bayesian updating captures the behavior of an agent who partitions the probability simplex, assigns a representative distribution to each cell, and adopts the representative of a cell as posterior if the Bayesian posterior belongs to that cell. This enables the following interpretations.

1. *Competing Theories.* Here, representative points constitute competing theories of the world, while the partition summarizes the agent's standard of proof for selecting among them. Such agents do not necessarily compute Bayesian posteriors of signals: they only need to know which cell contains the Bayesian posterior, so fully computing that posterior may not be necessary. Thus, agents do not deviate from Bayes' rule due to the "difficulty" of computing Bayesian posteriors; rather, they simplify the world by considering a restricted set of feasible theories, and analyze signals to the extent necessary to determine whether they should adopt a different theory.

2. *Signal Distortions.* Alternatively, one can interpret the procedure as the result of signal distortion: rather than selecting among feasible theories, agents mentally transform signals before applying Bayes' rule to update their prior. Thus, apparent deviations from Bayes' rule are the result of imperfect perception or attention. In the next section, I formalize such procedures and show that they are equivalent to Coarse Bayesian updating rules.

3. *Categorical Thinking.* In this interpretation, the agent reasons about categories of beliefs, each represented by a cell of the partition. The representative μ^P of cell P is an "archetype" of that category. When information arrives, the agent determines which category applies and adopts its archetype as posterior. Again, such agents need not fully compute Bayesian posteriors—they only glean enough information from a signal to figure out which category contains the Bayesian posterior. Signal distortion rules can also be interpreted as a form of categorical thinking: agents classify signals into categories, and update based on archetypes of those categories.

Although Coarse Bayesian updating accommodates many documented departures from Bayesian updating (see section 3), the model is not without its limitations. First, like Bayes' rule, the procedure requires an agent's posterior belief to depend only on the realized signal s. More precisely, only the ratios of entries in s can affect posterior beliefs. This eliminates behavior where agents are sensitive to the way information is framed, and rules out the possibility that agents might be affected by other signals that could have been generated by the underlying experiment σ.

Coarse Bayesian updating procedures are also not sensitive to "stakes." Like Bayesians, Coarse Bayesians update beliefs without regard to whatever decision problem they may be facing; their analysis of information is completely separated from whatever gains or losses they may incur as a result of that analysis.[3]

A more technical matter is that Coarse Bayesian updating rules are typically discontinuous in s. If continuity is an essential conceptual feature of some pattern of behavior—rather than a convenient technical assumption—then Coarse Bayesian updating procedures will, at most, provide an approximation to that behavior.

Finally, requiring cells of the partition to be convex might seem limiting. This convexity is driven by Cognizance and can be discarded by dropping that assumption. However, as explained above, Cognizance is potentially important in applications because it means agents correctly forecast their own updating behavior.

2.2 Signal Distortions

In this section, I provide an alternative representation of Coarse Bayesian updating rules. The idea, formalized by the next definition, is that non-Bayesian reactions are due to errors or biases in the agent's perception of information.

Definition 1. An updating rule μ has a **Signal Distortion Representation** if there is a function $d : S \to S$ (a **signal distortion**) such that $\mu^s = B(\mu^e | d(s))$ for all $s \in S$.

In a Signal Distortion Representation, an agent who observes signal s updates beliefs by applying Bayes' rule with a modified signal $d(s)$. Thus, the function d is a behavioral parameter capturing the agent's tendency to distort information. Such

[3] See Balzer and Young (2020) for a model of updating where, at an interim stage, an agent weighs costs of additional reasoning against benefits in a subsequent decision problem.

distortions could be due to imperfections in the agent's perception or, in some applications, may reflect the agent's beliefs about the accuracy or reliability of the information source.[4]

Note that if μ has a Signal Distortion Representation, then $d(e) \approx e$. This is the only substantive property of d implied by such a representation. Therefore, without additional restrictions, the concept of signal distortion can explain almost any updating behavior. The next definition provides three restrictions on d that are needed to establish an equivalence between Signal Distortion and Coarse Bayesian Representations.

Definition 2. A signal distortion d is

(i) **Homogeneous** if $d(s) \approx d(t)$ whenever $s \approx t$.

(ii) **Convex** if $d(s) \approx d(t)$ implies $d(\lambda s + (1-\lambda)t) \approx d(s)$ for all $\lambda \in [0, 1]$.

(iii) **Idempotent** if $d(d(s)) = d(s)$ for all s.

Homogeneity states that two distorted signals have common likelihood ratios if the original signals have common likelihood ratios. Convexity states that if two signals have common distorted likelihood ratios, then those ratios also come about from distorting mixtures of those signals. Idempotency requires the distortion process to be stable: the distortion of $d(s)$ is $d(s)$.

As illustrated by Figure 3, signal distortions d that are Homogeneous, Convex, and Idempotent effectively categorize signals and assign common distorted likelihood ratios to signals in the same category. In particular, d gives rise to a partition of S into convex cones, along with a representative ray for each cell. Signals in a given cell get distorted to points along the representative ray, ensuring common likelihood ratios.

Theorem 2. *An updating rule has a Coarse Bayesian Representation if and only if it has a Homogeneous, Convex, Idempotent Signal Distortion Representation. If* d *and* d' *are two such representations of a given updating rule, then* $d(s) \approx d'(s)$ *for all* $s \in S$.

[4] See also Aydogan et al. (2017), who propose a model of signal distortion similar in spirit to that of Rabin and Schrag (1999).

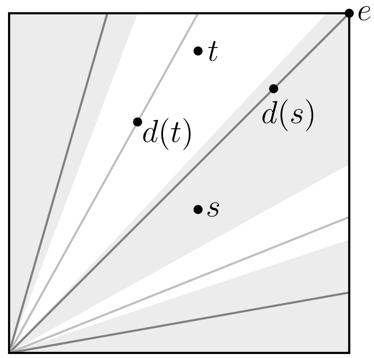

Figure 3: A Signal Distortion Representation on S for two states.

Proof. First, suppose μ has a Coarse Bayesian Representation. Note that for every $s \in S$ the signal $\frac{\mu^s/\mu^e}{\|\mu^s/\mu^e\|}$ is well-defined because μ^e has full support. Define $d : S \to S$ by

$$d(s) = \begin{cases} s & \text{if } \mu^s = B(\mu^e|s) \\ \frac{\mu^s/\mu^e}{\|\mu^s/\mu^e\|} & \text{otherwise} \end{cases}.$$

It is straightforward to verify that $\mu^s = B(\mu^e|d(s))$ for all s and that d is Homogeneous, Convex, and Idempotent.

Conversely, suppose μ has a Signal Distortion Representation with Homogeneous, Convex, and Idempotent d. Define a binary relation $\sim$ on S by $s \sim t$ if and only if $d(s) \approx d(t)$. Clearly, $\sim$ is an equivalence relation; thus, its equivalence classes partition S. Homogeneity and Convexity of d ensure each equivalence class is a convex cone. Thus, as in the proof of Theorem 1, each equivalence class is associated with a convex subset of Δ, and these subsets form a partition $\mathcal{P}$ of Δ. For each cell $P \in \mathcal{P}$, let $\mu^P := B(\mu^e|d(s))$ such that s belongs to the equivalence class associated with P. By Idempotency, $\mu^P \in P$. $\square$

Theorem 2 establishes the sought-after equivalence between Coarse Bayesian and Signal Distortion Representations. Thus, any updating rule satisfying Homogeneity, Convexity, and Confirmation has such a Signal Distortion Representation, and the distortion d is unique up to scalar transformation.

3 Models, Evidence, and Implications

Coarse Bayesian updating is related to a number of other theories of non-Bayesian updating, and accommodates a variety of experimental findings. In this section, I examine these relationships and explore some of the main implications of Coarse Bayesian updating.

3.1 Bias, Asymmetry, and Perception

1. Asymmetric Updating. Conservative updating, or under-reaction to information, is a well-documented behavior violating Bayes' rule.[5] On the other hand, many individuals also overreact to information in various settings. For example, the concept of base-rate neglect introduced and identified by Kahneman and Tversky (1973) exhibits a form of overreaction where individuals over-weight information relative to their priors. Rather than always under-reacting or always overreacting, individuals may respond asymmetrically to information. Eil and Rao (2011), for example, find that when information concerns personal attributes such as attractiveness, individuals under-react to negative signals but are approximately Bayesian when processing positive signals.

For a Coarse Bayesian, responsiveness to information depends on the set of feasible beliefs, their positions within their cells, and the "strength" of the observed signal. Thus, although under- and over-reaction are rather opposite phenomena, a Coarse Bayesian typically exhibits both behaviors: he under-reacts to some signals, but over-reacts to others. If, for example, the cell P containing the prior μ^e is not a singleton, then the agent will not update his beliefs for signals that yield Bayesian posteriors in P—an under-reaction to new information. However, signals that do result in belief revision typically yield posterior beliefs that do not coincide with the Bayesian posterior, often resulting in over-reaction.

Naturally, the concepts of over- and under-reaction make the most intuitive sense in two-state settings, where the probability simplex Δ can be represented by the unit interval. Figures 4a and 4b illustrate under- and over-reaction in such a setting. In 4a, the agent never over-reacts but typically under-reacts: his posterior belief is as

[5] See Phillips and Edwards (1966) and Edwards (1968) for early experiments on conservative updating.

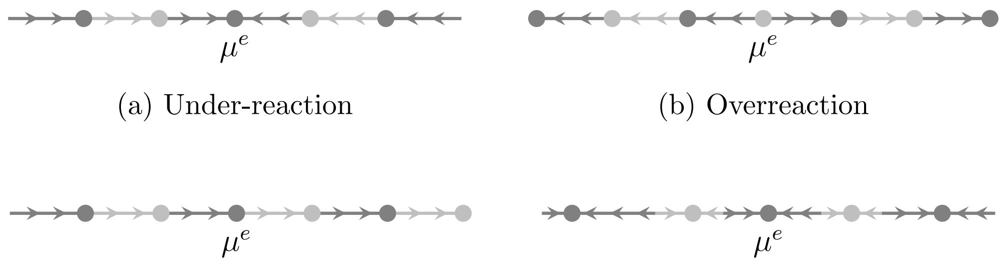

(a) Under-reaction

(b) Overreaction

(c) Upward Bias

(d) A "Typical" Coarse Bayesian

Figure 4: Four Coarse Bayesian Representations on $\Delta = [0, 1]$.

close as possible to μ^e given the partition of Δ into sub-intervals, resulting in conservative updating. In 4b, the agent never under-reacts but typically over-reacts: his posterior is farthest away from μ^e given the partition, resulting in a form of base-rate neglect. Figure 4c exhibits a biased agent who favors one state: posteriors typically assign higher probability to state 1 compared to the Bayesian posterior, but never less. Thus, it is relatively easy for this agent to revise his beliefs upward, but relatively difficult to revise them downward. This captures, for example, "motivated" reasoning where agents may place intrinsic value on their beliefs. Finally, Figure 4d depicts a "typical" Coarse Bayesian: representative points do not necessarily sit on the boundaries of cells, and therefore both over- and under-reaction occur.

2. Limited Perception and Extreme-Belief Aversion. Coarse Bayesian representations can also capture an agent's limited perception or attention. For example, consider Figure 5a. In this representation, the agent retains his prior μ^e unless the Bayesian posterior is sufficiently far away from μ^e, in which case he applies Bayes' rule. This captures the behavior of an agent who only notices signals that are sufficiently strong or provocative to yield a large shift in the Bayesian posterior (the associated Signal Distortion representation may be a more natural way of expressing and interpreting such behavior).[6]

Figure 5b exhibits rather the opposite behavior: the agent is Bayesian unless posterior beliefs are too "extreme"—that is, close to degenerate distributions representing

[6] The representation in Figure 5a need not be optimal in any sense, although one could endogenize it fairly easily. For example, one could introduce costs to shrinking the radius around the prior. Then, in a similar spirit to Brunnermeier and Parker (2005), such costs could be weighed against payoffs in a subsequent decision problem to yield an optimal representation. This would provide a different approach to rational inattention and/or costly information processing.

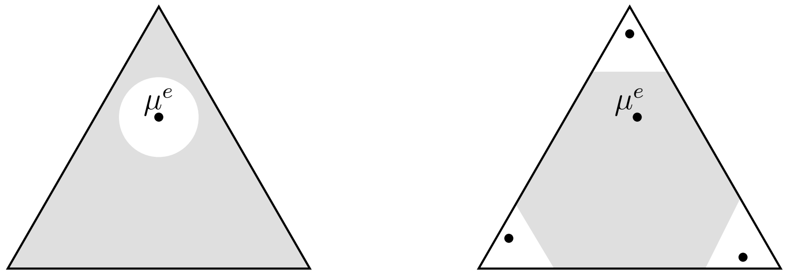

(a) Limited Perception (b) Extreme-Belief Aversion

Figure 5: Limited Perception and Extreme-Belief Aversion.

certainty about the state. Ducharme (1970) argues that such behavior may explain some of the experimental evidence on under-reaction (see also Benjamin et al. (2016), who introduce the term "extreme-belief aversion"). Indeed, a Coarse Bayesian employing the representation in Figure 5b would effectively under-react to signals that strongly support a particular state.

3. Logical Fallacies. Coarse Bayesian updating can give rise to—or make agents susceptible to—various logical or rhetorical fallacies that are common in real life. First, agents who consider only a small set of competing theories may perceive false dilemmas: they overlook plausible alternative explanations for the data, typically narrowing the options down to two alternatives. Political polarization, for example, is consistent with such behavior: in highly polarized environments, people may be quick to sort others into a small number of categories (eg, party affiliation) based on limited, noisy information about their views.

More generally, coarseness can lead to faulty generalizations: the partition represents the agent's tendency to falsely extrapolate, resulting in posteriors that need not be strongly supported by the evidence. This tendency to "jump to conclusions" can make agents susceptible to slippery-slope arguments.

Finally, Coarse Bayesian agents may be susceptible to "straw man" arguments: by providing evidence that refutes some particular theory, the agent may believe other theories are refuted as well. For example, suppose an agent's prior places high probability on state 1, and that a signal strongly refutes state 2 but not state 1. If δ_2 and μ^e are in the same cell, then the agent may end up rejecting his prior even though the signal does not strongly conflict with state 1. By refuting the "straw man" theory (δ_2), the signal has caused the agent to abandon a theory that did not conflict with the evidence.

3.2 Paradigm Shifts and Maximum-Likelihood Updating

For a Coarse Bayesian, the act of updating beliefs may resemble a "paradigm shift." If, for example, different feasible beliefs μ^P represent competing models or theories, then the act of revising beliefs may involve a dramatic shift in how the agent understands the world. In this interpretation, the partition $\mathcal{P}$ captures the agent's tolerance for conflicting evidence. Data that approximately support a theory μ^P do not cause paradigm shifts—such data yield Bayesian posteriors in cell P and, as such, are within the agent's subjective "margin of error." But data that strongly conflict with μ^P cause the agent to abandon theory μ^P in favor of some other $\mu^{P'}$.

Ortoleva (2012) proposes a Hypothesis-Testing (HT) model of belief updating. Under HT, an agent holds a subjective prior and applies Bayes' rule if the observed evidence is sufficiently likely under the prior (that is, above some threshold $\varepsilon > 0$, a subjective characteristic of the individual). For evidence that is "unexpected" (likelihood less than ε), the agent updates beliefs by applying a maximum-likelihood criterion to a *second-order prior*. In particular, the agent applies Bayes' rule to the second-order prior, then adopts as posterior a belief that has maximal probability under the new second-order distribution.

Coarse Bayesian updating accommodates similar behavior. For example, consider Figure 6. In this representation, the agent's prior places relatively high probability on a particular state. For signals that do not strongly conflict with that state, the agent applies Bayes' rule (each point in the shaded region is a cell of the partition). For other signals, however, the agent reacts in a non-Bayesian way. Thus, unexpected news yields non-Bayesian reactions, or paradigm shifts.[7]

As described above, Coarse Bayesian updating accommodates a fairly general notion of paradigm shifts where agents employ subjective tolerances for switching among competing theories. Under this interpretation, it is natural to wonder if the updating procedure can be reformulated in terms of second-order beliefs. Can a Coarse Bayesian agent be re-expressed as one who applies Bayes' rule to a second-order prior? To answer this question, I begin by extending the maximum-likelihood

[7] As Weinstein (2017) explains, the HT model allows essentially any updating to occur for unexpected news (ie, likelihood less than ε). As we shall see, extending maximum-likelihood updating procedures to the domain of noisy signals does rule out some updating behavior.

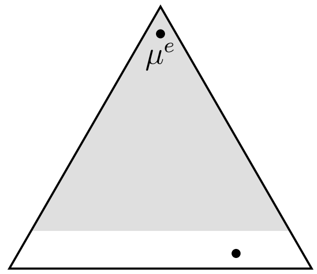

Figure 6: An agent who is Bayesian provided the signal is not too "surprising."

component of the HT model to the domain of noisy signals.

Definition 3. A Homogeneous, Convex updating rule μ has a **Maximum-Likelihood (ML) Representation** if there exists a probability distribution Γ over Δ (with density γ) such that

$$\mu^s \in \underset{\hat{\mu}\in\Delta}{\operatorname{argmax}}\, \gamma(\hat{\mu})\hat{\mu}\cdot s$$

for all $s \in S$. The function $L : \Delta \times S \to \mathbb{R}$ given by $L(\hat{\mu}|s) = \gamma(\hat{\mu})\hat{\mu}\cdot s$ is the **likelihood function**.

In a Maximum-Likelihood Representation, the agent has a second-order prior Γ that he updates (via Bayes' rule) upon arrival of signal s. Then, he selects a belief that has maximal probability under the new second-order distribution. This procedure selects among beliefs $\hat{\mu}$ that maximize the likelihood function at s.[8] Intuitively, ML updating captures the behavior of an agent who assigns prior degrees of confidence to competing theories, updates these values in a Bayesian fashion, and then selects the most-likely theory given available information.[9]

Proposition 1.

(i) Not every Maximum-Likelihood rule can be expressed as a Coarse Bayesian rule.

(ii) Not every Coarse Bayesian rule can be expressed as a Maximum-Likelihood rule.

[8] Notice that L is homogeneous (of degree 0) and convex in s. The restriction to homogeneous, convex updating rules, therefore, only takes effect when there are ties—multiple candidate beliefs that maximize L.

[9] There are other ways of reducing a second-order belief to a first-order belief. For example, one might use the second-order distribution to compute an average belief. However, such a procedure is continuous in s while Coarse Bayesian updating, in general, exhibits discontinuities in s.

(iii) If $N = 2$, then every Coarse Bayesian rule is a Maximum-Likelihood rule.

(iv) Bayesian updating is a special case of both Coarse Bayesian and Maximum-Likelihood updating. To express Bayesian updating as a Maximum-Likelihood rule, take

$$\gamma(\hat{\mu}) \propto \left\| \frac{\hat{\mu}}{\sqrt{\mu^e}} \right\|^{-1} \tag{2}$$

where $\sqrt{\mu^e} := (\sqrt{\mu^e_\omega})_{\omega \in \Omega}$.

Proposition 1 establishes that neither updating procedure subsumes the other—there exist updating rules that have Coarse Bayesian Representations but not ML Representations, and there exist updating rules that have ML Representations but not Coarse Bayesian Representations. These claims are demonstrated by Examples 1 and 2 below. Part (iii) establishes an important special case: if there are only two states, then every Coarse Bayesian rule can be expressed as a ML rule. Part (iv) asserts that standard Bayesian updating is a special case of both models. It is easy to see how Coarse Bayesian updating accommodates Bayesian updating—simply make each point of Δ a singleton cell. For proof that formula (2) generates Bayesian updating in the associated ML Representation, see the appendix.

Example 1. Not every ML rule can be expressed as a Coarse Bayesian rule. Suppose $|\Omega| = 2$ and consider the distribution γ such that $\gamma(\mu^1) = 3/4$ and $\gamma(\mu^2) = 1/4$, where $\mu^1 = (1/3, 2/3)$ and $\mu^2 = (3/4, 1/4)$. Observe that $L(\mu^1|e) = \gamma(\mu^1)\mu^1 \cdot e = \gamma(\mu^1) > \gamma(\mu^2) = \gamma(\mu^2)\mu^2 \cdot e = L(\mu^2|e)$; thus, $\mu^e = \mu^1$. It is easy to verify that $B(\mu^e|s) = \mu^2$ if and only if $s_1/s_2 = 6$. Therefore, to be consistent with a Coarse Bayesian updating rule, we must have $L(\mu^2|s) \geq L(\mu^1|s)$ whenever $s_1/s_2 = 6$. Take $s = (1, 1/6)$. Then $L(\mu^2|s) = 19/96 < 19/72 = L(\mu^1|s)$, so that the ML rule selects μ^1 at s. This means the rule is not Confirmatory, and therefore is inconsistent with Coarse Bayesian updating.

Example 2. Not every Coarse Bayesian rule can be expressed as a ML rule. Suppose $|\Omega| = 3$ and consider a Coarse Bayesian representation where $\mathcal{P}$ has two cells, P and P', with $\mu^P = \mu^e$ and $\mu^{P'} = \mu' \neq \mu^e$. The boundary between P and P' corresponds to a hyperplane, H, in S. We will choose H (hence, $\mathcal{P}$) in such a way that no distribution γ on Δ (with support $\{\mu^e, \mu'\}$) can generate the same updating behavior as $\langle \mathcal{P}, (\mu^P)_{P\in\mathcal{P}} \rangle$ under the ML procedure.

Observe that if γ generates the same updating behavior, then $L(\mu^e|s) = L(\mu'|s)$ for all $s \in H$. In particular, $[\gamma(\mu^e)\mu^e - \gamma(\mu')\mu'] \cdot s = 0$ for all $s \in H$. Thus, the line $\{\lambda[\mu^e + \mu'] - \mu' : \lambda \geq 0\}$ is orthogonal to the hyperplane H. Since $\mu^e \neq \mu'$, we may assume H strictly separates μ^e and μ'. Thus, we may perturb the hyperplane H to ensure it is not orthogonal to the line. Consequently, the resulting Coarse Bayesian Representation cannot be represented by any ML rule.

As demonstrated in Example 1 above, ML updating rules are not guaranteed to satisfy the Confirmation property and, therefore, may be incompatible with Coarse Bayesian updating. I show in appendix B that the categorical-thinking model of Mullainathan (2002) also violates Confirmation in some cases, and for a similar reason. Rather than employing a second-order prior γ to compute likelihoods and select posteriors, Mullainathan's model uses a particular formula to calculate "base rates" for candidate beliefs. Thus, the categorical-thinking model is similar in spirit to a ML procedure, and the particular functional form employed can produce violations of Confirmation.

3.3 Dynamics

This section examines some basic dynamic properties of Coarse Bayesian updating. Suppose an agent observes a sequence of signals $\vec{s} = (s^1, \ldots, s^n)$, where s^t is the signal generated in period t. How do properties of $\vec{s}$ affect the agent's final belief?

For standard Bayesians, the answer is quite simple: signals can be pooled and ordered in any fashion without impacting final beliefs. For example, consider a sequence $\vec{s} = (s^1, s^2, s^3)$. The terminal Bayesian belief is $B(\mu^e|s^1s^2s^3)$ regardless of whether the signals are arranged in a different order (eg. (s^2, s^1, s^3)), pooled differently (eg. (s^1, s^2s^3)), or both.[10]

For Coarse Bayesians, the answer is more nuanced. For example, whether an agent employs a Coarse Bayesian Representation $\langle \mathcal{P}, \mu^{\mathcal{P}} \rangle$ or its associated Signal Distortion Representation $d(\cdot)$ affects the types of history-dependence satisfied by the updating procedure. In addition, the behavior of an agent who incorporates the full history of signal realizations into current beliefs differs from one who performs signal-by-signal

[10] See Cripps (2018) for a general analysis of updating rules that are invariant to how an agent partitions histories of signals.

updating. For simplicity, I focus on "memoryless" agents who perform signal-by-signal updating.

Some additional terminology and notation is needed to proceed. A signal s is **interior** if $s_\omega > 0$ for all $\omega \in \Omega$. A **dynamic updating rule** associates a belief $\mu^{(s^1,\dots,s^n)}$ to every finite **history** $(s^1,\dots,s^n)$ of interior signals. Interpreting a signal s as a history of length 1, a dynamic updating rule clearly gives rise to an updating rule with prior μ^e (full support).

Definition 4. A dynamic updating rule μ is:

(i) A **Dynamic Coarse Bayesian** updating rule if there is a Coarse Bayesian Representation $\langle \mathcal{P}, \mu^{\mathcal{P}} \rangle$ for histories of length 1 such that, for every history $(s^1,\dots,s^n)$ of length $n \geq 2$, $\mu^{(s^1,\dots,s^n)} = \mu^P$ where $B(\mu^{(s^1,\dots,s^{n-1})}|s^n) \in P \in \mathcal{P}$.

(ii) A **Dynamic Signal Distortion** rule if there is a (Homogeneous, Convex, Idempotent) Signal Distortion Representation $d(\cdot)$ for histories of length 1 such that, for every history $(s^1,\dots,s^n)$ of length $n \geq 2$, $\mu^{(s^1,\dots,s^n)} = B(\mu^{(s^1,\dots,s^{n-1})}|d(s^n))$.

(iii) A **Dynamic Bayesian** updating rule if, for all histories $(s^1,\dots,s^n)$ of length $n \geq 1$, $\mu^{(s^1,\dots,s^n)} = B(\mu^e|s^1 s^2 \dots s^n)$.

A Dynamic Coarse Bayesian updating rule employs a fixed Coarse Bayesian Representation to perform signal-by-signal updating. Starting with prior μ^e, the agent applies $\langle \mathcal{P}, \mu^{\mathcal{P}} \rangle$ to yield some posterior μ^{s^1} after observing s^1. Then, treating μ^{s^1} as the prior, the agent applies the same representation $\langle \mathcal{P}, \mu^{\mathcal{P}} \rangle$ to reach posterior $\mu^{(s^1,s^2)}$ after observing s^2, and so on. Thus, the agent applies the same Coarse Bayesian Representation while processing signals one at a time, treating the current belief as the prior. A Dynamic Signal Distortion rule follows a similar procedure, substituting $d(\cdot)$ for $\langle \mathcal{P}, \mu^{\mathcal{P}} \rangle$. Finally, a Dynamic Bayesian updating rule applies Bayes' rule to history $(s^1,\dots,s^n)$ by updating prior μ^e with the pooled signal $s^1 \dots s^n$. As shown in Proposition 2 below, this is equivalent to applying Bayes' rule in a signal-by-signal fashion.

Definition 5. A dynamic updating rule μ is:

(i) **Invariant to signal ordering** if $\mu^{\vec{s}} = \mu^{\pi(\vec{s})}$ for all histories $\vec{s}$ and permutations $\pi(\vec{s})$ of $\vec{s}$.

(ii) **Invariant to signal pooling** if, for all histories $\vec{s} = (s^1, \ldots, s^n)$ of length $n \geq 2$ and all $k < n$, $\mu^{\vec{s}} = \mu^{(s^1, \ldots, s^{k-1}, s^k s^{k+1}, s^{k+2}, \ldots, s^n)}$.

Definition 5 formalizes two different notions of history independence. If a dynamic updating rule is invariant to signal ordering, then any history $\vec{s}$ can be reordered without affecting the final belief.[11] Invariance to signal pooling, by contrast, requires that any signal in a history can be pooled with its successor without affecting the final belief. Clearly, invariance to signal pooling implies invariance to signal ordering. However, as demonstrated by the next proposition, the converse implication does not hold.

Proposition 2. *Let μ^e have full support. Then:*

(i) The Dynamic Bayesian updating rule is invariant to signal ordering and pooling.

(ii) Dynamic Signal Distortion updating rules are invariant to signal ordering but not necessarily to signal pooling.

(iii) Dynamic Coarse Bayesian updating rules need not be invariant to signal ordering nor to signal pooling.

Proof. First, consider a Dynamic Signal Distortion rule μ. Observe that for every signal r, $\mu^r = B(\mu^e | d(r)) \approx d(r)\mu^e$. It follows immediately that $\mu^{(s,t)} \approx d(t)d(s)\mu^e \approx \mu^{(t,s)}$, so that μ is invariant to signal orderings. However, by the same logic, $\mu^{st} \approx d(st)\mu^e$; as shown in Example 3 below, this allows for the possibility that some Dynamic Signal Distortion rules need not be invariant to signal pooling. For (i), note that the Dynamic Bayesian updating rule is a special case of a Dynamic Signal Distortion rule with distortion function d_B given by $d_B(r) = r$ for all r. Thus, $d_B(st) = d_B(s)d_B(t)$, so that the Dynamic Bayesian updating rule is invariant to both signal orderings and signal merging. For (iii), see Example 4 and Proposition 3 below. □

Proposition 2 summarizes the ways in which Dynamic Coarse Bayesian and Signal Distortion rules may exhibit history (in)dependence. Notably, the dynamic setting

[11] Rabin and Schrag (1999) analyze a model of history-dependent updating where, at each time period, information is distorted to support the agent's current belief. Such a procedure would not be invariant to signal ordering.

introduces a wedge between Coarse Bayesian and Signal Distortion rules: Signal Distortion rules are always invariant to signal ordering, but Coarse Bayesian rules typically are not. Both rules, however, are typically not invariant to signal pooling.

Example 3. Dynamic Signal Distortion rules need not be invariant to signal pooling. For example, consider a model with two states and distortion function

$$d(s) = \begin{cases} (1/5, 4/5) & \text{if } \frac{s_2}{s_1} \geq 2 \\ e & \text{else} \end{cases}.$$

It is easy to verify that d is Homogeneous, Convex, and Idempotent. Let $s = (1/5, 4/5)$ and $t = (3/4, 1/4)$. Then $st = (3/20, 4/20)$, $d(st) = e$, $d(s) = (1/5, 4/5)$, and $d(t) = e$; thus, $d(s)d(t) = (1/5, 4/5) \neq e = d(st)$, so that $\mu^{(s,t)} \neq \mu^{st}$.

Example 4. Some (non-Bayesian) Dynamic Coarse Bayesian updating rules are invariant to signal ordering. For example, if $\mathcal{P}$ consists of a single cell (namely, $\{\Delta\}$), then $\mu^s = \mu^e$ for all $s \in S$. Less trivially, suppose $N = 2$ (so that Δ may be represented by the interval $[0,1]$) and consider $\mathcal{P} = \{[0,1), \{1\}\}$ with $\mu^{[0,1)} = 1/2$ and $\mu^{\{1\}} = 1$. It is straightforward to verify that this representation induces invariance to signal ordering.

Example 4 shows that there are cases where Dynamic Coarse Bayesian updating rules are invariant to signal ordering. However, there are many scenarios where they are not. The next proposition highlights a simple class of Coarse Bayesian Representations that result in such path-dependence.

A Coarse Bayesian Representation $\langle \mathcal{P}, \mu^{\mathcal{P}} \rangle$ is **regular** if, for every $P \in \mathcal{P}$, there is an open neighborhood $O \subseteq \Delta^0$ such that $\mu^P \in O \subseteq P$.

Proposition 3. *If $\langle \mathcal{P}, \mu^{\mathcal{P}} \rangle$ is regular and $\mathcal{P}$ has at least two cells, then the associated Dynamic Coarse Bayesian updating rule is not invariant to signal ordering.*

Proof. Let $P \in \mathcal{P}$ be a cell containing δ_1, where $\delta_1 \in \Delta$ assigns probability one to state ω_1. Let r be any signal such that $r\mu^e \approx \mu^P$. Choose any $P' \in \mathcal{P}$ such that $\mu^P \neq \mu^{P'}$. By regularity, μ^P and $\mu^{P'}$ have full support. Thus, there is a signal s such that $s\mu^P \approx \mu^{P'}$. It follows that $\mu^r = \mu^P$ and $\mu^{(r,s)} = \mu^{P'}$. Let t be a signal such that $t_1 = 1$ and $t_\omega = \varepsilon$ for all $\omega \neq 1$. By regularity, and the fact that $\delta_1 \in P$, there is an ε (sufficiently small) such that both $B(\mu^P|t) \in P$ and $B(\mu^{P'}|t) \in P$. Thus, $\mu^{(r,s,t)} = \mu^P$. However, $\mu^{(r,t,s)} = \mu^{P'}$ because $\mu^{(r,t)} = \mu^P$ and $B(\mu^P|s) = \mu^{P'}$. Thus, the updating rule is not invariant to signal ordering. □

4 Application: The Value of Information

Assessing the value of information is a fundamental part of decision making in many economic models. The classic characterization of Blackwell (1951) develops an informativeness ordering where one information structure is more informative than another if and only if it grants a Bayesian agent higher expected utility in all decision problems. In this section, I study the value of information for Coarse Bayesians, including its relationship to the Bayesian value of information, the Blackwell ordering, and a notion of cognitive sophistication.

Throughout this section, suppose μ is an updating rule with Coarse Bayesian Representation $\langle \mathcal{P}, \mu^{\mathcal{P}} \rangle$. Let $\mathcal{A}$ denote the set of all nonempty, compact subsets of $\mathbb{R}^\Omega$. Each $A \in \mathcal{A}$ is a **menu**, and elements $x = (x_\omega)_{\omega \in \Omega} \in A$ represent feasible **actions** the agent may take. An agent who chooses action $x \in A$ attains payoff x_ω in state ω. For each $A \in \mathcal{A}$ and $s \in S$, let $c^s(A) := \operatorname{argmax}_{x \in A} x \cdot \mu^s$; these are the actions in A that maximize expected utility under beliefs μ^s.

Definition 6. Let $A \in \mathcal{A}$.

(i) The **value of information** at A is given by the function $V^A : \mathcal{E} \to \mathbb{R}$ where

$$V^A(\sigma) := \max \sum_\omega \mu^e_\omega \sum_{s \in \sigma} s_\omega x^s_\omega \text{ subject to } x^s \in c^s(A) \tag{3}$$

for all $\sigma \in \mathcal{E}$.

(ii) The **Bayesian value of information** at A is given by the function $\overline{V}^A : \mathcal{E} \to \mathbb{R}$ where

$$\overline{V}^A(\sigma) := \max \sum_\omega \mu^e_\omega \sum_{s \in \sigma} s_\omega x^s_\omega \text{ subject to } x^s \in \operatorname*{argmax}_{x \in A} x \cdot \frac{s\mu^e}{s \cdot \mu^e} \tag{4}$$

for all $\sigma \in \mathcal{E}$.

Equation (3) expresses ex-ante expected utility for a Coarse Bayesian agent. Faced with a menu A and experiment σ, the agent calculates expected utility by applying weight μ^e_ω to the average payoff in state ω given that signals—and subsequent choices—are generated by σ. Consistent with the Cognizance assumption, the agent correctly forecasts his own signal-contingent beliefs and, hence, signal-contingent choices. Equation (4) expresses a similar formula for an agent with the same prior μ^e and who applies Bayes' rule: signal-contingent choices maximize expected utility at beliefs $B(\mu^e|s)$ instead of beliefs μ^s.

It will be convenient to express V^A in a slightly different form. For any $\hat{\mu} \in \Delta$ and $A \in \mathcal{A}$, let

$$c^{\hat{\mu}}(A) := \operatorname*{argmax}_{x \in A} x \cdot \mu^P \text{ subject to } \hat{\mu} \in P \tag{5}$$

and

$$v^A(\hat{\mu}) := \max_{x \in c^{\hat{\mu}}(A)} x \cdot \hat{\mu}. \tag{6}$$

Intuitively, $c^{\hat{\mu}}(A)$ consists of the actions in A that maximize expected utility for the Coarse Bayesian if the *Bayesian* posterior is $\hat{\mu}$ because the agent replaces $\hat{\mu}$ with μ^P if $\hat{\mu} \in P$. Similarly, $v^A(\hat{\mu})$ represents expected utility conditional on the Bayesian posterior being $\hat{\mu}$. These mappings are well-defined because $\mathcal{P}$ partitions Δ and each cell $P \in \mathcal{P}$ has a unique representative μ^P. For a standard Bayesian, analogous mappings are given by

$$\overline{c}^{\hat{\mu}}(A) := \operatorname*{argmax}_{x \in A} x \cdot \hat{\mu} \qquad \text{and} \qquad \overline{v}^A(\hat{\mu}) := \max_{x \in \overline{c}^{\hat{\mu}}(A)} x \cdot \hat{\mu}.$$

If $\sigma \in \mathcal{E}$ and $\hat{\mu} \in \Delta$, let $\tau^\sigma(\hat{\mu}) := \sum_{s \in \sigma : B(\mu^e|s) = \hat{\mu}} s \cdot \mu^e$; this is the total probability of generating Bayesian posterior $\hat{\mu}$ under information σ and prior μ^e. That is, given μ^e, σ generates a distribution of Bayesian posteriors where $\tau^\sigma(\hat{\mu})$ is the probability of posterior $\hat{\mu}$.

Proposition 4. *For all $A \in \mathcal{A}$ and $\sigma \in \mathcal{E}$, $V^A(\sigma) = \sum_{\hat{\mu} \in \Delta} \tau^\sigma(\hat{\mu}) v^A(\hat{\mu})$.*

Proof. First, observe that (3) can be rewritten as

$$V^A(\sigma) = \max \sum_{s \in \sigma} (s\mu^e) \cdot x^s \;\text{ subject to }\; x^s \in c^s(A).$$

It follows that

$$\begin{aligned} V^A(\sigma) &= \max \sum_{s\in\sigma} (s\cdot\mu^e)\frac{s\mu^e}{s\cdot\mu^e}\cdot x^s \ \text{ subject to } \ x^s \in c^s(A) \\ &= \max \sum_{s\in\sigma} (s\cdot\mu^e)B(\mu^e|s)\cdot x^s \ \text{ subject to } \ x^s \in c^s(A) \\ &= \sum_{\hat{\mu}\in\Delta} \tau^\sigma(\hat{\mu})v^A(\hat{\mu}), \end{aligned}$$

as desired. □

Proposition 4 establishes that V^A can be written in posterior-separable form. In particular, it is as if the agent associates value $v^A(\hat{\mu})$ to Bayesian posteriors $\hat{\mu}$, so that the distribution of Bayesian posteriors can be used to calculate expected utility. An immediate implication is that the techniques of Kamenica and Gentzkow (2011), for example, can be applied with Coarse Bayesian agents.

Intuitively, Proposition 4 holds because a Coarse Bayesian updating rule is a function of the Bayesian posterior: if one knows which Bayesian posterior has realized, then one knows which belief the Coarse Bayesian adopts. This is the fundamental assumption of de Clippel and Zhang (2019), who study persuasion with non-Bayesian agents. A similar argument also appears in Galperti (2019).

Proposition 5. *Coarse Bayesians have a lower value of information than Bayesians: for every $A \in \mathcal{A}$ and $\sigma \in \mathcal{E}$, $V^A(\sigma) \leq \overline{V}^A(\sigma)$.*

Proof. Observe that $\overline{V}^A$ can be expressed in posterior-separable form by replacing v^A with $\overline{v}^A$. Moreover, $\overline{v}^A(\hat{\mu}) \geq v^A(\hat{\mu})$ for all $A \in \mathcal{A}$ and $\hat{\mu} \in \Delta$ because $c^{\hat{\mu}}(A) \subseteq A$. Thus, $\overline{V}^A(\sigma) = \sum_{\hat{\mu}\in\Delta} \tau^\sigma(\hat{\mu})\overline{v}^A(\hat{\mu}) \geq \sum_{\hat{\mu}\in\Delta} \tau^\sigma(\hat{\mu})v^A(\hat{\mu}) = V^A(\sigma)$. □

Proposition 5 states that, compared to Bayesians, Coarse Bayesians experience lower ex-ante expected utility for all combinations of menu and experiment. Intuitively, this is driven by the fact that Coarse Bayesian agents tend to make suboptimal decisions conditional on available information. As illustrated in Figure 7, this implies $v^A(\hat{\mu}) \leq \overline{v}^A(\hat{\mu})$ for all $\hat{\mu}$, making $V^A(\sigma) \leq \overline{V}^A(\sigma)$ for all σ.

The next set of results examines whether and when Coarse Bayesians benefit from improvements to information. For experiments σ, σ', the relation $\sigma \sqsupseteq \sigma'$ indicates that σ is more informative than σ' in the sense of Blackwell (1951). This is a partial order

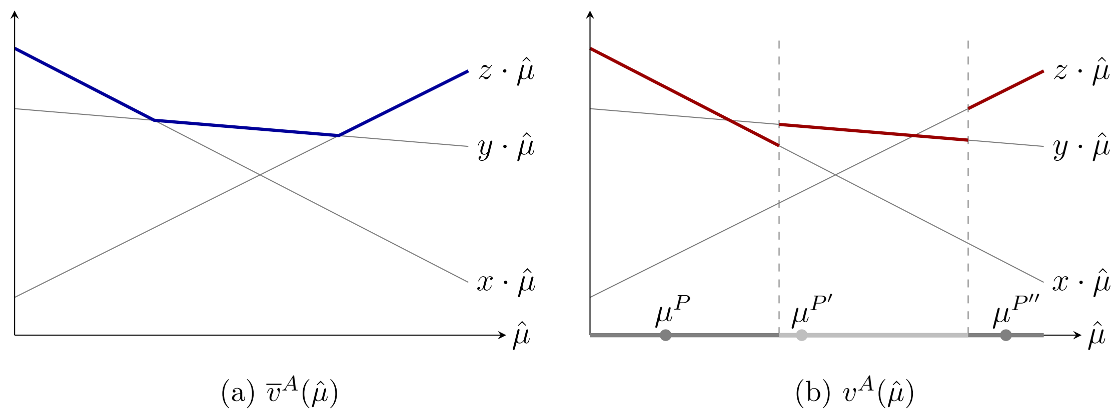

Figure 7: Bayesian vs. Coarse Bayesian value of information for $A = \{x, y, z\}$.

on $\mathcal{E}$. An experiment σ' is a **garbling** of σ if there is a stochastic matrix M such that $\sigma' = \sigma M$. As is well-known, $\sigma \sqsupseteq \sigma'$ if and only if σ' is a garbling of σ.

The function V^A **satisfies the Blackwell ordering** if $\sigma \sqsupseteq \sigma'$ implies $V^A(\sigma) \geq V^A(\sigma')$; if there exists $\sigma \sqsupseteq \sigma'$ such that $V^A(\sigma) < V^A(\sigma')$, then V^A **violates the Blackwell ordering**. An important part of Blackwell's characterization is that a Bayesian agent's value of information (that is, $\overline{V}^A$) satisfies the Blackwell ordering in all menus A—in fact, $\sigma \sqsupseteq \sigma'$ if and only if $\overline{V}^A(\sigma) \geq \overline{V}^A(\sigma')$ for all $A \in \mathcal{A}$. For Coarse Bayesians, this is not the case.

For every menu A and signal s, let $b^s(A) \subseteq A$ denote the Bayesian-optimal actions in A conditional on s. That is, $b^s(A) := \{x \in A : x \cdot \frac{s\mu^e}{s \cdot \mu^e} \geq y \cdot \frac{s\mu^e}{s \cdot \mu^e} \ \forall y \in X\}$. Let $c(A) = \bigcup_{s \in S} c^s(A)$ and $b(A) = \bigcup_{s \in S} b^s(A)$. Thus, $c(A)$ is the set of actions in A that are chosen by the Coarse Bayesian, while $b(A)$ is the set of actions chosen by the Bayesian. Observe that $c(A) \subseteq b(A)$.

Let Δ^0 denote the interior of Δ. For each $P \in \mathcal{P}$, endow P with the subspace topology of $\mathbb{R}^\Omega$ and let ∂P denote the boundary of P.[12]

Proposition 6. *Let $A \in \mathcal{A}$.*

(i) If $c^s(A) \cap b^s(c(A)) \neq \emptyset$ for all s, then V^A satisfies the Blackwell ordering.

(ii) If there is a cell $P \in \mathcal{P}$ and a point $\mu^* \in \partial P \cap \Delta^0$ such that $\mu^* \neq \mu^P$, then V^A violates the Blackwell ordering.

[12]More precisely, endow $\mathbb{R}^\Omega$ with the standard Euclidean topology. A set $Z \subseteq P$ is open in the subspace topology on P if and only if there is an open set $O \subseteq \mathbb{R}^\Omega$ such that $Z = O \cap P$.

Part (i) of Proposition 6 provides a sufficient and almost-necessary condition for the agent's value of information at A to satisfy the Blackwell information ordering. The condition states that, for every s, Coarse Bayesian choices overlap with Bayesian choices from the *restricted* menu $c(A) \subseteq A$. In other words, the condition states that if one restricts attention to actions in A that are chosen by the Coarse Bayesian for at least one signal realization, then Coarse Bayesian choices must be Bayesian optimal at all signal realizations. Statement (ii) provides a partial converse to (i). The condition in (ii) is satisfied by most Coarse Bayesian representations that violate (i). For example, any representation that is regular (as defined in section 3.3) but violates (i) will satisfy (ii). In fact, every Coarse Bayesian representation depicted in this paper satisfies (ii) for some A and, thus, would generate violations of the Blackwell ordering.

The final results examine how an agent's value of information vary with the following notion of cognitive sophistication:

Definition 7. Let μ and $\dot{\mu}$ be updating rules with full-support priors $\mu^e = \dot{\mu}^e$ and Coarse Bayesian Representations $\langle \mathcal{P}, \mu^{\mathcal{P}} \rangle$ and $\langle \mathcal{Q}, \dot{\mu}^{\mathcal{Q}} \rangle$, respectively. Then $\dot{\mu}$ is **more sophisticated** than μ if $\mu^{\mathcal{P}} \subseteq \dot{\mu}^{\mathcal{Q}}$ and every member of $\mathcal{P}$ is a union of members of $\mathcal{Q}$.

Definition 7 states that a Coarse Bayesian is more sophisticated if he employs both a larger set of feasible beliefs and a finer partition. Such an agent is "more Bayesian" because his updating rule better approximates the standard Bayesian updating rule.

Given a Coarse Bayesian representation $\langle \mathcal{P}, \mu^{\mathcal{P}} \rangle$, a menu A is $\mu^{\mathcal{P}}$**-decisive** if $c^s(A)$ is a singleton for all $s \in S$; that is, no μ^P makes the agent indifferent between two or more options in A.

Proposition 7. *Suppose $\langle \mathcal{P}, \mu^{\mathcal{P}} \rangle$ and $\langle \mathcal{Q}, \dot{\mu}^{\mathcal{Q}} \rangle$ are regular Coarse Bayesian Representations of μ and $\dot{\mu}$, respectively, and that $\mu^e = \dot{\mu}^e$. The following are equivalent:*

(i) $\langle \mathcal{Q}, \dot{\mu}^{\mathcal{Q}} \rangle$ is more sophisticated than $\langle \mathcal{P}, \mu^{\mathcal{P}} \rangle$.

(ii) If $\dot{V}^A(\sigma) = \dot{V}^A(\sigma')$ for all $\dot{\mu}^{\mathcal{Q}}$-decisive menus A, then $V^A(\sigma) = V^A(\sigma')$ for all $\dot{\mu}^{\mathcal{Q}}$-decisive menus A.

This result states that for regular Coarse Bayesians, greater sophistication means higher responsiveness to information: as sophistication increases, fewer pairs σ, σ' yield identical ex-ante expected utility for (almost) all menus A. The proof of Proposition 7 shows that the characterization holds even if one restricts attention to experiments σ, σ' that are Blackwell comparable. Thus, higher sophistication means greater responsiveness to *improvements* to information.

Does a more sophisticated agent always benefit from more information, or enjoy a higher value of information than a less sophisticated agent? In general, the answer to each question is negative. Nonetheless, there are simple conditions under which these intuitive relationships hold. For every $A \in \mathcal{A}$ and $x \in A$, let $B^{A,x} := \{\hat{\mu} \in \Delta : x \cdot \hat{\mu} \geq y \cdot \hat{\mu} \ \forall y \in A\}$; these are posterior beliefs that make x an optimal action in A. Let $\mathcal{B}^A := \{B^{A,x} : x \in A\}$.

Proposition 8. *Suppose μ and $\dot{\mu}$ are updating rules with Coarse Bayesian Representations $\langle \mathcal{P}, \mu^{\mathcal{P}} \rangle$ and $\langle \mathcal{Q}, \dot{\mu}^{\mathcal{Q}} \rangle$ and that $\dot{\mu}$ is more sophisticated than μ. Let $A \in \mathcal{A}$.*

(i) If for every $\dot{\mu}^Q \in \dot{\mu}^{\mathcal{Q}} \backslash \mu^{\mathcal{P}}$ there is a set $B \in \mathcal{B}^A$ such that $Q \subseteq B$, then $\dot{V}^A(\sigma) \geq V^A(\sigma)$ for all $\sigma \in \mathcal{E}$.

(ii) If V^A satisfies the Blackwell ordering and $c(A) = b(A)$, then $\dot{V}^A$ satisfies the Blackwell ordering.

Proof. For (i), the the stated hypotheses imply $\dot{v}^A(\hat{\mu}) \geq v^A(\hat{\mu})$ for all $\hat{\mu}$ because there is a larger set of Bayesian posteriors where $\dot{c}^{\hat{\mu}}(A)$ is Bayesian optimal. For (ii), the hypotheses imply $v^A = \overline{v}^A$. Then, since $\mathcal{Q}$ is finer than $\mathcal{P}$, we will also have $\dot{v}^A = \overline{v}^A$. ☐

Proposition 8 provides simple conditions under which greater sophistication leads to better decision making (hence, higher ex-ante value of information) and greater adherence to the Blackwell ordering. In general, a more sophisticated agent is not guaranteed to have a higher value of information or to adhere to the Blackwell ordering whenever a less-sophisticated agent does. Intuitively, a finer partition improves decisions at some signal realizations, but may worsen them at others. Whether this occurs depends on the menu under consideration, and the condition in part (i) ensures that the finer partition does not introduce any new points of "disagreement" with a Bayesian decision maker, thus increasing the value of information. The condition in (ii) ensures that a finer partition does not introduce any new choices from A. In particular, $c(A) = b(A)$ implies that every Bayesian-optimal action in A is chosen by the less-sophisticated Coarse Bayesian. This way, no new disagreement is generated by introducing a finer partition.

5 Conclusion

In this paper, I have introduced a new model of non-Bayesian updating, *Coarse Bayesian updating*, accommodating many documented types of behavior that are incompatible with standard Bayesian updating. Three testable and normatively appealing assumptions—*Homogeneity*, *Cognizance*, and *Confirmation*—characterize the updating procedure, the parameters of which are a partition of the probability simplex and a representative belief for each cell of the partition. A Coarse Bayesian agent can be interpreted as one who applies subjective criteria to select among competing theories, selectively distorts signals before applying Bayes' rule, or who engages in categorical thinking. I have analyzed how the model relates to existing models and evidence on non-Bayesian updating, and examined its implications in dynamic settings as well as in standard settings of decision under risk.

An advantage of my framework is that it employs standard primitives that frequently appear in applications. The use of noisy signals over an exogenous state space, for example, allows one to directly import Coarse Bayesian updating into familiar settings in economics and game theory. Exploring the implications of Coarse Bayesian updating in such settings may be a fruitful avenue for future research.

A Omitted Proofs

A.1 Proof of Proposition 1

Proof of part (iii). If every cell of $\langle \mathcal{P}, \mu^{\mathcal{P}} \rangle$ is a singleton, then the agent is Bayesian and the ML representation is established independently by the proof of part (iv) below. So, let $P^* \in \mathcal{P}$ be a non-singleton cell. Let I denote the set of all Coarse Bayesian representations $i = \langle \mathcal{Q}(i), \dot{\mu}^{\mathcal{Q}(i)} \rangle$ such that $\mathcal{Q}(i)$ is finite, $\mathcal{P}$ is finer than $\mathcal{Q}(i)$, $\dot{\mu}^{\mathcal{Q}(i)} \subseteq \mu^{\mathcal{P}}$, and $P^* \in \mathcal{Q}(i)$. Define a partial order $\geq_I$ on I by $i \geq_I i'$ if and only if $\mathcal{Q}(i)$ is finer than $\mathcal{Q}(i')$ and $\dot{\mu}^{\mathcal{Q}(i)} \supseteq \dot{\mu}^{\mathcal{Q}(i')}$ (that is, i is more sophisticated than i' in the sense of Definition 7). It is straightforward to verify that $\geq_I$ is a partial order and that for all $i, i' \in I$, there exists $i^* \in I$ such that $i^* \geq_I i$ and $i^* \geq_I i'$. Thus, $(I, \geq_I)$ is a directed set.

For each $\langle \mathcal{Q}, \dot{\mu}^{\mathcal{Q}} \rangle \in I$, define a function $\gamma : \Delta \to [0, \infty)$ as follows. Since $N = 2$, the (finite) set $\dot{\mu}^{\mathcal{Q}}$ can be arranged in decreasing order of state 1: $\dot{\mu}^{\mathcal{Q}} = \{\dot{\mu}^{Q_1}, \ldots, \dot{\mu}^{Q_M}\}$, where $\dot{\mu}_1^{Q_1} > \dot{\mu}_1^{Q_2} > \ldots > \dot{\mu}_1^{Q_M}$. Since $P^* \in \mathcal{Q}$, there exists m^* such that $\dot{\mu}^{Q_{m^*}} = \mu^{P^*}$. For $1 \leq m < M$, let $\dot{\mu}^m$ denote the (unique) belief belonging to $\partial Q_m \cap \partial Q_{m+1}$ and choose a signal s^m such that $B(\mu^e | s^m) = \dot{\mu}^m$. Now choose scalars $\alpha_m > 0$ such that, for all $1 \leq m < M$, $\alpha_m \mu^{Q_m} \cdot s^m = \alpha_{m+1} \mu^{Q_{m+1}} \cdot s^m$; taking $\alpha_{m^*} = 1$ pins down the α_m uniquely. Now define γ by

$$\gamma(\hat{\mu}) = \begin{cases} \alpha_m & \text{if } \hat{\mu} = \mu^{Q_m} \\ 0 & \text{otherwise} \end{cases}.$$

By construction, $\mu^{Q_m} \in \operatorname{argmax}_{\hat{\mu}} \gamma(\hat{\mu}) \hat{\mu} \cdot s$ (that is, μ^{Q_m} maximizes the likelihood function associated with γ) if and only if $B(\mu^e | s) \in Q_m$. Moreover, every point $\gamma(\hat{\mu})\hat{\mu}$, viewed as a point in $\mathbb{R}^2$, is contained in the half-space H bounded above by the line with normal s^* passing through μ^{P^*}, where s^* is any signal such that $B(\mu^e | s^*) = \mu^{P^*}$. Thus, there exists a scalar $\overline{\gamma} > 0$ such that $\gamma(\hat{\mu}) \in [0, \overline{\gamma}]$ for all $\hat{\mu}$. Observe that the bound $\overline{\gamma}$ is independent of i.

Having defined a function $\gamma^i : \Delta \to [0, \overline{\gamma}]$ for every $i \in I$, the family $\{\gamma^i\}_{i \in I}$ forms a net. Each γ^i is an element of the (compact) product set $[0, \overline{\gamma}]^\Delta$, so that $\{\gamma^i\}_{i \in I}$ has a convergent subnet. This means there is a directed set $(J, \geq_J)$ and a function $\iota : J \to I$ such that (a) $j \geq_J j'$ implies $\iota(j) \geq_I \iota(j')$, (b) for every $i \in I$, there exists $j \in J$ such that $\iota(j') \geq_I i$ for all $j' \geq_J j$, and (c) the net $\{\gamma^{\iota(j)}\}_{j \in J}$ converges to some γ^*. Thus, for every $\hat{\mu} \in \Delta$, $\gamma^{\iota(j)}(\hat{\mu})$ converges to a point $\gamma^*(\hat{\mu})$.

Let $P \in \mathcal{P}$. By definition of $(I, \geq_I)$ and properties (a) and (b) of $(J, \geq_J)$, there exists $j^P \in J$ such that $P \in \mathcal{Q}(\iota(j))$ and $\mu^P \in \dot{\mu}^{\mathcal{P}(\iota(j))}$ for all $j \geq_J j^P$. Suppose s satisfies $B(\mu^e | s) \in P$. By construction, μ^P maximizes the likelihood function associated with $\gamma^{\iota(j)}$ at s if $j \geq_J j^P$: for every $\hat{\mu} \in \Delta$, $\gamma^{\iota(j)}(\mu^P)\mu^P \cdot s \geq \gamma^{\iota(j)}(\hat{\mu})\hat{\mu} \cdot s$. Taking the limit of both sides with respect to j gives $\gamma^*(\mu^P)\mu^P \cdot s \geq \gamma^*(\hat{\mu})\hat{\mu} \cdot s$, so that μ^P maximizes the likelihood function associated with γ^* at s. □

Proof of part (iv). Notice that $B(\mu^e | s) = \mu'$ if and only if $s \approx \mu'/\mu^e := (\mu'_\omega / \mu^e_\omega)_{\omega \in \Omega}$. Thus, it will suffice to verify that $L(\cdot | s)$ is maximized at μ' for such signals s. This is done as follows. Let $s \in S$. Then, for any $\hat{\mu} \in \Delta$, we have

$$\begin{aligned}
L(\hat{\mu}|s) &= \gamma(\hat{\mu})\hat{\mu} \cdot s \\
&= \frac{\hat{\mu}}{\|\hat{\mu}/\sqrt{\mu^e}\|} \cdot s \\
&= \frac{\hat{\mu}/\sqrt{\mu^e}}{\|\hat{\mu}/\sqrt{\mu^e}\|} \cdot s\sqrt{\mu^e} \\
&= \left\| \frac{\hat{\mu}/\sqrt{\mu^e}}{\|\hat{\mu}/\sqrt{\mu^e}\|} \right\| \|s\sqrt{\mu^e}\| \cos\theta \\
&= \|s\sqrt{\mu^e}\| \cos\theta
\end{aligned}$$

where θ is the angle (in radians) between $\hat{\mu}/\sqrt{\mu^e}$ and $s\sqrt{\mu^e}$. Thus, $L(\cdot|s)$ is maximized by choosing $\hat{\mu}$ such that $\hat{\mu}/\sqrt{\mu^e} \approx s\sqrt{\mu^e}$ (because then $\theta = 0$), implying $\hat{\mu} \approx s\mu^e \approx \frac{\mu'}{\mu^e}\mu^e = \mu'$. □

A.2 Proofs for Section 4

A.2.1 Proof of Proposition 6

For part (i), observe that if $c^s(A) \cap b^s(c(A)) \neq \emptyset$ for all s, then every Coarse Bayesian choice from A is Bayesian-optimal in the menu $A' = c(A)$. Since Coarse Bayesian choices from A are identical to those from A', it follows that $V^A(\sigma) = V^{A'}(\sigma) = \overline{V}^{A'}(\sigma)$ for all σ. That is, V^A coincides with the Bayesian value of information in some menu, and therefore satisfies the Blackwell ordering.

The remainder of this section proves part (ii). Let $S^0 := \{s \in S : s_\omega < 1 \; \forall \omega \in \Omega\}$. For any signal $s \in S^0$, let $\sigma^s := [s, (1-s_1)e^1, \ldots, (1-s_N)e^N] \in \mathcal{E}$, where $e^\omega \in S$ such that $e^\omega_\omega = 1$ and $e^\omega_{\omega'} = 0$ for $\omega' \neq \omega$.

Lemma 1. *Let $s, t \in S^0$ such that $s_\omega \leq t_\omega$ for all ω. Then σ^t is a garbling of σ^s.*

Proof. Since $s_\omega \leq t_\omega$ for all ω, there exists a vector δ such that $t = s + \delta$ and $\delta_\omega \geq 0$ for all ω. Let $\alpha_\omega := \frac{\delta_\omega}{1-s_\omega}$. Notice that $\alpha_\omega \in [0,1]$ because $s_\omega < 1$ and $s_\omega + \delta_\omega = t_\omega < 1$.

Let

$$M = \begin{bmatrix} 1 & 0 & 0 & \cdots & 0 \\ \alpha_1 & 1-\alpha_1 & 0 & \cdots & 0 \\ \alpha_2 & 0 & 1-\alpha_2 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \alpha_N & 0 & 0 & \cdots & 1-\alpha_N \end{bmatrix}.$$

Clearly, M is a stochastic matrix. Moreover,

$$\begin{aligned} \sigma^s M &= \left[s + \sum_{\omega \in \Omega} \alpha_\omega (1 - s_\omega) e^\omega, (1-\alpha_1)(1-s_1)e^1, \ldots, (1-\alpha_N)(1-s_N)e^N \right] \\ &= \left[s + \sum_{\omega \in \Omega} \delta_\omega e^\omega, (1 - s_1 - \delta_1)e^1, \ldots, (1 - s_N - \delta_N)e^N \right] \\ &= \left[s + \delta, (1-t_1)e^1, \ldots, (1-t_N)e^N \right] \\ &= \sigma^t. \end{aligned}$$

Thus, σ^t is a garbling of σ^s. □

To prove Proposition 6, suppose $\mu^* \in \partial P \cap \Delta^0$ and $\mu^* \neq \mu^P$. There are two cases.

Case 1: $\mu^* \notin P$. There is a hyperplane that strictly separates μ^P and μ^*. Therefore there is a menu $A = \{x, y\}$ such that $\mu^P \cdot (x - y) > 0$ and $\mu' \cdot (y - x) < 0$ for all μ' sufficiently close to μ^*. Therefore, there exists a sequence of signals $s^n \to s^*$ such that such that $B(\mu^e|s^*) = \mu^*$, $B(\mu^e|s^n) \in P$ for all n, and $B(\mu^e|s^n) \cdot (y - x) > 0$ for all n. Since $\mu^* \in \Delta^0$, we have $s^*_\omega > 0$ for all ω. Combined with the fact that $B(\mu^e|\tilde{s}) = B(\mu^e|\lambda\tilde{s})$ for all $\lambda > 0$ such that $\lambda\tilde{s} \in S$, we may assume that $s^n_\omega \leq s^*_\omega < 1$ for all n and ω. Thus, there is a sequence of vectors $\delta^n \to 0$ such that $s^* = s^n + \delta^n$ for all n. Let $\sigma^n = [s^n, \lambda^n_1 e^1, \ldots, \lambda^n_N e^N]$ and $\sigma^* = [s^*, \lambda^*_1 e^1, \ldots, \lambda^*_N e^N]$ as in Lemma 1. Then σ^* is a garbling of σ^n. Since $(s^n \mu^e) \cdot x \to (s^* \mu^e) \cdot x < (s^* \mu^e) \cdot y$, it follows that $\lim_{n\to\infty} V^A(\sigma^n) < V^A(\sigma^*)$. Thus, for large enough n, $V^A(\sigma^n) < V^A(\sigma^*)$, which is a reversal of the Blackwell ordering.

Case 2: $\mu^* \in P$. Choose a cell $P' \neq P$ as follows. If there exists $P'' \neq P$ such that μ^* belongs to the closure of P'', take $P' = P''$. Otherwise, there exists $P'' \neq \mu^*$ such that every point on the line connecting μ^* and $\mu^{P''}$ is the representative of a (singleton) cell. Take P' to be any such P''. There is a hyperplane strictly separating μ^P and $L := \text{co}\{\mu^*, \mu^{P'}\}$. Thus, there is a menu $A = \{x, y\}$ such that $\mu^P \cdot (x - y) > 0$ and $\mu' \cdot (y - x) > 0$ for all μ' sufficiently close to μ^*, including all $\mu' \in L$. In a similar fashion to the previous case, this means we can construct a sequence $\sigma^n \to \sigma^*$ where $B(\mu^e|s^n) \in L$, $B(\mu^e|s^*) = \mu^*$, and σ^n is a garbling of σ^*. The sequence satisfies $\lim_{n\to\infty} V^A(\sigma^n) > V^A(\sigma^*)$, so that V^A does not satisfy the Blackwell ordering.

A.2.2 Proof of Proposition 7

For any $\langle \mathcal{P}, \mu^{\mathcal{P}} \rangle$ and $P \in \mathcal{P}$, let $S^P := \{s \in S : B(\mu^e|s) \in P\}$. For any σ, let $s^{P,\sigma} := \sum_{s \in \sigma \cap S^P} s$. Then σ and σ' are **$\mathcal{P}$-equivalent** if $s^{P,\sigma} = s^{P,\sigma'}$ for all $P \in \mathcal{P}$.

Lemma 2. *Suppose $\langle \mathcal{P}, \mu^{\mathcal{P}} \rangle$ is regular and let $\sigma, \sigma' \in \mathcal{E}$. Then σ and σ' are $\mathcal{P}$-equivalent if and only if $V^A(\sigma) = V^A(\sigma')$ for every $\mu^{\mathcal{P}}$-decisive menu A.*

Proof. Suppose σ and σ' are $\mathcal{P}$-equivalent. Observe that for every $\mu^{\mathcal{P}}$-decisive menu A and experiment $\hat{\sigma}$, $V^A(\hat{\sigma}) = \sum_{P \in \mathcal{P}} (\mu^e s^{P,\hat{\sigma}}) \cdot c^{\mu^P}(A)$ because $c^{\mu^P}(A)$ is a singleton for all $P \in \mathcal{P}$. Thus, $V^A(\sigma) = V^A(\sigma')$ because $s^{P,\sigma} = s^{P,\sigma'}$ for all $P \in \mathcal{P}$.

For the converse, suppose σ and σ' are not $\mathcal{P}$-equivalent. We construct a $\mu^{\mathcal{P}}$-decisive menu A such that $V^A(\sigma) \neq V^A(\sigma')$. For each $P \in \mathcal{P}$, let $\delta^P := s^{P,\sigma} - s^{P,\sigma'}$. Since experiments consist of finitely many signals, there are finitely many (but at least two) cells P such that $\delta^P \neq 0$. Let $\mu^\delta := \{\mu^P : \delta^P \neq 0\}$ and let μ^{P^*} be an extreme point of the convex hull of μ^δ. Since μ^δ is finite, μ^{P^*} can be strictly separated from the convex hull of $\mu^\delta \backslash \{\mu^{P^*}\}$; that is, there exists x such that $x \cdot \mu^{P^*} > 0 > x \cdot \mu^{P'}$ for all $\mu^{P'} \in \mu^\delta \backslash \{\mu^{P^*}\}$. By regularity, we may assume that x is such that the menu $A = \{x, 0\}$ is $\mu^{\mathcal{P}}$-decisive. Then $V^A(\sigma) - V^A(\sigma') = \sum_{P \in \mathcal{P}} (\mu^s \delta^P) \cdot c^{\mu^P}(A) = (\mu^e \delta^{P^*}) \cdot x$ because $c^{\mu^P}(A) = 0$ for all $\mu^P \in \mu^\delta \backslash \{\mu^{P^*}\}$. Thus, $V^A(\sigma) \neq V^A(\sigma')$ provided $(\mu^e \delta^{P^*}) \cdot x \neq 0$. Since the separation is strict (and $\langle \mathcal{P}, \mu^{\mathcal{P}} \rangle$ is regular), we may perturb x if necessary to ensure $(\mu^e \delta^{P^*}) \cdot x \neq 0$. □

Proof that (i) implies (ii). Suppose $\dot{V}^A(\sigma) = \dot{V}^A(\sigma')$ for all $\mu^{\mathcal{Q}}$-decisive A. By Lemma 2, σ and σ' are $\mathcal{Q}$-equivalent. Since $\mathcal{Q}$ is finer than $\mathcal{P}$, it follows that σ and σ' are $\mathcal{P}$-equivalent. Since $\mu^{\mathcal{P}} \subseteq \dot{\mu}^{\mathcal{Q}}$, this implies $V^A(\sigma) = V^A(\sigma')$ for all $\dot{\mu}^{\mathcal{Q}}$-decisive A.

Proof that (ii) implies (i). Let $Q \in \mathcal{Q}$ and suppose $s, t \in S^Q$. Let $\sigma = [s, t, e - s - t]$ (if necessary, scale s and t down by a factor $\lambda > 0$ to make σ well-defined), and let

$\sigma' = [s+t, e-s-t]$. By Convexity, $s+t \in S^Q$ and, thus, σ and σ' are $\mathcal{Q}$-equivalent. By Lemma 2 and the hypothesis of (ii), this implies σ and σ' are $\mu^{\mathcal{P}}$-equivalent. Thus, there exists $P \in \mathcal{P}$ such that $s, t \in S^P$ (otherwise, there are distinct cells $P', P'' \in \mathcal{P}$ such that $s \in P'$ and $t \in P''$; but then σ and σ' are not $\mathcal{P}$-equivalent, as $s+t$ belongs to a single cell). We have shown that any two signals belonging to a common S^Q ($Q \in \mathcal{Q}$) belong to a common S^P ($P \in \mathcal{P}$). Thus, $\mathcal{Q}$ is finer than $\mathcal{P}$.

We now verify that for every $P \in \mathcal{P}$, there exists $Q \in \mathcal{Q}$ such that $\mu^P = \dot{\mu}^Q$. Suppose toward a contradiction that $\mu^P \neq \dot{\mu}^Q$ for all $Q \in \mathcal{Q}$. Let $A = \{x, 0\}$ be a $\dot{\mu}^{\mathcal{Q}}$-decisive menu such that $x \cdot \mu^P = 0$ (that is, the agent is indifferent between x and 0 at beliefs μ^P). By regularity, and the fact that $\mathcal{Q}$ is finer than $\mathcal{P}$, there exists $Q \in \mathcal{Q}$ and $s, t \in S^Q \subseteq S^P$ such that $(s\mu^e) \cdot x > 0 > (t\mu^e) \cdot x$. Let $\sigma = [s, t, e-s-t]$ and $\sigma' = [s+t, e-s-t]$ (again, scale s and t down if necessary). Then σ and σ' are $\mathcal{Q}$-equivalent. By definition of V^A, and by our choice of s and t, we have $V^A(\sigma) = V^A(s) + V^A(t) + V^A(e-s-t) > V^A(s+t) + V^A(e-s-t) = V^A(\sigma')$, contradicting $\mathcal{Q}$-equivalence of σ and σ'.

B Relationship to Mullainathan (2002)

In a working paper, Mullainathan (2002) develops a model of categorical thinking sharing several features of Coarse Bayesian updating. In this appendix, I show that the categorical thinking model (adapted to my framework of states and signals) satisfies Homogeneity and Cognizance but not necessarily Confirmation.

Mullainathan works with a type space T and prior p where $p(t)$ is the prior probability of type $t \in T$. The analogous components in my model are the state space Ω and prior μ^e, where μ^e_ω is the prior probability of state $\omega \in \Omega$. Data d in Mullainathan's model can be expressed by conditional probabilities $p(d|t)$ indicating the probability of observing the data given type t; in my model, data corresponds to a signal realization s, and s_ω (the probability of observing s in state ω) plays the role of $p(d|t)$.

A set C of probability distributions over T constitutes a set of "categories." These are feasible beliefs that the agent can hold in Mullainathan's model. Thus, the set C is analogous to the set $\{\mu^P : P \in \mathcal{P}\}$ in my model. For a category c and data d, $p(d|c)$ is the probability of generating data d in category c; this is analogous to $s \cdot \mu^P$, which is the probability of observing signal s if μ^P is the true probability law.

Finally, Mullainathan defines $p(c) := \int_t p(t)c(t)$ to be the "base rate" of category c.[13] In my model, the analogous rate is $\mu^e \cdot \mu^P$.

Like Coarse Bayesians, agents in Mullainathan's model partition the probability simplex and assign posterior beliefs as a function of the cell containing the Bayesian posterior. Any set of C of categories (feasible posteriors) is permitted; however, the partition is pinned down by C and an optimality criterion resembling that of Maximum-Likelihood rules in section 3.2. In particular, let $c^*(d) \in C$ denote the agent's posterior after observing data d. Mullainathan requires that

$$c^*(d) \in \operatorname*{argmax}_{c \in C} p(d|c)p(c). \tag{7}$$

In my framework, the analogous condition is

$$\mu^s \in \operatorname*{argmax}_{\hat{\mu} \in \hat{C}} (s \cdot \hat{\mu})(\mu^e \cdot \hat{\mu}), \tag{8}$$

where $\hat{C} \subseteq \Delta$ is some set of feasible posteriors. This is very similar to maximization of the likelihood function specified in section 3.2; the main difference is that my likelihood functions use a second-order belief γ instead of the base rate $p(c)$ proposed by Mullainathan.

Thus, Mullainathan's model works by specifying a set C of categories (feasible posteriors) from which the criterion (7) selects posteriors after observing data d. Because of the functional forms employed, it is as if there is a partition of the probability simplex such that the agent's selected posterior only depends on the Bayesian posterior.

Unlike Coarse Bayesians, categorical thinkers need not satisfy Confirmation because condition (7) does not guarantee that beliefs $c^*(d)$ belong to the cell containing the Bayesian posterior associated with data d.[14] Below, I prove these claims in my framework (in particular, employing condition (8)).

First, let $\hat{C}$ be a nonempty set of feasible posteriors. Suppose that some $\mu^* \in \hat{C}$ is a solution to the maximization problem in (8) for both s and t. That is, μ^* solves

[13] I have modified the notation slightly; Mullainathan writes $q_c(\cdot)$ instead of $c(\cdot)$ to indicate the probability distribution over T associated with category $c \in C$.

[14] Note that the partitions in Mullainathan's model will typically have convex cells. Convexity only fails if the maximization problem in (7) has more than one solution and the agent's tie-breaking criterion is not convex.

both

$$\max_{\hat{\mu}\in\hat{C}}(s\cdot\hat{\mu})(\mu^e\cdot\hat{\mu}) \qquad \text{and} \qquad \max_{\hat{\mu}\in\hat{C}}(t\cdot\hat{\mu})(\mu^e\cdot\hat{\mu}).$$

Then, if $\alpha, \beta \geq 0$, it follows at μ^* solves

$$\max_{\hat{\mu}\in\hat{C}}((\alpha s+\beta t)\cdot\hat{\mu})(\mu^e\cdot\hat{\mu}).$$

It follows that the map $s \mapsto \text{argmax}_{\hat{\mu}\in\hat{C}}(s\cdot\hat{\mu})(\mu^e\cdot\hat{\mu})$ is measurable with respect to a partition of S into convex cones. As demonstrated in the proof of Theorem 1, such convex cones can be associated with convex subsets of Δ by mapping signals s to Bayesian posteriors $B(\mu^e|s)$.

Thus, any updating rule satisfying (8) satisfies Homogeneity and Cognizance if one restricts attention to signals that yield unique solutions to the optimization problem. For signals that involve ties, Homogeneity and/or Cognizance may be violated if the agent's tie-breaking selection is not Homogeneous or Convex.

A more substantive difference between Mullainathan's model and Coarse Bayesian updating is that condition (8) does not guarantee that the updating rule satisfies Confirmation. To see this, suppose $|\Omega| = 2$ and let $\mu^e = (\frac{1}{3}, \frac{2}{3})$. Suppose that $\hat{\mu}, \hat{\mu}' \in \hat{C}$ where $\hat{\mu} = (\frac{1}{4}, \frac{3}{4})$ and $\hat{\mu}' = (\frac{1}{5}, \frac{4}{5})$. Let $s = (\frac{3}{8}, \frac{9}{16})$. It follows that $B(\mu^e|s) = \hat{\mu}$; so, Confirmation requires $\hat{\mu}$ to solve

$$\max_{\tilde{\mu}\in\hat{C}}(s\cdot\tilde{\mu})(\mu^e\cdot\tilde{\mu}).$$

However,

$$(s\cdot\hat{\mu})(\mu^e\cdot\hat{\mu}) = \frac{77}{256} < \frac{63}{200} = (s\cdot\hat{\mu}')(\mu^e\cdot\hat{\mu}').$$

Thus, $\hat{\mu}$ is not selected at s, violating Confirmation.

References

Ambuehl, S. and S. Li (2018). Belief updating and the demand for information. *Games and Economic Behavior 109*, 21–39.

Aydogan, I., A. Baillon, E. Kemel, and C. Li (2017). Signal perception and belief updating. *Working paper*.

Balzer, B. and B. Young (2020). A theory of intuition and contemplation. *Working paper*.

Barberis, N., A. Shleifer, and R. Vishny (1998). A model of investor sentiment. *Journal of financial economics 49*(3), 307–343.

Benjamin, D. J. (2019). Errors in probabilistic reasoning and judgment biases. In *Handbook of Behavioral Economics: Applications and Foundations 1*, Volume 2, pp. 69–186. Elsevier.

Benjamin, D. J., M. Rabin, and C. Raymond (2016). A model of nonbelief in the law of large numbers. *Journal of the European Economic Association 14*(2), 515–544.

Blackwell, D. (1951). Comparison of experiments. In *Proceedings of the second Berkeley symposium on mathematical statistics and probability*, Volume 1, pp. 93–102.

Brunnermeier, M. K. and J. A. Parker (2005). Optimal expectations. *American Economic Review 95*(4), 1092–1118.

Camerer, C. (1995). Individual decision making. *Handbook of experimental economics*.

Cripps, M. W. (2018). Divisible updating. *Working Paper*.

de Clippel, G. and X. Zhang (2019). Non-bayesian persuasion. *Working Paper*.

Ducharme, W. (1970). Response bias explanation of conservative human inference. *Journal of Experimental Psychology 85*, 66–74.

Edwards, W. (1968). Conservatism in human information processing. *Formal representation of human judgment*, 17–52.

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*(2), 114–38.

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

Epstein, L. G., J. Noor, and A. Sandroni (2008). Non-Bayesian updating: A theoretical framework. *Theoretical Economics 3*(2), 193–229.

Fryer, R. G., P. Harms, and M. O. Jackson (2019). Updating beliefs when evidence is open to interpretation: Implications for bias and polarization. *Journal of the European Economic Association 17*(5), 1470–1501.

Galperti, S. (2019). Persuasion: The art of changing worldviews. *American Economic Review 109*(3), 996–1031.

Gennaioli, N. and A. Shleifer (2010). What comes to mind. *The Quarterly journal of economics 125*(4), 1399–1433.

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

Jakobsen, A. (2020). An axiomatic model of persuasion. *Working Paper*.

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

Kahneman, D. and A. Tversky (1973). On the psychology of prediction. *Psychological review 80*(4), 237.

Kamenica, E. and M. Gentzkow (2011). Bayesian persuasion. *American Economic Review 101*(6), 2590–2615.

Kovach, M. (2020). Conservative updating. *Working Paper*.

Mullainathan, S. (2002). Thinking through categories. *Working Paper*.

Mullainathan, S., J. Schwartzstein, and A. Shleifer (2008). Coarse thinking and persuasion. *The Quarterly Journal of Economics 123*(2), 577–619.

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

Phillips, L. D. and W. Edwards (1966). Conservatism in a simple probability inference task. *Journal of experimental psychology 72*(3), 346.

Rabin, M. (1998). Psychology and economics. *Journal of economic literature 36*(1), 11–46.

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

Savage, L. J. (1972). *The foundations of statistics*. Courier Corporation.

Tversky, A. and D. Kahneman (1974). Judgment under uncertainty: Heuristics and biases. *science 185*(4157), 1124–1131.

Weinstein, J. (2017). Bayesian inference tempered by classical hypothesis testing. *Working paper*.

Wilson, A. (2014). Bounded memory and biases in information processing. *Econometrica 82*(6), 2257–2294.

Zhao, C. (2016). Pseudo-bayesian updating. *Working paper*.